A Lightweight Defeasible Description Logic in Depth

Quantification in Rational Reasoning and Beyond

MAXIMILIAN PENSEL

DISSERTATION

zur Erlangung des akademischen Grades
Doktoringenieur (Dr.-Ing.)

vorgelegt an der
Technischen Universität Dresden
Fakultät Informatik

GutachterIn:
PD. Dr.-Ing. Anni-Yasmin Turhan
Technische Universität Dresden

Prof. Dr. Piero Andrea Bonatti
Università degli Studi di Napoli Federico II

Oktober 2019

ABSTRACT

Description Logics (DLs) are increasingly successful knowledge representation formalisms, useful for any application requiring implicit derivation of knowledge from explicitly known facts. A prominent example domain benefiting from these formalisms since the 1990s is the biomedical field. This area contributes an intangible amount of facts and relations between low- and high-level concepts such as the constitution of cells or interactions between studied illnesses, their symptoms and remedies. DLs are well-suited for handling large formal knowledge repositories and computing inferable coherences throughout such data, relying on their well-founded first-order semantics. In particular, DLs of reduced expressivity have proven a tremendous worth for handling large ontologies due to their computational tractability. In spite of these assets and prevailing influence, classical DLs are not well-suited to adequately model some of the most intuitive forms of reasoning.

The capability for abductive reasoning is imperative for any field subjected to incomplete knowledge and the motivation to complete it with typical expectations. When such default expectations receive contradicting evidence, an abductive formalism is able to retract previously drawn, conflicting conclusions. Common examples often include human reasoning or a default characterisation of properties in biology, such as the *normal* arrangement of organs in the human body. Treatment of such defeasible knowledge must be aware of exceptional cases—such as a human suffering from the congenital condition *situs inversus*—and therefore accommodate for the ability to retract *defeasible* conclusions in a non-monotonic fashion. Specifically tailored non-monotonic semantics have been continuously investigated for DLs in the past 30 years. A particularly promising approach, is rooted in the research by Kraus, Lehmann and Magidor for preferential (propositional) logics and Rational Closure (RC). The biggest advantages of RC are its well-behaviour in terms of formal inference postulates and the efficient computation of defeasible entailments, by relying on a tractable reduction to classical reasoning in the underlying formalism. A major contribution of this work is a reorganisation of the core of this reasoning method, into an abstract framework formalisation. This framework is then easily instantiated to provide the reduction method for RC in DLs as well as more advanced closure operators, such as Relevant or Lexicographic Closure. In spite of their practical aptitude, we discovered that all reduction approaches fail to provide *any* defeasible conclusions for elements that only occur in the relational neighbourhood of the inspected elements. More explicitly, a distinguishing advantage of DLs over propositional logic is the capability to model binary relations and describe aspects of a related concept in terms of existential and universal quantification. Previous approaches to RC (and more advanced closures) are *not* able to derive typical behaviour for the concepts that occur within such quantification.

The main contribution of this work is to introduce stronger semantics for the lightweight DL $\mathcal{EL}_\bot$ with the capability to infer the expected entailments, while maintaining a close relation to the reduction method. We achieve this by introducing a new kind of first-order interpretation that allocates defeasible information on its elements directly. This allows to compare the level of typicality of such interpretations in terms of defeasible information satisfied at elements *in the relational neighbourhood*. A typicality preference relation then provides the means to single out those sets of models with *maximal typicality*. Based on this notion, we introduce two types of *nested rational semantics*, a sceptical and a selective variant, each capable of deriving the missing entailments under RC for arbitrarily nested quantified concepts. As a proof of versatility for our new semantics, we also show that the stronger Relevant Closure, can be imbued with typical information in the successors of binary relations. An extensive investigation into the computational complexity of our new semantics shows that the sceptical nested variant comes at considerable additional effort, while the selective semantics reside in the complexity of classical reasoning in the underlying DL, which remains tractable in our case.

ACKNOWLEDGEMENTS

After a long journey in the pursuit of knowledge and cognitive challenge I can honestly say that I have never faced a task more demanding than doing a PhD. Fortunately, I did not have to make this experience alone as I would not have gotten to this point without the people and relationships that surrounded me throughout this time.

I am grateful to the research training group QuantLA for three years of no-strings-attached funding. It allowed me the freedom to nurture my own interests, follow the research questions that I deemed important and investigate them at my own pace. Being encouraged to collaborate among peers has not only promoted my own research but has created some incredible friendships over the past years.

I thank Rafael Peñaloza for guiding me into the field of Description Logics and originally incepting me with the idea to do a PhD. Without your encouragement I would not have believed in myself enough to even consider this path.

To Franz Baader I am thankful for his professional support before and during my time with QuantLA. In spite of your busy schedule, your door was always open for technical and personal inquiries, with particular avail of the latter during the finalisation of my project.

My supervisor, Anni-Yasmin Turhan, has (often subconsciously) taught me the importance of making my thoughts accessible to the outside of my mind. You have shaped my ideas, presentation and work ethic, encouraging me to develop a superior form of professionalism which helped me to succeed in this project.

For the completion of this thesis I am grateful to the comments I have received from Gerd Brewka, David Hübner, Antoine Mottet and Pavlos Marantidis.

Not coming from a mathematical background, I am especially thankful for the continuous patient support by Antoine and Pavlos, lending me their expertise where mine was lacking. On a more personal note, your friendship over the past years was indispensable and I am not sure that I would have managed to hold on without your unrestricted sympathetic ears.

David, I bet that 25 years ago we did not imagine to be exchanging each other's PhD theses, let alone in the same year. Thank you for being a reliable and long-term, true friend who unknowingly has encouraged me to reach for the challenge and exceed my own expectations.

To my parents Ute and Thomas I am immeasurably grateful for everything you have done, every step of the way. Your support, understanding, pride and encouragement have allowed me to be where, and who I am today. To my other parents, Martina and Uwe I am grateful for your continued selfless support and appreciation, as well as your effortless acceptance of me as a part of this family.

Finally, my loving wife, you are the anchor that kept me sane and level for all these years. You have given my life a new sense of direction and made me into a stronger and better version of myself. I genuinely believe that together we can accomplish anything and overcome any obstacle. Just say the word, and we'll take on the world.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ALGORITHMS

LIST OF ABBREVIATIONS

1 INTRODUCTION

Any modern knowledge-based system necessitates to represent their data in a formal and machine-readable way. The field of *knowledge representation and reasoning* (KR) ([HLP'08]) studies logical foundations for applications where not all knowledge is explicitly known, but inferable from general rules. One of the most basic rules is logical implication, deriving truth of some state (consequent) from truth of another (antecedent). Given appropriate semantics, such rules are used to generate implicit conclusions (reasoning) from explicitly known facts (representation). Among the oldest, most well-known KR formalisms, count the propositional and *first-order logic* (FOL). While those classical foundations with their formal semantics are well-behaved and reliable in terms of obtainable consequences, this reliability quickly becomes a *liability* in non-static KR scenarios. A common example is the formalisation of human cognition. It is in our very nature to act with respect to incomplete knowledge and to develop or adjust our set of facts, or (more appropriately) beliefs over time. We *normally* drive on the right, *unless* we are in the UK (or others), we assume someone is busy if they are not answering their phone, a suspect is assumed innocent until proven guilty, and so forth. More meaningful perhaps are the many exceptional cases in natural sciences. The density of water behaves *unusual* below 4°C, elements further down in the periodic table have a higher electronegativity, *except* for those in group 11, etc.

The reliability of classical KR formalisms is due to the expectation of the development of knowledge, to *preserve truth*. At an early stage of an investigation, a suspect must be found innocent if they cannot be proven guilty. Classical formalisms forcibly preserve this conclusion in light of any new evidence. However, such preservation must be contested if an assumption, supplementing an incomplete set of facts, turns out to be false. The resulting *monotonic* behaviour of these formalisms makes them subject to the principle of explosion. Encountering mutually exclusive situations forces them to derive anything, rendering their supported inferences meaningless. When knowledge develops over time, Nute [Nut'01] explains, that it is not preservation of truth we expect, but preservation of *justification*.

> "[...] any reasoning system that preserves truth must be monotonic, [...] [but] a reasoning system that preserves *justification* will *not* be monotonic."
>
> *[Nut'01, p. 152]*

Extending knowledge, while preserving justification rather than truth, allows to falsify a previously drawn conclusion. The former inference is retracted in a *non-monotonic* fashion, but the description of the justification for deriving it, remains intact. If someone is suspected of a crime and all evidence is circumstantial, they must be found innocent. When substantial evidence is discovered, the innocence verdict must be retracted, without contesting the rule that circumstantial evidence alone is insufficient for a guilty verdict; perhaps this new evidence is later found inconclusive as well.

The paradigm of non-monotonic reasoning is intrinsically separate from the underlying KR formalism. Studies in this area, including this one, detach forms of knowledge representation from classical formalisms and merge them with a more powerful semantics, capable of modelling the expected behaviour.

1.1 NON-MONOTONIC REASONING

Classical Tarskian model-theoretic semantics ([TV'57]) are well-defined and widely accepted. If *all reasonable views* (models) on the world support a statement, it must be true. Monotonicity in such semantics is caused by the set of all models decreasing monotonically with increasing constraints (knowledge) on the world. Once truth of a statement is supported by a set of models, this truth is never retracted by considering subsets of those models. However, such retraction is very intuitive in the way we obtain our own conclusions on a daily basis. Among the most popular toy examples in this area is modelling the classification of penguins as birds, i.e. $\mathfrak{penguin} \to \mathfrak{bird}$. It is legitimate for any human (or otherwise) to assume birds to fly in general ($\mathfrak{bird} \to \mathfrak{fly}$). Only when learning that penguins, as specific birds, do not fly ($\mathfrak{penguin} \to \neg\mathfrak{fly}$), we end up with a set of mutually unsatisfiable characteristics. Unlike in classical semantics, if retraction of one of these statements is on the table, it is not immediately clear how a formalism should react. This predicament is even more evident from the famous Nixon diamond. Capturing that quakers are pacifists and republicans are not pacifists is problematic for Nixon, who is considered to be a quaker *and* a republican. According to these rules, Nixon is supposed to be a pacifist and not a pacifist at the same time. Where the penguin example might indicate an intuitive preference for more specific information to prevail (i.e. knowledge about penguins overriding general assumptions on birds), such specificity is not present in the Nixon diamond.

The study of *non-monotonic reasoning* (NMR) explores semantics capable of the expected consequence retraction. Ever since Sandewall [San'72] introduced a first kind of *default implication*, and Hewitt [Hew'72] studied a form of negation as failure to derive truth, the approaches to achieve non-monotonic semantics are plentiful and manifold. Sophisticated ideas such as McCarthy's circumscription [McC'80], explicitly modelling and minimising exceptional behaviour, and Reiter's default logic [Rei'80], furnishing the idea of default implication and completion of incomplete knowledge, are among the foundations for many modern reasoning systems.

Particularly promising, and continuously investigated since the early 1990s, is the preferential approach by *Kraus, Lehmann and Magidor* (KLM) [KLM'90]. They study the inferential capabilities of a weaker version of logical implication, e.g. $\mathtt{bird} \mathrel{|\!\!\approx} \mathtt{fly}$ (birds typically fly), allowing to explicitly model default behaviour. The appeal in their discussion on non-monotonic inference is the well-behaviour of a class of entailment relations that are characterised through a set of formal postulates. KLM carefully weaken the inferential properties of classical semantics, such as monotonicity (if α entails β, then so does $\alpha \wedge \gamma$), so as to capture entailment relations that are not necessarily monotonic. For an example, *rational monotonicity* allows to inherit the property β (implied by α) for the more specific sentence $\alpha \wedge \gamma$, if $\alpha \wedge \gamma$ does *not* imply the negation of β. These postulates can be understood as a guarantee to the behaviour of an entailment relation, capable of retraction. Consequently, the intuitively described preference of more specific attributes (e.g. penguins not flying) is inherent for what KLM call *rational entailment relations*.

In search of a practical answer to the intricate question of how to treat default knowledge, an extended study by Lehmann and Magidor [LM'92] introduces the rational extension of a body of default knowledge, as the *Rational Closure* (RC) of said knowledge. They claim that the sensible conclusions of a knowledge base are *at least* those contained in its Rational Closure. Enthusiasm for adaptations of RC to more sophisticated KR formalisms comes from a simple algorithmic characterisation of its entailments, in terms of classical semantics. In the past decade, this notion of rational consequence has been extensively studied, in particular, for a variety of modern KR formalisms known as Description Logics.

1.2 DESCRIPTION LOGICS

Description Logics (DLs) are KR formalisms whose strength lies in their formal semantics, computational properties and range of application opportunities. The first is essentially provided from first-order logic, the formalism that most DLs are fragments of. This fragmentation of FOL is motivated by the two other strengths, as essential goals in the study of DLs are to (1) provide a large range of expressive means to represent and reason over knowledge, and (2) fully map the borders of decidability and computational complexity in regards of such expressivity. Contributions, including this one, are often a combination of both. From an application's point of view, a variety of differently expressive and complex logics, each with a variety of supported inference mechanisms, provides a sizeable toolbox from which to select an appropriate formalism for almost any conceivable KR task.

Among the most notable applications of DLs is the *Web Ontology Language* (OWL) [HPH'03], a standard for formal representation of ontological knowledge in the semantic web, introduced by the World Wide Web Consortium. Advanced application areas such as biochemical engineering and medicine rely on OWL and DL ontologies to query vast amounts of data for logical consequences that no human could possibly be asked to derive [RBG+'97; Spa'00].

KR WITH DESCRIPTION LOGICS. The central notion in DLs is that of a *concept*. On a semantic level, concepts intuitively capture classes of objects by logical combinations of their primitive attributes. Such primitive properties include atomic characterisations of classes of objects (*concept names*), and binary relations between one object and another (*role names*). Primitive elements such as the set of all $\mathsf{Cats}/\mathsf{Dogs}$ or the directed relation characterising two elements as $\mathsf{friends}$ can then be combined with known logical connectives in the following manner. The concept

$$\mathsf{Cat} \sqcap \exists \mathsf{friend}.(\mathsf{Dog} \sqcup \mathsf{Cat}) \sqcap \neg \forall \mathsf{eats}.\mathsf{CatFood} \tag{1.1}$$

describes the set of all cats that befriend either a cat or (disjunction) a dog and (conjunction) *not* only eat cat food. Such concept constructors are evaluated in terms of sets, where for instance the elements belonging to $\mathsf{Cat} \sqcap \mathsf{Smart}$ are those that belong to the intersection of all Cat and all Smart elements. Note that DLs use a specific variant of quantification, capturing classes for elements in the relational neighbourhood, along the binary relations that are described by role names (friend, eats). More formally, in terms of first-order quantification, DLs quantify over the *range* of the binary relations that are represented as role names. An element d belongs to $\exists \mathsf{friend}.\mathsf{Cat}$, if there is at least some element e, related to d via the role friend, such that e belongs to Cat.

From concepts alone, one can derive simple consequences. For example, the cats described in (1.1) have at least some friend ($\exists \mathsf{friend}.\top$), or they eat *something* besides cat food ($\exists \mathsf{eats}.(\neg \mathsf{CatFood})$). A concept can also be analysed for *consistency*, i.e. checking whether any element can possibly satisfy this combination of properties. A trivial example for an inconsistent concept would be the unfortunate modelling of humans that are behaving inhumane with $\mathsf{Human} \sqcap \neg \mathsf{Human}$. When quantification is involved, such inconsistencies are not so obvious.

More complex forms of knowledge representation are provided by describing universal relations between concepts. It could be captured that all cats befriend only cats ($\mathsf{Cat} \sqsubseteq \forall \mathsf{friend}.\mathsf{Cat}$), or that, if a cat befriends a dog, it cannot be very smart ($\mathsf{Cat} \sqcap \exists \mathsf{friend}.\mathsf{Dog} \sqsubseteq \neg \mathsf{Smart}$), etc. Such general statements about classes of objects are called *terminological axioms*. They resemble classical implication, which, on a set-level, corresponds to inclusion. The consequences that can be derived from such terminological knowledge are also of terminological nature. The corresponding reasoning service is called *subsumption*. Formally, one asks if (new) *concept inclusions* can be implicitly derived from the explicitly modelled terminological information.

In addition to the (inclusion-)hierarchical organisation of entire classes of objects, DLs allow to capture explicitly named individuals and place them within this hierarchical knowledge on concepts. This placement is formally achieved with assertional axioms, for instance, capturing explicit individuals such as daisy, to belong to the class of all cats, or befriend another explicitly named individual molli:

$\mathit{Cat}(\mathit{daisy})$, or $\mathit{friend}(\mathit{daisy}, \mathit{molli})$

In light of such assertive axioms, implicit class membership of individuals can be concluded, respecting the terminological knowledge as well. For instance, from the concept inclusions above, we would conclude daisy to befriend only cats, and thus derive for molli to be a cat, i.e.

$(\forall \mathit{friend}.\mathit{Cat})(\mathit{daisy})$, and $\mathit{Cat}(\mathit{molli})$.

This reasoning service is called *instance checking*. Other questions about assertional knowledge are asking for the *set* of all individuals belonging to a certain concept (instance retrieval), or even formulating a complex query as conjuncts of several retrieval questions (conjunctive query answering), just to name a few. However, our work is exploring means of reasoning that are orthogonal to the types of queries asked. Therefore, we are only concerned with the very basic services of subsumption and instance checking, implicitly relying on an analysis of consistency w.r.t. terminological and assertional knowledge. This form of consistency check is more meaningful than its concept-only variant. Considering the example from Section 1.1, modelling birds to fly ($\mathit{Bird} \sqsubseteq \mathit{Fly}$), and penguins, as specific birds ($\mathit{Penguin} \sqsubseteq \mathit{Bird}$), not to fly ($\mathit{Penguin} \sqsubseteq \neg\mathit{Fly}$), leads to the conclusion that $\mathit{Penguins}$ are inconsistent with this knowledge, while birds are not.

STUDYING DESCRIPTION LOGICS. As we have just shown, Description Logics are composed of many separate aspects. Investigating different compositions of concept constructors, knowledge representation axioms, reasoning services or even non-standard semantics (e.g. multi-valued, fuzzy, non-monotonic), is the heart of Description Logics research. The most fundamental variations, resulting in the characterisation of different formalisms altogether, are considered to be on the concept constructor level. A handy nomenclature is used to distinguish DLs in terms of their allowed constructors. The boolean spectrum together with the DL quantification that we have used above, constitutes the *attribute language with complement*, $\mathcal{ALC}$. The inclusions of other constructors or even axioms are usually described by such a sequence of labels. For instance, allowing to express *sub-role* relationships and essentially organise role names in a *hierarchical* fashion, is signified with the additional letter $\mathcal{H}$. In fact, the full OWL 2 standard is actually based on the DL $\mathcal{SROIQ}$, a powerful extension of $\mathcal{ALC}$.[1] At the same time, even the basic services subsumption and instance checking are already intractable in $\mathcal{ALC}$ (when including the types of axioms from above). Motivated by the goal of tractable reasoning, it is also very appealing to move below the expressivity of $\mathcal{ALC}$.

1 The letter $\mathcal{S}$ is an abbreviation for $\mathcal{ALC}$ with transitive roles.

Renouncing universal quantification, disjunction and negation yields the existential language $\mathcal{EL}$. $\mathcal{EL}$ and sub-boolean extensions (or fragments) of it are considered prominent, *lightweight* DLs with tractable reasoning capabilities. While computational superiority is gained at the cost of expressivity, not all application areas rely on the full boolean spectrum per se. This is endorsed by large biomedical ontologies such as SNOMED [Spa'00] considering such lightweight DLs as sufficient modelling languages.

LIGHTWEIGHT DLS. In the existential language $\mathcal{EL}$, only conjunction and existential restriction are allowed as concept constructors. Intuitively, the possible concepts in $\mathcal{EL}$ are of a constructive nature. Rather than capturing restrictions of properties, existence of related elements and collections of properties can be modelled, e.g.

$$\mathit{Cat} \sqcap \exists \mathit{friend}.(\mathit{Dog} \sqcap \mathit{Cute}) \sqcap \exists \mathit{eats}.\mathit{CatFood}\,.$$

Terminological axioms are positive in the sense of Horn-logic. Satisfaction of some properties can only lead to the entailment of *more* (positive) properties. Consequently, tractable reasoning algorithms employ the technique of knowledge completion, i.e. the consequences of every concept/individual can be derived deterministically and iteratively.

As expected, some reasoning services become trivial when certain features are not expressible in a DL. For example, consistency checking without any support for negation is insignificant. Essentially, any negation-free concept can always be satisfied, regardless of the (negation-free) background terminological or assertional knowledge. However, there are means of modelling negative constraints without introducing full negation, and therefore remaining below $\mathcal{ALC}$. A small extension of $\mathcal{EL}$ is the DL $\mathcal{EL}_\bot$, allowing to use $\bot$ as a primitive and always *unsatisfiable* concept. Formally, in terms of sets, $\bot$ is understood as the empty set. Due to the constructive nature of conjunction and existential restriction, any $\mathcal{EL}$ concept containing $\bot$ is trivially unsatisfiable on its own. Nevertheless, utilising $\bot$ in terminological axioms provides the means to express simple forms of negation. For example, the bird-penguin predicament could be expressed entirely in $\mathcal{EL}_\bot$, by introducing a concept name Walk, in $\mathit{Penguin} \sqsubseteq \mathit{Walk}$. Then, the terminological *disjointness*

$$\mathit{Fly} \sqcap \mathit{Walk} \sqsubseteq \bot,$$

is used to ensure that no element flies and walks at the same time.

1.3 RATIONAL CLOSURE IN DESCRIPTION LOGICS

Exploring NMR for DLs is very appealing, due to the intuitive capabilities of non-monotonicity and the computational properties and practical applicability of DLs. As a matter of fact, the biomedical domain, already employing DLs as a modelling formalism, could benefit greatly from the means to model non-monotonic behaviour, because exceptions in biology are everywhere.

Aside from the classification of penguins as birds, the condition of situs inversus—where the organs of the affected human appear in a mirrored arrangement to the *normal* case—or, all vertebrates carrying red blood cells—*except* crocodile icefish (Channichthyidae)—are more meaningful examples; the list is inexhaustible. Unfortunately, non-monotonic reasoning in DLs is not established enough to have brought forth a practical ontology employing their capabilities.

Many NMR approaches from the past century have been applied to differently expressive DLs, including the practical algorithm to capture entailments under Rational Closure. Default terminological knowledge in DLs is called *defeasible inclusion*, and as for KLM, it weakens the notion of terminological implication. Defeasible inclusions allow to capture normality statements such as $\mathsf{Bird} \mathrel{\underset{\sim}{\sqsubset}} \mathsf{Fly}$ or $\mathsf{Penguin} \mathrel{\underset{\sim}{\sqsubset}} \neg\mathsf{Fly}$, reading for instance "birds usually fly".

The algorithm producing defeasible entailments based on such knowledge, employs a reduction technique. Entailment of a defeasible consequence such as $\mathsf{Penguin} \mathrel{\underset{\sim}{\sqsubset}} \mathsf{Feathered}$, is determined by enriching the query concept ($\mathsf{Penguin}$) with defeasible inclusions in the form of *material-implications*. As in propositional logic, the material implication $\mathsf{Bird} \rightarrow \mathsf{Feathered}$ is equivalent to the concept $\neg\mathsf{Bird} \sqcup \mathsf{Feathered}$. The defeasible query can then be transformed into a classical query, using such concepts to restrict the queried elements to satisfy certain defeasible statements. Specifically, the defeasible information $\mathsf{Bird} \mathrel{\underset{\sim}{\sqsubset}} \mathsf{Feathered}$ can be appended as a material-implication conjunct to the left-hand side of a query, as

$$(\neg\mathsf{Bird} \sqcup \mathsf{Feathered}) \sqcap \mathsf{Penguin} \sqsubseteq \mathsf{Feathered}\,.$$

In this particular case, knowing that all $\mathsf{Penguins}$ are Birds, this classical subsumption would be concluded from the strict terminological knowledge. To obtain meaningful conclusions, defeasible information is only *materialised* into the query, if it remains consistent together with the query concept. For example

$$(\neg\mathsf{Bird} \sqcup \mathsf{Fly}) \sqcap (\neg\mathsf{Penguin} \sqcup \neg\mathsf{Fly}) \sqcap \mathsf{Penguin}$$

is inconsistent (empty set) w.r.t. the strict terminological knowledge. By the principle of explosion, its consequences become worthless (technically, the empty set is included in *every* set). Hence, a set of defeasible statements that are consistent with the query must be determined prior to materialisation. This determination implicitly models the knowledge retraction demanded in Section 1.1, as $\mathsf{Bird} \mathrel{\underset{\sim}{\sqsubset}} \mathsf{Fly}$ and $\mathsf{Penguin} \mathrel{\underset{\sim}{\sqsubset}} \neg\mathsf{Fly}$ are consistent with Bird, but not with the more specific bird, $\mathsf{Penguin}$. The difficulty in this procedure is to determine such consistent sets. Different methods have been studied [CMMN'14; CS'10; CS'11; CS'12], leading to differently *strong* closure operators for defeasible knowledge. The study of stronger closures is motivated to alleviate the drawback of RC by KLM known as *inheritance blocking*. Intuitively, once a class is identified as exceptional (e.g. $\mathsf{Penguin}$ is exceptional w.r.t. Bird), it is blocked from inheriting *any* defeasible information about its more general superclass under RC.

Notwithstanding the above, all closures that employ materialisation fail spectacularly when considering DL quantification. It is very easily illustrated that the materialisation of $\mathtt{Bird} \mathrel{\sqsubseteq\!\!\!\sim} \mathtt{Fly}$ has no effect on concepts in the relational neighbourhood:

$$(\neg\mathtt{Bird} \sqcup \mathtt{Fly}) \sqcap \exists\mathtt{friend}.\mathtt{Bird}$$

describes the set of all elements that befriend some bird, and *themselves* are either not birds or capable of flying. No restrictions are imposed on the class of birds related via the role $\mathtt{friend}$. Quantification is effectively *neglected* in forms of defeasible reasoning that are based on materialisation. Surprisingly, this issue has not been addressed in terms of rational closure until our advancements in 2017 ([PT'17a; PT'17b]).

Resolving this fatal drawback of materialisation-based rational (and stronger) entailment is the main contribution of this thesis. Materialisation-based entailments can be viewed as *propositional* in nature, because defeasible information is only applied to the top level of query concepts. We formally introduce a more meaningful, *nested coverage* of default information and redefine the reasoning service of defeasible subsumption and instance checking for $\mathcal{EL}_\bot$. The full spectrum of the semantics we investigate with a model-theoretic characterisation is captured by two parameters.

STRENGTH: Initially, we consider the basic determination of consistent sets of statements, resulting in *rational strength*. To resolve the issues of inheritance blocking and quantification neglect at the same time, we also study consequences of *relevant strength*.

COVERAGE: As a result of aligning our semantics closely with previous approaches to capture Rational and Relevant Closure, we are able to recreate the original *propositional coverage*, in addition to two different types of *nested coverage*. *Sceptical nested coverage* has debuted in [PT'17a], while the *selective nested coverage* is introduced here, providing even more entailments than the sceptical variant.

We formally show equivalence of reasoning with propositional coverage and the materialisation-based reduction, as well as superiority of nested reasoning over propositional coverage (and thus, materialisation).

1.4 STRUCTURE OF THE THESIS

This thesis is divided into three major parts. As the second and third part present the majority of our contributions, they are each preceded by their own introductions and structural overviews.

PART I. The research question we address is composed of two aspects. The foundations and preliminary notation of Description Logics, in particular $\mathcal{ALC}$ and $\mathcal{EL}_\bot$, are given in Chapter 2. Chapter 3 provides a broader introduction into the area of non-monotonic reasoning, briefly covering related approaches and their combinations with DLs, with a more technical focus on reasoning under Rational Closure. PART I is purposefully not denoted as "preliminaries", because in addition to providing basic notations, several original contributions, as results for classical DLs, are presented in Section 2.4. These results provide advanced foundations, but foundations nonetheless, for PART II and III.

PART II. We reorganise previous materialisation-based approaches to RC in DLs and unify them in an abstract framework in Chapter 4. The versatility of this framework for materialisation-based defeasible entailment is exemplified in Chapter 6 with instantiations for the more expressive Relevant and Lexicographic Closure. The drawbacks and insufficiencies of the materialisation framework, motivating the final part of this thesis, are discussed thoroughly in Chapter 5.

PART III. Finally, our main contribution in the form of stronger, model-theoretic semantics, is covered in full technical detail in PART III. Our new model formalism is introduced and analysed in terms of expressivity in Chapter 7. This analysis is complemented with an in-depth investigation of this new formalism's computational properties in Chapter 8, including a full algorithmic characterisation in Section 8.1. We conclude with a brief discussion on open problems and potential directions for future investigations, alongside an overview of all contributions of this thesis in Chapter 9.

Part I

FOUNDATIONS

2 DESCRIPTION LOGICS

In this chapter we introduce the required foundations of Description Logics, on which the remainder of this work is built. We begin by formally presenting the basic notions of DLs in general (Section 2.1). These are exemplified with the two prominent members $\mathcal{ALC}$ and $\mathcal{EL}_\bot$, which are the bedrock of our investigations. Building upon these basics, we continue to introduce the standard notions for representation of, and reasoning over knowledge expressed in these DLs (Section 2.2). A small outlook, fitting the non-standard notion of defeasible reasoning into the context of classical DLs, is enclosed in Sec. 2.2. The final additions to this chapter are a brief exposition on fundamental results (Sec. 2.3), and an investigation of more advanced notions that are required in later parts of this work (Sec. 2.4). The latter are somewhat unconventional in the broad area of DLs. They are tailored specifically to our studies, and as such, contain original thoughts and contribute preliminary results.

For a full introduction to DLs, consult Baader et al. [BHLS'17] or [BCM+'10]. For more detailed fundamental investigations of $\mathcal{ALC}$ and $\mathcal{EL}_\bot$ in particular, consider Schmidt-Schauß and Smolka [SS'91] and Baader et al. [BBL'05; BLB'08], respectively.

2.1 SYNTAX AND SEMANTICS

The central syntactic construct in DLs is a concept. It is inductively defined in terms of *concept constructors*, starting from the primitive building blocks that are *concept names* and *role names*. Formally concept and role names are taken from two disjoint sets $\mathsf{N_C}$ and $\mathsf{N_R}$, respectively. By convention, concept names are denoted either with upper case descriptive words, e.g. $\mathit{Cat}, \mathit{Dog}$ or in abstract examples with upper case letters near the beginning of the alphabet (A, B). Dually, role names are denoted with either lower case descriptive words $\mathit{friend}, \mathit{eats}$ or lower case letters, typically starting from and succeeding r. The most basic connectors to allow for more complex constructions of concepts, are the boolean conjunction $\sqcap$, disjunction $\sqcup$ and negation $\neg$, whose meaning is aligned with that of first-order logic. Intuitively, concepts capture classes of objects and role names describe binary relations between objects. Using roles, quantified concepts capture sets of elements that are related to some (existential) or only (universal) elements of the *quantified* concept. In particular, the concept $\exists r.C$ (*existential restriction*) gathers those objects with at least some binary relation of type r to an element fitting the description C. Conversely, the concept $\forall r.C$ (*value restriction*) collects those objects for which *all* successors through the relation r fit the description C. Existential and value restriction are

Constructor	Syntax	Semantics
Conjunction	$C \sqcap D$	$C^{\mathcal{I}} \cap D^{\mathcal{I}}$
Disjunction	$C \sqcup D$	$C^{\mathcal{I}} \cup D^{\mathcal{I}}$
Negation	$\neg C$	$\Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$
Existential Restriction	$\exists r.C$	$\{d \in \Delta^{\mathcal{I}} \mid \exists e \in \Delta^{\mathcal{I}}.(d, e) \in r^{\mathcal{I}} \wedge e \in C^{\mathcal{I}}\}$
Value Restriction	$\forall r.C$	$\{d \in \Delta^{\mathcal{I}} \mid \forall e \in \Delta^{\mathcal{I}}.(d, e) \in r^{\mathcal{I}} \implies e \in C^{\mathcal{I}}\}$

Table 2.1: Basic DL concept constructors.

what move DLs truly beyond the expressive power of propositional logic. It is worth noting that this form of quantification can be expressed by first-order formulas [Bor'96]. However, their syntactical variant in DLs is not to be mistaken for a formula quantifying over a variable r, but rather over elements in the *range* of the relation r. The border cases for the objects of concern, namely *every object* and *no object*, can be captured uniformly using the special primitive concepts top $\top$ and bottom $\bot$. Intuitive uses of $\top$ or $\bot$ on a concept level are for instance $\forall r.\bot$, describing the class of all objects without any r-successors, or $\exists r.\top$, the class of all objects with some r-successor.

Different members of the DL-family are typically distinguished by the set of concept constructors that are allowed to express their concepts. This is also a fundamental distinguishing aspect, in terms of computability and expressivity of consequences in the resulting logic. When results or definitions are independent of the underlying DL, we use the generic $\mathcal{L}$ to denote *any* DL. The two members that are of interest in this work are $\mathcal{ALC}$, using the full spectrum of the above constructors, and its fragment $\mathcal{EL}_{\bot}$.

Definition 2.1 (Syntax of $\mathcal{ALC}$ and $\mathcal{EL}_{\bot}$). For a DL $\mathcal{L}$, the set of all concepts in $\mathcal{L}$ is denoted as $\mathfrak{C}(\mathcal{L})$. Concepts in $\mathfrak{C}(\mathcal{L})$ are defined in the following inductive fashion. For any DL, $N_C \subseteq \mathfrak{C}(\mathcal{L})$. For $\mathcal{ALC}$, if $r \in N_R$, and $C, D \in \mathfrak{C}(\mathcal{ALC})$ then

$$C \sqcap D, C \sqcup D, \neg C, \exists r.C, \forall r.C, \top, \bot \in \mathfrak{C}(\mathcal{ALC}).$$

For the lightweight DL $\mathcal{EL}_{\bot}$, if $r \in N_R$, and $C, D \in \mathfrak{C}(\mathcal{EL}_{\bot})$ then

$$C \sqcap D, \exists r.C, \top, \bot \in \mathfrak{C}(\mathcal{EL}_{\bot}).$$

If nothing is known or expected of given concepts in $\mathfrak{C}(\mathcal{L})$, it is the convention to denote them with upper case letters starting from and succeeding C, even when they appear nested in another concept. For example, we say C is quantified in the concept $D \sqcap \exists r.C$, where neither C nor D are expected to be primitive concept names.

The meaning of concepts that was described intuitively as the description of classes of objects, is formally captured by DL semantics. Because DLs

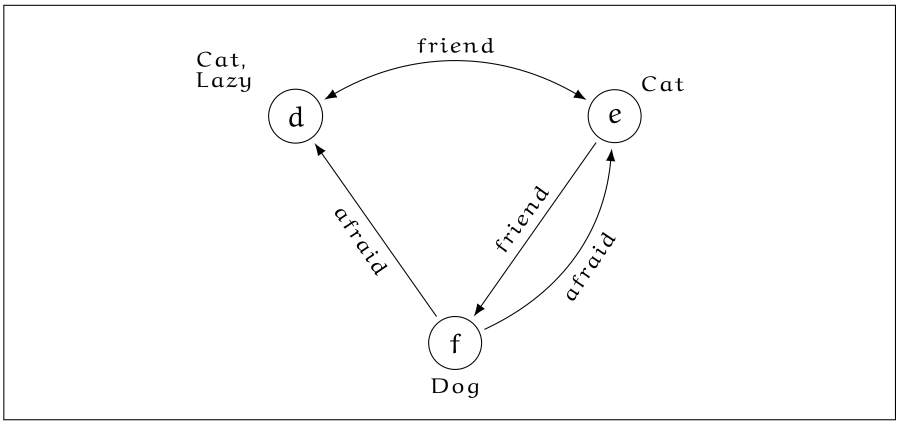

Figure 2.1: Graph visualisation of $\mathcal{I}$ in Example 2.3.

are designed as fragments of first-order logic, their interpretation semantics are akin.

Definition 2.2 (Semantics of $\mathcal{ALC}$). An *interpretation* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ is comprised of a non-empty interpretation *domain* (short: domain) $\Delta^{\mathcal{I}}$ and an *extension mapping* $\cdot^{\mathcal{I}}$. Concept names $A \in N_C$ are mapped to subsets of the domain under the extension mapping, i.e. $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$, while role names $r \in N_R$ are mapped to binary relations over the domain, i.e. $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. The special primitive concepts $\top$ and $\bot$ are mapped to $\top^{\mathcal{I}} = \Delta^{\mathcal{I}}$ and $\bot^{\mathcal{I}} = \emptyset$ under any interpretation $\mathcal{I}$. The mapping $\cdot^{\mathcal{I}}$ is inductively extended to non-primitive concepts as in Table 2.1.

Given the notion of interpretations, it becomes apparent that every concept (whether primitive or not) can be seen as a set of objects, more explicitly, as a set of *domain elements*. We frequently refer to these sets, e.g. for a concept $C^{\mathcal{I}}$, as *the extension of* C *under* $\mathcal{I}$. By convention, domain elements are often denoted with lower case letters starting from, and succeeding d. Note that in general, there are no restrictions on what *kind* of elements belong to a domain. As a matter of fact, for reasoning in $\mathcal{EL}_\bot$, it is common practice to let $\Delta^{\mathcal{I}} \subseteq \mathfrak{C}(\mathcal{EL}_\bot)$ (cf. Sec. 2.4).

Interpretations can be described directly as a collection of subsets of the domain (one for each concept name) and binary relations (one for each role name). It is common practice to visualise them as directed graphs, labelling vertices (domain elements) with sets of concept names and edges with sets of role names. Sometimes we require to treat specific *role edges* within an interpretation. For that, we use the construct $r(d, e)$ for $r \in N_R$ and $d, e \in \Delta^{\mathcal{I}}$, calling d the *predecessor* and e the *successor*.

Example 2.3. Consider our simple and comprehensible application context about cats and dogs. Let $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ with $\Delta^{\mathcal{I}} = \{d, e, f\}$. Suppose the concept names Cat, Dog, $\mathsf{Lazy} \in N_C$, and the role names friend, $\mathsf{afraid} \in N_R$, are assigned to sets and relations in $\mathcal{I}$ as follows.

$\mathsf{Cat}^\mathcal{I} =$	$\{d, e\}$	$\mathsf{friend}^\mathcal{I} =$	$\{(d, e), (e, d), (e, f)\}$
$\mathsf{Dog}^\mathcal{I} =$	$\{f\}$	$\mathsf{afraid}^\mathcal{I} =$	$\{(f, e), (f, d)\}$
$\mathsf{Lazy}^\mathcal{I} =$	$\{d\}$		

The interpretation $\mathcal{I}$ is visualised, and much more easily grasped, in Figure 2.1. From $\mathcal{I}$, we can derive more complex extensions, such as the set of all lazy cats $(\mathsf{Lazy} \sqcap \mathsf{Cat})^\mathcal{I} = \{d\}$ or those elements having at least some friend that is a Dog, $(\exists \mathsf{friend}.\mathsf{Dog})^\mathcal{I} = \{e\}$. Such consequences are also more easily understood from visualisations such as Fig. 2.1. Thus, in many examples, we will present the labelled graph representation of an interpretation, rather than its set-theoretic definition.

Visual representations of interpretations, will often use $\bullet$ for domain elements, if their name is irrelevant or clear from the context. The interpretation $\mathcal{I}$ in Example 2.3 is one potential view on relations and aspects of this particular application context. Infinitely many others, including ones with infinite domains exist alongside $\mathcal{I}$. The essence of representing knowledge and reasoning over it in Description Logics, is to *explicitly* define sets of axioms (representation) that restrict the set of all interpretations to those in compliance with the given axioms. *Implicit* knowledge is then derived from commonalities that are shared among those interpretations (reasoning).

2.2 KNOWLEDGE REPRESENTATION AND REASONING IN DLS

While investigating the properties of a concept in terms of interpretation semantics can be seen as a representation and reasoning scenario already, it is a rather inexpressive one. Formulating, for example, relations between concepts that are expected to be satisfied in general, is a more meaningful type of represented knowledge. One could generalise the fact that every Cat is Smart, or that everything whose friend of a friend is a Dog should be friendly to some Dog directly. A constraint could also be to require the disjointness of two concepts such as Dog and Cat. We separate the syntactic aspect of such knowledge representation from the semantic repercussions they have on interpretations, before discussing the actual reasoning services that are supported by this syntax and semantics.

SYNTAX OF DL KNOWLEDGE REPRESENTATION. The examples above are of terminological nature, that is, describing relations between entire classes of objects, similar to logical implication. Other types of axioms are of assertional nature, allowing to persist named elements, such as daisy or molli, throughout all interpretations, and *asserting* properties and relations onto it. This is achieved by introducing explicit names, called *individuals*, taken from the set N_I, which is disjoint from N_C and N_R. In abstract scenarios, the convention is to use lower case letters $a, b \in N_I$. In the presence of individuals, every interpretation $\mathcal{I}$ is required to assign a domain element to these individuals, i.e. $a^\mathcal{I} \in \Delta^\mathcal{I}$. Continuing on Example 2.3, $\mathcal{I}$ could assign for $\mathsf{molli}, \mathsf{daisy} \in N_I$, $\mathsf{daisy}^\mathcal{I} = d$ and

$\mathsf{molli}^{\mathcal{I}} = e$. In such cases, we usually visualise the domain elements d, e with the individual name they represent. If $\mathcal{I}$ is extended in precisely this way, then f would be considered an *anonymous* element of the concept Dog, in contrast to the *named* Cat elements daisy and molli.

There are two types of assertive axioms. One imposes a concept C onto an individual a, intuitively expecting interpretations to assign a to a domain element in the extension of C. The other is explicitly forcing named relations (role edges) between two individuals.

Definition 2.4 (Terminological and Assertional Axioms)**.** Let $a, b \in N_I$ be individuals, $C, D \in \mathfrak{C}(\mathcal{L})$ concepts and $r \in N_R$ a role name.

- A *general concept inclusion* (GCI) is an axiom $C \sqsubseteq D$.
- A *concept assertion* is an axiom $C(a)$.
- A *role assertion* is an axiom $r(a, b)$.

A finite set of GCIs is called a *TBox* $\mathcal{T}$, and a finite set of concept and role assertions is called an *ABox* $\mathcal{A}$.

An ABox and a TBox are the usual constituents of a *knowledge base* (KB), signifying the representation of knowledge. Formally, a *knowledge base* $\mathcal{K} = (\mathcal{A}, \mathcal{T})$ is represented as a tuple. The following methods for discussing the syntax of KBs (or concepts/axioms) are heavily used in the remainder of this work.

Definition 2.5. For a KB $\mathcal{K} = (\mathcal{A}, \mathcal{T})$, the set of *subconcepts* $\mathsf{sub}(\mathcal{K})$ in $\mathcal{K}$ is inductively defined, starting with $C, D \in \mathsf{sub}(\mathcal{K})$ for $C \sqsubseteq D \in \mathcal{T}$ or $C(a) \in \mathcal{A}$. For the constructors in Table 2.1, $C \sqcap D, C \sqcup D \in \mathsf{sub}(\mathcal{K})$ implies $C, D \in \mathsf{sub}(\mathcal{K})$ and $\neg C, \exists r.C, \forall r.C \in \mathsf{sub}(\mathcal{K})$ implies $C \in \mathsf{sub}(\mathcal{K})$. The set of *quantified concepts* $\mathsf{Qc}(\mathcal{K})$ in $\mathcal{K}$ is

$$\mathsf{Qc}(\mathcal{K}) = \{C \in \mathsf{sub}(\mathcal{K}) \mid \exists r.C \in \mathsf{sub}(\mathcal{K}) \text{ or } \forall r.C \in \mathsf{sub}(\mathcal{K})\}.$$

The *signature* of $\mathcal{K}$ is $\mathsf{sig}(\mathcal{K}) = \mathsf{sig}_C(\mathcal{K}) \uplus \mathsf{sig}_R(\mathcal{K}) \uplus \mathsf{sig}_I(\mathcal{A})$, with

$$\begin{aligned}
\mathsf{sig}_C(\mathcal{K}) &= \mathsf{sub}(\mathcal{K}) \cap N_C, \\
\mathsf{sig}_R(\mathcal{K}) &= \{r \in N_R \mid \exists r.C \in \mathsf{sub}(\mathcal{K}) \text{ or } \forall r.C \in \mathsf{sub}(\mathcal{K}) \\
&\qquad \text{or } r(a, b) \in \mathcal{A}\}, \text{ and} \\
\mathsf{sig}_I(\mathcal{A}) &= \{a, b \in N_I \mid C(a) \in \mathcal{A} \text{ or } r(a, b) \in \mathcal{A}\}.
\end{aligned}$$

Remark 2.6. All of the functions in Definition 2.5 extend to axioms (e.g. $\mathsf{sig}_C(C \sqsubseteq D)$) and concepts (e.g. $\mathsf{Qc}(C)$) in the natural way. On occasion, we will also apply these notions to arbitrary tuples, e.g. $\mathsf{sub}(\mathcal{K}, C, D, \cdots)$.

Looping back to the intuitively expressed GCIs from the beginning of this section, their formal syntactic representation would be $\mathsf{Cat} \sqsubseteq \mathsf{Smart}$ (all cats are smart), $\exists\mathsf{friend}.(\exists\mathsf{friend}.\mathsf{Dog}) \sqsubseteq \exists\mathsf{friend}.\mathsf{Dog}$ (befriending a dog by proxy requires to befriend a dog directly), and $\mathsf{Cat} \sqcap \mathsf{Dog} \sqsubseteq \bot$ (cats cannot be dogs and vice versa). We continue to supply meaning to such axioms in terms of interpretation semantics.

	Syntax	Semantics	
Concept	C	$\mathcal{I} \models C$	*iff* $C^{\mathcal{I}} \neq \emptyset$
Role edge	$r(d, e)$	$\mathcal{I} \models r(d, e)$	*iff* $(d, e) \in r^{\mathcal{I}}$
GCI	$C \sqsubseteq D$	$\mathcal{I} \models C \sqsubseteq D$	*iff* $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$
Concept assertion	$C(a)$	$\mathcal{I} \models C(a)$	*iff* $a^{\mathcal{I}} \in C^{\mathcal{I}}$
Role assertion	$r(a, b)$	$\mathcal{I} \models r(a, b)$	*iff* $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in r^{\mathcal{I}}$

Table 2.2: Satisfaction of DL axioms, concepts and role edges under an interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, for $C, D \in \mathfrak{C}(\mathcal{L})$, $d, e \in \Delta^{\mathcal{I}}$, $r \in N_R$ and $a, b \in N_I$.

SEMANTICS OF DL KNOWLEDGE REPRESENTATION. So far we presented only the syntactical constructs for representing knowledge in DLs, accompanied with some *intuition* on their semantics. Formally, the axioms in a KB characterise a set of interpretations that are called *models* of the KB. As in many logics, the notion of a model is tightly linked to the notion of *satisfaction*. The satisfaction of axioms and other DL-structures, is defined as follows.

Definition 2.7 (Satisfaction in DLs)**.** An interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ *satisfies* a concept, role edge, GCI, concept assertion or role assertion according to the condition in the third column of Table 2.2. An interpretation satisfies a TBox ($\mathcal{I} \models \mathcal{T}$), ABox ($\mathcal{I} \models \mathcal{A}$) or KB ($\mathcal{I} \models \mathcal{K}$) *iff* it satisfies all axioms they contain. The set of all models of a KB $\mathcal{K} = (\mathcal{A}, \mathcal{T})$ is

$$\texttt{Mod}(\mathcal{K}) = \{\mathcal{I} \mid \mathcal{I} \models \mathcal{A} \text{ and } \mathcal{I} \models \mathcal{T}\}.$$

Many times it is also beneficial to single out a particular domain element and discuss properties that it satisfies specifically. This could be in relation to a concept, saying e satisfies C in $\mathcal{I}$ ($e \in C^{\mathcal{I}}$), but also a GCI $C \sqsubseteq D$, i.e. $e \in C^{\mathcal{I}} \implies e \in D^{\mathcal{I}}$. We distinguish between an element $e \in \Delta^{\mathcal{I}}$ *actively* satisfying a GCI with $e \in C^{\mathcal{I}} \cap D^{\mathcal{I}}$, or *passively* satisfying it, when e does not satisfy C in the first place.

Recall the interpretation $\mathcal{I}$ from Example 2.3. The axiom $\mathsf{Cat} \sqsubseteq \mathsf{Smart}$ is clearly not satisfied by $\mathcal{I}$, because neither d nor e satisfy the concept Smart. For $\mathsf{Cat} \sqcap \mathsf{Dog} \sqsubseteq \bot$ to be satisfied, the intersection of the extensions of Cat and Dog needs to be empty (see Table 2.1 and 2.2), which is the case in $\mathcal{I}$. Less intuitively, e does satisfy $\exists \mathsf{friend}.(\exists \mathsf{friend}.\mathsf{Dog}) \sqsubseteq \exists \mathsf{friend}.\mathsf{Dog}$ in $\mathcal{I}$, not because it belongs to the extension of the right-hand side, but because it does *not* belong to the left-hand side, and thus passively satisfies it.

CLASSICAL REASONING IN DESCRIPTION LOGICS. One of the elementary reasoning tasks that is typically considered for DLs, is consistency checking. Intuitively, this means asking the question "Does the knowledge we represented even allow satisfaction in any way?". This notion can be extended to check consistency of constructs, such as concepts or GCIs, *with respect to* a KB.

Definition 2.8 (Consistency of and with Knowledge Bases). A KB $\mathcal{K} = (\mathcal{A}, \mathcal{T})$, is *consistent iff* $\mathtt{Mod}(\mathcal{K}) \neq \emptyset$. For a concept $C \in \mathfrak{C}(\mathcal{L})$, C is consistent with the KB $\mathcal{K}$ *iff* there is some $\mathcal{I} \in \mathtt{Mod}(\mathcal{K})$ such that $\mathcal{I} \models C$. A KB is QC-consistent *iff* all $E \in \mathrm{Qc}(\mathcal{K})$ are consistent with $\mathcal{K}$.

Consistency, in particular QC-consistency, can also be seen as a property of a KB, rather than an informative reasoning task. QC-consistency is most useful for $\mathcal{EL}_\perp$ KBs $\mathcal{K}$, because if $E \in \mathrm{Qc}(\mathcal{K})$ is not consistent with $\mathcal{K}$, then $(\exists r.E)^\mathcal{I} = \perp^\mathcal{I} = \emptyset$ holds for all $\mathcal{I} \in \mathtt{Mod}(\mathcal{K})$. Thus, replacing all existential quantifications over inconsistent quantified concepts by $\perp$ always yields an equivalent QC-consistent KB $\mathcal{K}'$ (i.e. $\mathtt{Mod}(\mathcal{K}) = \mathtt{Mod}(K')$).

As described in the end of Section 2.1, the most interesting task is to find commonalities among all models of the KB. Such commonalities can be a variety of different properties. The most basic such properties, those that are the subject of this work, are GCIs and concept assertions. It can be readily seen, that an explicit representation of the fact that $Dog \sqsubseteq \exists afraid.Cat$ and $Cat \sqsubseteq Smart$ results in *all* models of this KB to *implicitly* satisfy $Dog \sqsubseteq \exists afraid.(Cat \sqcap Smart)$. Similarly, asserting the property Cat to the individual $daisy$ ($Cat(daisy)$) and generally expecting all Cats to be $Smart$ ($Cat \sqsubseteq Smart$), allows to draw the conclusion that $Smart(daisy)$ is satisfied in all models of the KB.

Definition 2.9 (Subsumption and Instance Checking). Let $\mathcal{K} = (\mathcal{A}, \mathcal{T})$ be a KB and α a query. The tuple $(\mathcal{K}, \alpha)$ is called an *inference problem* (or reasoning task). Given an inference problem, the projections $\mathtt{QI}((\mathcal{K}, \alpha)) = \alpha$ and $\mathtt{KB}((\mathcal{K}, \alpha)) = \mathcal{K}$ are used to extract the query and KB, respectively. An inference problem is a *classical subsumption query*, if α is a GCI and a *classical instance check*, if α is a concept assertion. The classical inference problem $(\mathcal{K}, \alpha)$ is true under *classical semantics iff*

$$\forall \mathcal{I} \in \mathtt{Mod}(\mathcal{K}). \mathcal{I} \models \alpha. \tag{2.1}$$

For a query $C \sqsubseteq D$ or $D(a)$, C and a are denoted as the *query subject*, respectively.

Referring to an inference problem as a tuple of a KB and query, allows many of our results to remain formally clean. This will play an important role in PART II, where we discuss reductions allowing to decide one inference problem, by transforming it to another. At the same time, we will also make use of the more widely adopted notation, extending the satisfaction operator $\models$ that we used with single interpretations. More explicitly, for classical entailment we shall use the shorthand $\mathcal{K} \models \alpha$ (say $\mathcal{K}$ classically entails α), when $(\mathcal{K}, \alpha)$ is true under classical semantics. The main contributions of this work introduce semantics going beyond classical, formally defining different variants of the entailment operator "$\models$", to capture more expressive queries than classical subsumption and instance checks.

DEFEASIBLE REPRESENTATION AND REASONING. Before stepping into more advanced aspects of classical Description Logics, we extend the standard notions from above by a more vague form of knowledge representation. As argued in the beginning, knowledge is often neither fixed, nor complete. It may develop over time and introduce aspects that contradict previously posed axioms. If all axioms are strict, e.g. GCIs, such contradictory knowledge can lead to the inconsistency of a concept w.r.t. a KB, or even worse, the KB itself. From Definition 2.9, it becomes clear why an inconsistent KB is counterproductive for any KR scenario. If $\mathtt{Mod}(\mathcal{K}) = \emptyset$, then (2.1) is trivially satisfied for any classical query α, rendering the represented knowledge ineffectual. Likewise, if $C^{\mathcal{I}} = \emptyset$ in all models $\mathcal{I} \in \mathtt{Mod}(\mathcal{K})$, then $\mathcal{K} \models C \sqsubseteq D$ holds trivially for any concept D. There are numerous approaches and ideas to treat such a development of knowledge adequately. An overview is given in Chapter 3. Here we only introduce the basic notions that are required for the semantics we investigate.

Instead of expressing all knowledge as strict and irrefutable GCIs or assertions, Defeasible Description Logics allow for the modelling of defeasible axioms. Intuitively, such axioms can be understood to describe the properties of a *typical member* of a concept. As such, instances of this concept are expected to satisfy also these typical properties. They are exempt of this, only if they satisfy more specific properties, contradicting said typical behaviour. In a manner of speaking, this exemption can be seen as *defeating* typical information.

Formally, a *defeasible concept inclusion* (DCI) is an axiom $C \sqsubsetsim D$, reading, "elements of C are *usually* also elements of D". DCIs are gathered in a finite set, called the DBox, typically denoted with $\mathcal{D}$. From here on out, a KB is, in general, a triple $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$. When any of those boxes is empty, we can freely describe $\mathcal{K}$ as the remaining pair, or even denote any single box as the knowledge base as well. If required, we explicitly introduce a KB as *defeasible*, if $\mathcal{D} \neq \emptyset$, and as *strict*, if $\mathcal{D} = \emptyset$. For a (general) KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, its strict part is formally captured with $\mathcal{K}_{\mathtt{strict}} = (\mathcal{A}, \mathcal{T})$. The syntactic notions of Definition 2.5 carry over to DCIs in the natural way (essentially from GCIs). In particular, the notion of a domain element actively/passively satisfying a DCI, is paramount for comparing "typicality" of elements w.r.t. the represented defeasible knowledge.

To separate classical from defeasible knowledge, also in terms of inferences, a KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ is expected to have the same *classical* consequences as $\mathcal{K}_{\mathtt{strict}}$. Therefore, a knowledge base must be explicitly queried for the entailment of defeasible knowledge, which should be determined by taking the DBox into account. The queries we consider here, are the defeasible variants of subsumption and instance checking. Formally, a *defeasible* subsumption query is presented in the form of a DCI $C \sqsubsetsim D$ and the defeasible variant of an instance check is described as $C\{a\}$. The separation of classical and defeasible instance checking is only precedented in [CMVM'13], while earlier approaches made no such distinction, somewhat blending them together [CS'10; CS'12]. An inference problem $(\mathcal{K}, \alpha)$ is called defeasible, if the query α is either a defeasible subsumption or a defeasible instance check.

Some of the original studies for defeasible logics by Loui [Lou'87] and Nute [Nut'01] introduce strict knowledge, defeasible rules and *defeaters*. The purpose of defeaters is to disable defeasible rules under the characterised conditions. The role of strict knowledge in defeasible DLs is of course taken on by GCIs and assertive axioms. DCIs from the DBox act as both defeasible rules and defeaters. The motivation for defeating a rule is to maintain consistency of the query subject, i.e. exclude contradicting DCIs.

For a brief illustration, consider the defeasible assumption that cats are usually smart, $\mathtt{Cat} \sqsubsim \mathtt{Smart}$, and that cats who befriend dogs are usually not smart, $\mathtt{Cat} \sqcap \exists \mathtt{friend}.\mathtt{Dog} \sqsubsim \neg\mathtt{Smart}$. These DCIs can never be both actively satisfied by any domain element in any interpretation. Hence, deciding an appropriate query in a meaningful way, requires one or the other to prevail[1]. How to determine this interaction between DCIs and produce meaningful defeasible entailments of a KB is the main subject of this thesis. A continuously studied approach from the literature is fully formalised in a new and abstract way in PART II. Our novel, even stronger model-theoretic semantics are introduced and analysed in full in PART III.

2.3 FUNDAMENTAL RESULTS

In the following we present several general results on classical reasoning in DLs, that are essential in the technical parts of this work. These results are considered folklore in the area of Description Logics, hence, we shall be brief in their presentation. For more thorough details, the reader is politely referred to the introductory literature once more, specifically Chapters 2,3 in [BCM+'10] and Chapter 6 in [BHLS'17].

DISJOINT MODEL UNION PROPERTY. The DL $\mathcal{ALC}$ and its sublogics all enjoy the so-called *disjoint model union property* (DMUP), a notion that is quite common in modal logics[2]. The intuition behind this property is that two models can easily be joined into a new interpretation that also satisfies the given knowledge base. The practical benefit is that this disjoint union interpretation retains some of the features that each of the original models exhibited (cf. Example 2.10). The most common application of this property is in proofs, when it is necessary to derive a single model lacking or satisfying two (or more) features, which can originally be only assumed separately for two distinct models.

Formally, the *disjoint union* of two interpretations $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, $\mathcal{J} = (\Delta^{\mathcal{J}}, \cdot^{\mathcal{J}})$ is defined as $\mathcal{I} \uplus \mathcal{J} = (\Delta^{\mathcal{I}} \uplus \Delta^{\mathcal{J}}, \cdot^{\mathcal{I} \uplus \mathcal{J}})$ with $A^{\mathcal{I} \uplus \mathcal{J}} = A^{\mathcal{I}} \uplus A^{\mathcal{J}}$ and $r^{\mathcal{I}} \uplus r^{\mathcal{J}}$. When individuals are involved in the satisfaction of a KB

1 If an element satisfies $\mathtt{Cat}$ and is not related to any $\mathtt{Dog}$ with the role $\mathtt{friend}$, then it can satisfy *both* DCIs, the second one passively. For argument's sake we consider this prevalence of the first DCI.

2 Schild [Sch'91b] showed that $\mathcal{ALC}$ is a syntactic variant of the modal logic **K**, and thus inherits its properties, such as the DMUP.

$\mathcal{K}$, their extension under $\mathcal{I} \uplus \mathcal{J}$ can match either that of $\mathcal{I}$ or $\mathcal{J}$, but has to be chosen consistently. If the DMUP is satisfied for a DL $\mathcal{L}$, and the interpretations $\mathcal{I}$, $\mathcal{J}$ satisfy an $\mathcal{L}$ KB $\mathcal{K} = (\mathcal{A}, \mathcal{T})$, then the disjoint union of $\mathcal{I}$ and $\mathcal{J}$ also satisfies $\mathcal{K}$.

Example 2.10. Suppose a KB $\mathcal{K}$ does not entail $C \sqsubseteq D$ and $C \sqsubseteq E$. A priori, we can only assume that there exists some model $\mathcal{I}$ with $d \in C^{\mathcal{I}} \setminus D^{\mathcal{I}}$ and some $\mathcal{J}$ with $e \in C^{\mathcal{J}} \setminus E^{\mathcal{J}}$. If $\mathcal{K}$ is an $\mathcal{ALC}$ KB (or below), then $\mathcal{I} \uplus \mathcal{J}$ is a model of $\mathcal{K}$ and $d, e \in C^{\mathcal{I} \uplus \mathcal{J}}$, s.t. $d \notin D^{\mathcal{I} \uplus \mathcal{J}}$ and $e \notin E^{\mathcal{I} \uplus \mathcal{J}}$ are both satisfied in $\mathcal{I} \uplus \mathcal{J}$.

REASONING IN $\mathcal{ALC}$. Algorithms for reasoning in DLs are often tailored specifically to the set of concept constructors and axioms that are available in this DL. For $\mathcal{ALC}$ specifically, a form of search and backtracking is required to investigate the different options for satisfaction that are opened up by disjunction ($\sqcup$). For subsumption and instance checking w.r.t. general concept inclusions in $\mathcal{ALC}$, tableaux algorithms are typically employed [DM'00]. Such algorithms run in exponential time, matching the lower bound for such reasoning in $\mathcal{ALC}$, as proved by Schild [Sch'91a]. In our own complexity investigation in Chapter 8, the ExpTime-completeness of $\mathcal{ALC}$ is our only concern, which is why we leave the comments on tableaux algorithms at that.

REASONING IN $\mathcal{EL}_{\bot}$. The lightweight Description Logic $\mathcal{EL}$ and siblings of the so-called $\mathcal{EL}$-family are predominantly studied for their computational prowess. Baader et al. [BBL'05] and Brandt [Bra'04] show that deciding classical entailment of subsumption remains P-complete, even in the logic $\mathcal{EL}^{++}$ with GCIs. This logic encompasses $\mathcal{EL}_{\bot}$ and certain concept constructors, allowing to model assertive knowledge—as introduced here within the ABox—relying only on GCIs. Therefore, this tractability trivially transfers also to the problem of deciding classical instance checks in $\mathcal{EL}_{\bot}$.

An additional product of [BBL'05; Bra'04] is the *canonical model property* that members of the $\mathcal{EL}$-family possess. Recall that the standard definition of entailment in DLs relies on satisfaction of a query by *all* models of the underlying knowledge base. It turns out that, in such simple DLs, there exist models that are *canonical* for the set of all models, in terms of their shared entailments. To be more precise, a model of a KB is considered canonical (w.r.t. a certain signature), if it can be homomorphically embedded into *every* other model of the KB. Consequently, if an entailment is supported by a canonical model, it must be entailed by all models. A more formal account of this canonicity is presented in Section 2.4, in a slightly unconventional, but equivalent approach. The conventional formalisation of canonical models utilises *graph simulations*. For a thorough exposition on that subject, consider Lutz and Wolter [LW'10].

As a consequence of [BBL'05; Bra'04], such models can be computed in polynomial time, and are of polynomial size in the original input. This allows for very intuitive algorithmic solutions, in particular for computing non-standard inferences such as most specific concept [Neb'90; SL'83], least common subsumer [CBH'92; Tur'07], or concept similarity [Eck'17; EPT'15], as well as defeasible entailment [PT'17a; PT'18] (Section 8.1), among others.

2.4 ADVANCED NOTIONS

We close the introductory part on Description Logics with a formal account of two notions that are specifically tailored to this thesis. As such, they are rarely adopted in other areas of research in DLs, and count towards the contributions of this work. The first is a straightforward extension of the set operations intersection or subset, to interpretations over a shared domain. For the second, we present a view on canonicity in classical $\mathcal{EL}_\bot$, which is used as the main motivation for the formal construction of the defeasible semantics introduced in Chapter 7.

SET-LIKE INTERPRETATION OPERATIONS. The approach we take in Chapter 7 to achieve non-monotonic semantics relies heavily on interpretations sharing the same domain. Such interpretations behave very well when treated as tuples of sets. Standard set operations can be lifted to these interpretations, applying the operation component-wise to the extension of every concept and role name.

Definition 2.11 (Interpretation-Operations). Let $\mathcal{I} = (\Delta, \cdot^\mathcal{I})$, $\mathcal{J} = (\Delta, \cdot^\mathcal{J})$ be two interpretations over a shared domain Δ.

1. $\mathcal{I} \subseteq \mathcal{J}$ *iff* $\forall A \in N_C, r \in N_R.(A^\mathcal{I} \subseteq A^\mathcal{J} \land r^\mathcal{I} \subseteq r^\mathcal{J})$
   and $\forall a \in N_I. a^\mathcal{I} = a^\mathcal{J}$
2. $\mathcal{I} \cap \mathcal{J} = (\Delta, \cdot^{\mathcal{I}\cap\mathcal{J}})$ with $A^{\mathcal{I}\cap\mathcal{J}} = A^\mathcal{I} \cap A^\mathcal{J}$, $r^{\mathcal{I}\cap\mathcal{J}} = r^\mathcal{I} \cap r^\mathcal{J}$ *if* $a^\mathcal{I} = a^\mathcal{J}$ (for all $A \in N_C, r \in N_R, a \in N_I$)

If $\mathcal{I}$ and $\mathcal{J}$ disagree on the extensions of individuals, then their intersection is undefined.

A simple consequence of Definition 2.11 is that the subset relation between two interpretations over a shared domain translates in the straightforward way to their extension of $\mathcal{EL}_\bot$ concepts.

Lemma 2.12. *For two interpretations $\mathcal{I}, \mathcal{J}$ over a shared domain Δ, it holds that*

$$\mathcal{I} \subseteq \mathcal{J} \text{ implies } C^\mathcal{I} \subseteq C^\mathcal{J} \text{ for all } C \in \mathfrak{C}(\mathcal{EL}_\bot).$$

Proof. We prove the claim by induction on the structure of C. The cases $C = A \in N_C$ and $C = E \sqcap F$ are trivial (under the hypothesis that the claim holds for E and F). Let $C = \exists r.E$, for $r \in N_R$ and assume the claim holds for E. By definition, $(\exists r.E)^\mathcal{I} = \{d \in \Delta \mid \exists (d,e) \in r^\mathcal{I}. e \in E^\mathcal{I}\}$. From $\mathcal{I} \subseteq \mathcal{J}$, i.e. $r^\mathcal{I} \subseteq r^\mathcal{J}$ and $E^\mathcal{I} \subseteq E^\mathcal{J}$, we can directly conclude

$$(\exists r.E)^\mathcal{I} \subseteq \{d \in \Delta \mid \exists (d,e) \in r^\mathcal{J}. e \in E^\mathcal{J}\}. \qquad \square$$

This preliminary result comes immediately into action in the following.

CONTEXT RESTRICTION AND REPRESENTABILITY IN $\mathcal{EL}_\perp$. Intuitively, every domain element of a canonical model (in $\mathcal{EL}_\perp$) represents an entire class of objects (or a specific named individual). As such, these *class-representatives* share precisely the information that all elements of this class share in all models of the underlying knowledge base. One could say, the canonical model "knows everything that the KB entails about this class". On the flipside, canonical models can only be used to determine entailments, if they accommodate for *enough* information. More precisely, deciding entailment of a subsumption query $C \sqsubseteq D$ or an instance check $D(a)$, is only possible if the canonical model contains a domain element representing the concept C, or respectively, the individual a. Luckily, for every (consistent) query subject, an appropriate canonical model can be constructed (given that the KB is consistent). From an application's point of view, it also makes sense to fix, a priori, a set of concepts/individuals of interest. Then, the corresponding canonical model needs to be computed only once, and can be repeatedly queried for the entailments it supports. Formally, this set of concepts/individuals of interest can be generalised as a *context* of relevant terms.

Definition 2.13 (Relevant Context). A *relevant context* $\mathbb{C}$, $\mathbb{O}$ (short: context) is comprised of a set of individuals $\mathbb{O} \subseteq N_I$, and a set of concepts $\mathbb{C} \subseteq \mathfrak{C}(\mathcal{L})$, over the (present) DL $\mathcal{L}$. For a KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, a context $\mathbb{C}$, $\mathbb{O}$,

- *contains* $\mathcal{K}$ *iff* $\mathrm{Qc}(\mathcal{K}) \subseteq \mathbb{C}$ and $\mathrm{sig}_I(\mathcal{A}) \subseteq \mathbb{O}$,
- is *consistent* with $\mathcal{K}$ *iff* $\forall C \in \mathbb{C}.\mathcal{K} \not\models C \sqsubseteq \perp$, and
- is *quantification closed iff* $\mathrm{Qc}(\mathbb{C}) \subseteq \mathbb{C}$.

Furthermore, a query $\alpha \in \{C \sqsubseteq D, C \precsim D, D(a), D\{a\}\}$ is said to be *over* a context $\mathbb{C}$, $\mathbb{O}$, if $C, D \in \mathbb{C}$ and $a \in \mathbb{O}$.

Note that we often introduce "a consistent context containing $\mathcal{K}$", implicitly associating its consistency with the KB $\mathcal{K}$.

Requiring a canonical model to contain certain concepts (such as query subjects), is common practice ([LW'10]). This is generalised with a relevant context. Another widely accepted use of a context is to simplify a problem by restricting it to a reasonable (e.g. finite) context, rather than dealing with an unbounded number of relevant subjects (e.g. [Bon'19]). We use the context to show that reasoning in $\mathcal{EL}_\perp$ can be reduced to reasoning over models *only* representing concepts and individuals given in a predefined context, if this context is sufficiently large. This is a slightly weaker result than the canonical model property, but its exhibition is prototypical for the construction in Chapter 7. The representation of concepts and individuals begins with a direct association of the context in terms of an interpretation domain.

Definition 2.14 (Representative Domain)**.** For a context $\mathbb{C}$, $\mathbb{O}$, the *representative domain* is defined as

$$\Delta^{\mathbb{C},\mathbb{O}} = \mathbb{C} \cup \mathbb{O}.$$

A domain element $d \in \Delta^{\mathbb{C},\mathbb{O}}$ is called

- a *concept representative*, if $d \in \mathbb{C}$, and
- an *individual representative*, if $d \in \mathbb{O}$.

Interpretations over a representative domain are called representative interpretations, or models, if they satisfy a given KB. Note that all the notations introduced in Section 2.1 and 2.2 clearly also apply to interpretations over representative domains. The elements in a representative domain are given meaning in representative interpretations, by securing a form of well-behaviour, for such interpretations.

Definition 2.15 (Standard Interpretations)**.** For a context $\mathbb{C}$, $\mathbb{O}$, an interpretation $\mathcal{I}$ over the representative domain $\Delta^{\mathbb{C},\mathbb{O}}$ is *standard iff*

1. $C \in C^{\mathcal{I}}$ for all $C \in \mathbb{C}$,
2. $a = a^{\mathcal{I}}$ for all $a \in \mathbb{O}$, and
3. $d \in (\exists r.C)^{\mathcal{I}}$ implies $(d, C) \in r^{\mathcal{I}}$ for all $d \in \Delta^{\mathbb{C},\mathbb{O}}$, $C \in \mathbb{C}$.

The standard property of representative interpretations ensures several things. When a domain element is associated with a concept C (or an individual a), the standard property will ensure that an interpretation satisfying a KB $\mathcal{K}$ maintains at least the information that is entailed for C (or a) by $\mathcal{K}$, within the respective domain element. The final aspect of standard interpretations ensures a sort of uniform behaviour for witnesses to the satisfaction of existential restrictions. This allows to intersect interpretations in terms of Definition 2.11 and expect this intersection to adhere to Lemma 2.12. It is essentially a technical attribute that simplifies a plethora of the results throughout this entire thesis.

As promised, we capture a set of models of a KB that is restricted to standard models over a given representative domain.

Definition 2.16 (Standard Models)**.** For a quantification closed context $\mathbb{C}$, $\mathbb{O}$, containing $\mathcal{K}$, the set of all *standard models* of $\mathcal{K}$ for $\mathbb{C}$, $\mathbb{O}$ is

$$\texttt{Mod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) = \{\mathcal{I} \in \texttt{Mod}(\mathcal{K}) \mid \mathcal{I} = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\mathcal{I}}) \wedge \mathcal{I} \text{ is standard}\}. \quad (2.2)$$

At this point, it is clear what we meant by a "large enough" context. To fully ensure well-behaviour in terms of Def. 2.15, it must be guaranteed that the required witnesses to any potentially satisfied existential restriction (those occurring in $\mathcal{K}$) also occur in the context.

The reduction of classical entailment in $\mathcal{EL}_\perp$ to entailment under standard representative models (over an appropriate context) is contingent on one essential result. We show that every arbitrary interpretation $\mathcal{I}$ can be converted to a standard representative interpretation.[3] This "normalised" standard interpretation, is shown to maintain the information that $\mathcal{I}$ holds for all concepts and individuals of the context. In particular, if $\mathcal{I}$ satisfies a KB, it follows that the constructed standard interpretation does so as well.

Lemma 2.17. *For a knowledge base $\mathcal{K}$ and a consistent, quantification closed context $\mathbb{C}, \mathbb{O}$ containing $\mathcal{K}$ and any arbitrary interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, let $\mathcal{J}(\mathcal{I}) = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\mathcal{J}(\mathcal{I})})$ with*

$$
\begin{aligned}
A^{\mathcal{J}(\mathcal{I})} &= \{C \in \Delta^{\mathbb{C},\mathbb{O}} \mid \mathcal{I} \models C \sqsubseteq A\} \\
&\quad \cup \{a \in \Delta^{\mathbb{C},\mathbb{O}} \mid \mathcal{I} \models A(a)\} \\
r^{\mathcal{J}(\mathcal{I})} &= \{(C, D) \in \mathbb{C} \times \mathbb{C} \mid \mathcal{I} \models C \sqsubseteq \exists r.D\} \\
&\quad \cup \{(a, D) \in \mathbb{O} \times \mathbb{C} \mid \mathcal{I} \models (\exists r.D)(a)\} \\
&\quad \cup \{(a, b) \in \mathbb{O} \times \mathbb{O} \mid (a^{\mathcal{I}}, b^{\mathcal{I}}) \in r^{\mathcal{I}}\} \\
a^{\mathcal{J}(\mathcal{I})} &= a
\end{aligned}
$$

for all $A \in N_C$, $r \in N_R$, $a \in N_I$. It holds that

1. $C \in D^{\mathcal{J}(\mathcal{I})}$ iff $\mathcal{I} \models C \sqsubseteq D$ *(for $C, D \in \mathbb{C}$),*
2. $a \in D^{\mathcal{J}(\mathcal{I})}$ iff $\mathcal{I} \models D(a)$ *(for $D \in \mathbb{C}$, $a \in \mathbb{O}$), and*
3. $\mathcal{I} \in \mathtt{Mod}(\mathcal{K})$ *implies* $\mathcal{J} \in \mathtt{Mod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$

Proof.

CLAIM 1 AND 2. We prove Claim 1 and 2 by induction on the concept D. The base case for $D = A \in N_C$ follows trivially by definition of $\mathcal{J}(\mathcal{I})$ for both $C \in \mathbb{C}$ and $a \in \mathbb{O}$. For the induction step $D = \exists r.E$, we treat the first two claims separately, under the assumption that the claims hold for E. The case of $D = E \sqcap F$ is trivial.

$C \in (\exists r.E)^{\mathcal{J}(\mathcal{I})}$ implies that there is some $(C, F) \in r^{\mathcal{J}(\mathcal{I})}$ with $F \in E^{\mathcal{J}(\mathcal{I})}$ (the successor of this edge can only be some $F \in \mathbb{C}$, by definition of $\mathcal{J}(\mathcal{I})$). From the induction hypothesis, it follows that $\mathcal{I} \models F \sqsubseteq E$. By construction, $(C, F) \in r^{\mathcal{J}(\mathcal{I})}$ implies $\mathcal{I} \models C \sqsubseteq \exists r.F$, collectively proving $\mathcal{I} \models C \sqsubseteq \exists r.E$. For the other direction, note that, $E \in \mathbb{C}$ is guaranteed by $D \in \mathbb{C}$ and the context being quantification closed (i.e. $Qc(D) \subseteq \mathbb{C}$). Thus, $\mathcal{I} \models C \sqsubseteq \exists r.E$ directly implies $(C, E) \in r^{\mathcal{J}(\mathcal{I})}$ by construction of $\mathcal{J}(\mathcal{I})$. The induction hypothesis implies $E \in E^{\mathcal{J}(\mathcal{I})}$ (because trivially, $\mathcal{I} \models E \sqsubseteq E$), and therefore, $C \in (\exists r.E)^{\mathcal{J}(\mathcal{I})}$.

The only difference for Claim 2 is that $a \in (\exists r.E)^{\mathcal{J}(\mathcal{I})}$ could be the result of $(a, b) \in r^{\mathcal{J}(\mathcal{I})}$ with $b \in E^{\mathcal{J}(\mathcal{I})}$. In this case, the induction hypothesis is $\mathcal{I} \models E(b)$, which allows for two insights:

3 This transformation can be somewhat regarded as the reverse of a *homomorphic embedding* of a standard representative model, into an arbitrary model of the KB.

1. $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in r^{\mathcal{I}}$ implies $\mathcal{I} \models (\exists r.E)(a)$, proving the claim.
2. Additionally, $\mathcal{I} \models (\exists r.E)(a)$ implies $(a, E) \in r^{\mathcal{J}(\mathcal{I})}$, showing that every named witness (individual representative) to $a \in (\exists r.E)^{\mathcal{J}(\mathcal{I})}$ is accompanied by the appropriate anonymous witness E, if $E \in \mathbb{C}$.

CLAIM 3. First of all, the final consideration of the proof for Claim 1 and 2 (together with Claim 1 and 2), shows that $\mathcal{J}(\mathcal{I})$ satisfies all properties of standard interpretations (Def. 2.15). Every $C \in \mathbb{C}$ belongs to $C^{\mathcal{J}(\mathcal{I})}$ (Claim 1), every $a \in \mathbb{O}$ is assigned to $a^{\mathcal{J}(\mathcal{I})}$ (by construction), and, regardless of the successor element witnessing $C \in (\exists r.E)^{\mathcal{J}(\mathcal{I})}$ (or $a \in (\exists r.E)^{\mathcal{J}(\mathcal{I})}$), $(C, E) \in r^{\mathcal{J}(\mathcal{I})}$ is satisfied ($(a, E) \in r^{\mathcal{J}(\mathcal{I})}$, respectively).

Assume $\mathcal{J}(\mathcal{I})$ does not satisfy some GCI or ABox assertion from $\mathcal{K}$. The construction of $\mathcal{J}(\mathcal{I})$, together with Claim 1 and 2, shows that $\mathcal{I}$ cannot be a model of $\mathcal{K}$. □

Remark 2.18. Note that for the *if*-direction of Claim 1 and 2 in Lem. 2.17 it is enough to rely on $Qc(D) \subseteq \mathbb{C}$. This property is naturally satisfied when applying those claims to GCIs and concept assertions of a KB that is contained in the context. The *only-if*-direction of the first two claims in Lem. 2.17 does not rely on $D \in \mathbb{C}$ at all.

We rely on Remark 2.18 to prove the desired main result without explicitly making restrictions on the context as in Claim 1 and 2 of Lem. 2.17. Finally, this result shows that, in classical $\mathcal{EL}_{\bot}$, a restriction to a relevant context is (almost) no restriction at all.

Theorem 2.19. *For a KB $\mathcal{K}$ and a consistent, quantification closed context $\mathbb{C}$, $\mathbb{O}$, containing $\mathcal{K}$, it holds for every $C \in \mathbb{C}$, $a \in \mathbb{O}$, that*

1. $\mathcal{K} \models C \sqsubseteq D$ iff $C \in D^{\mathcal{I}}$ *for all* $\mathcal{I} \in \mathtt{Mod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$, *and*
2. $\mathcal{K} \models D(a)$ iff $a \in D^{\mathcal{I}}$ *for all* $\mathcal{I} \in \mathtt{Mod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$.

Proof. The *only-if*-direction follows trivially from $\mathtt{Mod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \subseteq \mathtt{Mod}(\mathcal{K})$ and Definition 2.15 for both claims.

For the *if*-direction of both claims, assume there is a model $\mathcal{I} \in \mathtt{Mod}(\mathcal{K})$ that does not satisfy $C \sqsubseteq D$ (or $D(a)$). Constructing $\mathcal{J}(\mathcal{I})$ as in Lemma 2.17 supplies a model in $\mathtt{Mod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ such that $C \notin D^{\mathcal{J}(\mathcal{I})}$ (or $a \notin D^{\mathcal{J}(\mathcal{I})}$). $\mathcal{J}(\mathcal{I})$ is a model of $\mathcal{K}$, because we can utilise the entirety of Lem. 2.17 and $C \notin D^{\mathcal{J}(\mathcal{I})}$ (or $a \notin D^{\mathcal{J}(\mathcal{I})}$) relies only on the *only-if*-direction of Claim 1 and 2 in the preceding lemma (i.e. not relying on $D \in \mathbb{C}$). □

The only downside of this context restriction, is that Theorem 2.19 applies only to queries whose subjects are part of the given context. Nevertheless, this idea holds theoretical value, not only for the main body of this work. Among others:

1. Standard representative models are uniform, which allows for easy comparison via $\subseteq$ and combination via $\cap$.

2. As a matter of fact, using Lemma 2.12, it can be shown that standard models are closed under intersection, allowing to derive a $\subseteq$-smallest model in $\mathtt{Mod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$. This smallest model is effectively an alternative characterisation of a canonical model for the KB $\mathcal{K}$.

3. Domain elements can represent entire concepts, allowing to determine subsumption entailments by checking single elements, rather than all members of a concept extension.

4. In many cases, it suffices to consider a finite context, resulting in a finite representative domain. Due to finiteness of $\mathcal{K}$ (and thus of $\mathtt{sig}(\mathcal{K})$), this trivially guarantees the set $\mathtt{Mod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ to be finite.

We abstain from formally showing Point 2 of the above. A variant of this line of arguments will be presented in Section 7.2 for an extension of representative models that allows treatment of defeasible information and characterisation of non-monotonic semantics for $\mathcal{EL}_{\bot}$.

3 NON-MONOTONIC REASONING

"I propose to consider the question, 'Can machines think?'"

Turing [Tur'50, p. 433]

Ever since Turing discussed machine intelligence, different methods have been investigated to automate reasoning over formal knowledge (KR). What we consider to be classical approaches, generally exhibit a *monotonic* behaviour in the inferences that they are able to derive. More formally, for knowledge bases $\mathcal{K}, \mathcal{K}'$ and any query α:

$$\mathcal{K} \models \alpha \text{ implies } \mathcal{K}' \models \alpha \text{ for } \mathcal{K} \subseteq \mathcal{K}'. \quad (3.1)$$

While satisfaction of this property is proof of a reliable, stable, and reasonably objective KR formalism, monotonic reasoning semantics are not suited to adequately model some very intuitive scenarios. Knowledge is often subjective and evolving over time. This easily results in the need to formally represent contradictory information, be it because of discrepancies among different sources (e.g. semantic web), or the need to express generic behaviour in the light of incomplete information, expecting to encounter exceptions of such behaviour in the future. Reconsider our playful example from Chapter 2, stating that all Cats are Smart. When learning of single individuals, such as molli befriending a Dog (which, for argument's sake, we shall consider to be evidence of not being Smart), this general statement ceases to be satisfiable. More explicitly, if no interpretation can satisfy two such axioms (say α, α') simultaneously, then a KB containing them ($\alpha, \alpha' \in \mathcal{K}$) has no model. In classical formalisms, such contradictory information, together with monotonicity, renders the extended KB meaningless, *ex falso quodlibet*. A formalism that is not committed to monotonicity, could resolve such inconsistencies by retracting either α or α', thereby maintaining consistency of the KB and the value of its entailments. The main difficulty in devising such non-monotonic formalisms has always been to provide reasonable arguments (as a form of well-behaviour) for the fundamental choice between α and α'.

Since the 60s, many different formalisms have emerged. While this is by no means a survey on non-monotonic semantics, we would like to briefly present the fundamental idea of three different categories of formalisms that are capable of such knowledge retraction. These are not necessarily considered as non-monotonic in the literature. For a full overview, consult Brewka et al. [BDK'97] or a more brief version in [Bre'89].

ABDUCTIVE REASONING describes the study of entailment relations that allow specifically for the individual retraction of previously drawn conclusions, to maintain consistency within the current set of inferences. A cornerstone of such reasoning mechanisms is that consistency is maintained without altering previously represented knowledge. These formalisms are most commonly referred to as non-monotonic.

BELIEF REVISION studies the evolution of knowledge bases, while relying on classical semantics to compute entailments ([Gär'92]). Whenever contradictory knowledge is learned, the KB is analysed for its inconsistencies and *revised* to produce a consistent variant of this KB. To not venture too far from the entailments of the original KB, formal properties are imposed on such a revision, an idea originating from Alchourrón et al. [AGM'85]. Motivated by the evolution of knowledge in the semantic web, revision has received and continues to receive a lot of attention also in DL research. For a brief survey, consider Qi and Yang [QY'08].

INCONSISTENCY TOLERANT REASONING is studied with a slightly different purpose ([BHS'05]). In reasoning under *repair semantics*, the evolution of knowledge is not to be tempered with as done in belief revision. In light of contradictory knowledge in the KB, information is not being rewritten, retracted or changed in any way. However, to obtain a useful and consistent set of consequences, consistent subsets of the original KB are determined. It is then studied how to supply a user with sound consequences based on a potential variety of distinct, so-called *repairs* of the KB. This approach is particularly well studied for Description Logics [BR'13; Bou'16; LLR+'10], due to their application to (often contradictory) collaborative sources of knowledge, as most commonly encountered in the semantic web. Orthogonal approaches to repair semantics consider multivalued logics (origin: [Bel'77], DL: [MMH'13]), to accommodate for inconsistent knowledge.

Note that these categories are somewhat independent of the underlying representation formalism. As such, they are not only remnants of the past century but are of great influence for modern formalisms such as Description Logics ([LLR+'10; MMH'13; QY'08]). In particular, ours is a study of abductive reasoning in DLs. Therefore, we discuss different approaches to achieve such reasoning in slightly more detail. The list we present in Section 3.1 is a selection of the fundamental ideas that have already been adopted to DLs in the past, thus, classifying them as related work for the present contribution. This list is leading up to the immediate foundations of our own semantics, the so-called *preferential logics*, which we will present in more detail in Section 3.2.

3.1 COMMON APPROACHES TO ABDUCTIVE REASONING

The following is a collection of noteworthy approaches to non-monotonic reasoning in DLs, which are selected for two reasons:

1. For their prominence among related approaches, evidenced by their persistence and continued adaptation to new formalisms, and

2. For their successful coalescence with the KR formalism of Description Logics.

CIRCUMSCRIPTION. The technique of circumscription originates from McCarthy [McC'80], circumscribing predicates and axioms in classical first-order logic. The central intuition is to explicitly allow elements in the domain to contradict certain axioms, but to impose a form of minimality on such exceptions.

To be slightly more concrete, we move to a more intuitive version of circumscription, not coincidentally extending classical DLs as per Bonatti et al. [BLW'06]. Strict implications, such as $\mathit{Cat} \sqsubseteq \mathit{Smart} \sqcup \mathit{Abnormal}$ would be modelled with exceptional cases in mind. A *circumscription pattern* supplements the KB to state which concept or role names are to be minimised. In this short example, it would be intuitive to minimise the number of $\mathit{Abnormal}$ instances. This would allow individuals such as molli befriending a Dog to remain consistent with this strict statement. Minimisation forces more generic individuals, e.g. daisy, to satisfy Smart, for lack of a sufficient reason to not be Smart. It is worth noting that the adaptation of circumscription to Description Logics exhibits an increase in reasoning complexity [BLW'09] over classical reasoning, also in lightweight DLs [BFS'11]. This is not always the case in adaptations of non-monotonic semantics to DLs.

DEFAULT LOGIC. The origin of default logic is typically pointed to Reiter [Rei'78], with a more advanced state in [Rei'80]. The essence of default reasoning is not the acceptance of exceptional cases, but rather a form of consistent completion of incomplete knowledge. A new form of axiomatic implication is introduced as *default rules*, allowing to express "A implies B if it is plausible to assume C", e.g.

$$\frac{\mathit{Cat}(x) : \mathit{Smart}(x)}{\mathit{Smart}(x)}$$

That is, if x is a Cat and x satisfying Smart would be consistent with the remaining knowledge, then gain the conclusion that x is Smart. Difficulties in default reasoning arise when the consequents of two default rules invalidate each other's consistency checks, i.e.

$$\frac{\alpha : \neg\beta}{\gamma} \text{ and } \frac{\alpha : \neg\gamma}{\beta}$$

Both of these rules are sound, as long as the other is not "used". This phenomenon is argued to be desired, because of the assumptive nature of default rules. It leads to several *extensions* of the KB, which are considered to be different coherent sets of beliefs about the world. The main question for default reasoning is then how to combine distinct, consistent extensions of a KB to provide meaningful entailments. The immediate options that come to mind are formally introduced as

SCEPTICAL: derive only consequences supported by all extensions, and

CREDULOUS: derive all consequences supported by any extension.

Adaptations of Reiter's default logic for Description Logics were first investigated by Baader and Hollunder [BH'92] with more advanced results in [BH'95a]. This adaptation did not come without difficulties. Known drawbacks from its foundation, such as a lack of precedence among default rules, retained their severity in the DL setting. An improved variant of default semantics in DLs has appeared in [BH'95b], aiming to resolve this deficiency by consulting the hierarchical structure of DL concepts.

AUTO-EPISTEMIC LOGIC. An idea going back to Moore [Moo'85] is to employ an epistemic operator, allowing to capture the *belief in truth* in addition to factual truth. If there is no evidence to derive the truth of a sentence p, then it shall be *believed* that p is false, which is distinguishable from the *fact* that p is false by the epistemic operator. An extension of the knowledge base, providing evidence for the truth of p, forces to retract such consequences that were only supported by the belief about $\neg p$. Such semantics—their adaptations to DLs included [DLN+'98; DLN+'92; DNR'97]—rely on a technique known as *MKNF (minimal knowledge, negation as failure)*, a paradigm that is common also to logic programming [AB'94; Cla'78].

A continuous wave of research towards non-monotonic semantics in DLs and related formalisms (e.g. modal logics [BV'18a]) appears every year. Seeing as many adapted semantics suffer from crucial drawbacks (see Chapter 5), it is perhaps more promising to approach non-monotonicity in DLs directly. A noteworthy approach by Bonatti et al. [BFPS'15; BFS'10; BS'17] is introducing a notion of *overriding* for Description Logics. Essentially, *defeasible inclusions*, describing prototypical attributes of a *normality concept*, can be overridden in accordance with a *priority relation*, determining prevalence among conflicting statements.

Last, but certainly not least, on this list of non-monotonic semantics are *preferential logics*. With their roots in the work by Kraus et al. [KLM'90], they provide the foundation for the semantics we introduce and discuss in PART II and PART III.

3.2 PREFERENTIAL REASONING

The study of preferential reasoning revolves around the idea to axiomatise the expected behaviour of an agent, in a setting where it is necessary to draw valuable conclusions from incomplete knowledge. In case of a contradiction, premature conclusions must be allowed to be retracted or overridden in a non-monotonic fashion. This dynamic mechanism to draw conclusions, based on the assumptions represented in a KB, is often denoted as *defeasible reasoning*. Preferential reasoning was originally introduced for the propositional calculus around 1990 [KLM'90; LM'92]. The axioms proposed by these authors are famously referred to as the KLM-postulates. Many studies since then build on their notion of well-behaved non-monotonic reasoning, adopting their approach to different logics [Bou'94; BMV'11b; Del'98; LM'90].

Technically, the KLM-postulates characterise a multitude of consequence relations. The initial set of postulates proposed by KLM classifies the set of *preferential* entailment relations. Not much later, Lehmann and Magidor refined this set of axioms, introducing a weak form of monotonicity, called *rational monotonicity*, and in doing so, classifying a set of *rational* entailment relations. The study of such relations (in any logic formalism) usually revolves around two aspects. For one, a semantic characterisability for each of the accepted entailment relations is investigated. This is usually accomplished with *preferential semantics*, extending classical semantics with some form of preference relation on the treated entities. In [KLM'90; LM'92] specifically, this is a preference relation on different truth value assignments, or worlds (Def. 2.4 in [LM'92]). This aspect is usually leading up to representation results of the form "this (single) preferential interpretation satisfies certain properties *iff* the entailment relation it describes is preferential/rational". While this characterisation of entailment relations through semantic structures is inherently valuable, it lacks practicability for concrete scenarios where some user is presented with defeasible, or (as denoted by Lehmann and Magidor) *conditional knowledge*, and seeks valuable answers to queries over such knowledge. The second aspect of preferential reasoning emerged in [LM'92], fittingly introducing a formal notion of preference over all rational entailment relations. Lehmann and Magidor claim that any sensible non-monotonic entailment relation must derive *at least* those entailments that are contained in the *most preferred* rational entailment relation (Thesis 5.25 in [LM'92]). This most preferred entailment relation is famously known as the Rational Closure of the knowledge base. Their thesis lays the ground work for numerous investigations around the notion of Rational Closure in more expressive formalisms, including Description Logics.

To provide a better understanding for the foundation of our work, we formally introduce the basic notions of the KLM-approach in the following. As ours is a study revolving around Rational Closure, we refer the interested reader to [KLM'90; LM'92] for more details on representability of preferential and rational entailment relations.

3.2.1 *What does a Conditional Knowledge Base Entail?*

Lehmann and Magidor [LM'92] propose a form of knowledge representation that weakens the notion of logical implication into what they call *conditional assertions*. Such statements can express that "α *normally* entails β", formally $\alpha \mathrel{|\!\sim} \beta$. The conditionality of these assertions can be intuitively paraphrased

(Ref)	$\alpha \mathrel{|\!\approx} \alpha$	(LLE)	$\dfrac{\models \alpha \leftrightarrow \beta \quad \alpha \mathrel{|\!\approx} \gamma}{\beta \mathrel{|\!\approx} \gamma}$
(And)	$\dfrac{\alpha \mathrel{|\!\approx} \beta \quad \alpha \mathrel{|\!\approx} \gamma}{\alpha \mathrel{|\!\approx} \beta \wedge \gamma}$	(RW)	$\dfrac{\alpha \mathrel{|\!\approx} \beta \quad \models \beta \rightarrow \gamma}{\alpha \mathrel{|\!\approx} \gamma}$
(Or)	$\dfrac{\alpha \mathrel{|\!\approx} \gamma \quad \beta \mathrel{|\!\approx} \gamma}{\alpha \vee \beta \mathrel{|\!\approx} \gamma}$	(CM)	$\dfrac{\alpha \mathrel{|\!\approx} \beta \quad \alpha \mathrel{|\!\approx} \gamma}{\alpha \wedge \gamma \mathrel{|\!\approx} \beta}$
(RM)	$\dfrac{\alpha \mathrel{|\!\approx} \beta \quad \alpha \mathrel{|\!\not\approx} \neg\gamma}{\alpha \wedge \gamma \mathrel{|\!\approx} \beta}$		

Figure 3.1: Propositional KLM-postulates for preferential/rational entailment relations.

as "the assertion holds true under the condition that it is sound". That is, in light of a contradiction, the reasoning mechanism is allowed to disregard these implications. This process can be seen as a form of *defeasible* reasoning (cf. [Nut'01]), where the duty of defeaters is imposed directly on the defeasible (here: conditional) rules. Similarly, this formulation of the defeasibility condition has a strong resemblance to Reiter's default rules [Rei'80], with the difference that conditional assertions are evaluated with an inherent specificity-based[1] precedence. The lack of precedence is sometimes criticised for default logics [BH'95b].

What we describe as the reasoning mechanism, is formally captured with entailment relations $\mathrel{|\!\approx}$, which are essentially sets of conditional assertions, e.g. $\alpha \mathrel{|\!\sim} \beta \in \mathrel{|\!\approx}$. We often write $\mathrel{|\!\approx} \alpha \mathrel{|\!\sim} \beta$ or directly $\alpha \mathrel{|\!\approx} \beta$.

Definition 3.1 (Preferential and Rational Entailment Relations)**.** An entailment relation $\mathrel{|\!\approx}$ is *preferential iff* it satisfies the postulates *reflexivity* (Ref), *left logic equivalence* (LLE), (And), (Or), *right weakening* (RW), and *cautious monotonicity* (CM) (Fig. 3.1). An entailment relation $\mathrel{|\!\approx}$ is *rational iff* it is preferential and satisfies the additional postulate *rational monotonicity* (RM) (Fig. 3.1).

[LM'92] is more practically oriented than their original contribution [KLM'90], in the sense that the authors inspect a full KR scenario, starting from the representation of conditional knowledge and proposing a single entailment relation (extending the input knowledge base), to provide meaningful answers to queries over the represented knowledge.

> "The question asked in the title [...] has no simple answer and has probably no unique answer good for everyone in every situation."
>
> *[LM'92, p. 3]*

Attempts to answer this question include the definition of *preferential entailment*, i.e. to draw those conclusions that are supported by all preferential

1 Loui [Lou'87] distinguishes specificity, which is inherently syntax-dependent, from a syntax-independent form of precedence that he calls superior evidence.

entailment relations as well as the analogous characterisation of *rational entailment*. Lehmann and Magidor illustrate inadequacies, showing that both preferential and rational entailment are hardly more effective than classical semantics, seeing as they remain monotonic. Consequently, they argued that among all rational entailment relations, there would be one allowing to draw the most sensible conclusions.

Formally, Lehmann and Magidor introduce a preference relation $\ll$ over all rational entailment relations. To not overwhelm with the technical details of Def. 5.5 in [LM'92], we reiterate only the intuition that the authors describe in terms of a discussion between two rational agents. Preference of one agent over the other is determined by him attacking a conclusion that his opponent derives and that the opponent is not able to defend. There is no better way to describe this conversation than through the words of Lehmann and Magidor.

> "Suppose two agents, who agree on a common knowledge base, are discussing the respective merits of two rational relations [...] [$\approx_0$] and [...] [$\approx_1$]. A typical attack would be: your relation contains an assertion [...] [$\alpha \mathrel{|}\!\approx_1 \beta$], that mine does not contain (and therefore contains unsupported assertions). A possible defense against such an attack could be: yes, but your relation contains an assertion [...] [$\gamma \mathrel{|}\!\approx_0 \delta$] that mine does not, and you yourself think that γ refers to a situation that is more usual than the one referred to by α. Such a defense must be accepted as valid."

[LM'92, p. 31]

A key aspect in this intuition is the use of the term "more usual", something that ensures precedence of more specific conditional statements over more general ones. Rational Closure is then formalised as a collection of conditional assertions, extending the given KB.

Definition 3.2 (Rational Closure)**.** For a conditional knowledge base $\mathcal{K}$, if there is a rational entailment relation $\mathrel{|}\!\approx$ with $\mathcal{K} \subseteq \mathrel{|}\!\approx$, that is preferable (according to $\ll$) to all other rational relations extending $\mathcal{K}$, then the *Rational Closure* of $\mathcal{K}$ is

$$\mathsf{RC}(\mathcal{K}) = \{\alpha \mathrel{|}\!\sim \beta \mid \alpha \mathrel{|}\!\approx \beta\}. \tag{3.2}$$

The final part in [LM'92] proposes a simple algorithm to compute entailments under RC, relying on a transformation of conditional assertions $\alpha \mathrel{|}\!\approx \beta$ to their material implication $\alpha \to \beta$, or (equivalently) $\neg\alpha \vee \beta$. With this algorithm, the problem of deciding entailments under RC is effectively reduced to classical reasoning in propositional logics.

The simplicity of *materialisation*-based RC has inspired numerous syntactic translations to more expressive formalisms, in particular DLs [CMMN'14;

CMM+'15; CMMV'13; CMVM'13; CS'10; CS'12; CS'13; GGOP'15]. However, such adaptations are dangerous, as features driving a formalism beyond the expressivity of propositional logic need to be accounted for very carefully. Most notably, the need to treat first-order quantification explicitly, is already conveyed in [LM'90; LM'92]. We discuss the drawbacks of such immediate adaptations thoroughly in Section 5.2. It becomes clear that stronger formalisms, in particular those allowing for quantification, need more attention and a specifically tailored semantics, if meaningful entailments are to be produced. We present our solution to this problem for rational reasoning in $\mathcal{EL}_\bot$ with new, model-theoretic semantics in PART III.

Part II

RATIONAL CLOSURE IN DESCRIPTION LOGICS

INTRODUCTION AND OVERVIEW

In Chapter 3, we have given a glimpse on the variety of approaches to non-monotonic reasoning and their adaptations to DLs. The current part continues where Section 3.2.1 left off, investigating the algorithmic characterisation of RC through materialisation in the DL case.

Adaptations of Rational Closure to Description Logics are first and foremost a matter of definition. Many nuances of such a definition have appeared in the literature. In order to provide a DL reasoning mechanism worthy of the name Rational Closure, a close relation to its original notion must be maintained. This relation is typically established by an translation of the KLM postulates to defeasible variants of reasoning problems in DLs, such as subsumption or instance checking. This translation allows to adopt the basic notions of preferential or rational entailment relations (Def. 3.1) to, e.g., defeasible subsumption relations $\sqsubsetsim$. Consequently, defeasible reasoning in DLs has adopted all three aspects of KLM-style research:

- Model-theoretic semantic representations of preferential and rational entailment relations for defeasible subsumption: [BCM+'13; BMV'11b; BV'16; BV'17a; BV'18b; GGOP'07; GGOP'10b; GOGP'09; Var'18]
- Different characterisations of an ideal rational entailment relation: [BCM+'13; BMV'11b; BV'18b; CMMN'14; CMS'18; GG'18; GGOP'12]
- An algorithmic characterisation of RC, relying on an efficient reduction from a defeasible subsumption query (or instance check) to an appropriate classical subsumption query (or instance check): [CMMN'14; CMM+'15; CMMV'13; CMVM'13; CS'10; CS'12; CS'13; GGOP'15]

Most adaptations of all three aspects to DLs in the literature is remarkably close to the original formalisation, perhaps too close, considering their divergence in expressivity. The closest non-algorithmic characterisation of RC was studied by Britz et al. [BMV'11b]. It is essentially a syntactic translation of the preference order over rational entailment relations (Sec. 3.2.1 and [LM'92]), based on the notion of rational subsumption relations characterised through DL postulates, as exemplified in Figure 3.2. Britz et al. [BMV'11b] divert this characterisation to a more intuitive one, relying on the computation of concept ranks (Def. 4.12), based on a concepts (in-)consistency in the presence of defeasible knowledge. Since then, numerous contributions have adopted this concept rank-based characterisation as the definition of RC in DLs [Bon'19; BCM+'13; BMV'11b; BV'18b; CMMN'14; CMS'18; GG'18; GGOP'12]. To extend the chain of equivalent characterisations, algorithmic approaches through materialisation are many

$$\text{(RW)} \qquad \frac{C \mathrel{\tilde{\sqsubseteq}} D \quad D \sqsubseteq E}{C \mathrel{\tilde{\sqsubseteq}} E}$$

$$\text{(CM)} \qquad \frac{(\mathcal{A}, \mathcal{T}, \mathcal{D}) \mathrel{|\!\approx} C\{a\} \quad (\mathcal{A}, \mathcal{T}, \mathcal{D}) \mathrel{|\!\approx} D\{b\}}{(\mathcal{A} \cup \{D(b)\}, \mathcal{T}, \mathcal{D}) \mathrel{|\!\approx} C\{a\}}$$

Figure 3.2: Two examples for the translation of KLM-postulates, adjusted to the inference problems of defeasible subsumption and defeasible instance checking, respectively [CS'10].

times proven to capture RC as defined through concept ranks [CMMN'14; CMM+'15; CMMV'13; CMVM'13; CS'10; CS'12; CS'13; GGOP'15].

We see these proofs of equivalence as an invitation to keep things simple and streamlined in this part of our work. That is, we will move one step further and define Rational Closure directly through materialisation and cover concept ranks only briefly, and model-theoretic characterisations not at all. To incorporate the resulting semantics into our *strength-coverage* identification scheme, we shall assign to entailment relations based on materialisation the coverage `mat`, i.e. for the most basic case of RC, introducing (`rat`,`mat`)-semantics.

The contributions of this part are a mix of reorganising approaches from the literature and supplementing them by the additions that we originally published in [PT'17a; PT'17b]. The structure of this part is essentially twofold.

A FRAMEWORK FOR MATERIALISATION. The materialisation-based characterisation of Rational Closure as well as several proposed extensions of it [CMMN'14; CS'12; CS'13] share so many commonalities, that we can generalise these approaches into a single framework that is able to (a) capture all of those extensions, (b) vary on the underlying DL, in particular covering the special case of sub-boolean materialisation in $\mathcal{EL}_\bot$, and (c) propose a clean setup for future implementations. The generic foundations of this framework are introduced in Chapter 4, accompanied directly by the necessary instantiations to decide defeasible subsumption and instance checks in $\mathcal{ALC}$. The flexibility of this framework is demonstrated in Section 4.4 and Chapter 6. There, we show how to achieve reduction algorithms relying on sub-boolean materialisation to remain entirely in the DL $\mathcal{EL}_\bot$ and how to exchange single modules of this framework to obtain much stronger entailment relations, that have been discussed as extensions of RC.

DISCUSSION OF RC IN DLS. Materialisation allows us to provide simple illustrations for some severe drawbacks of RC. We will discuss the well-known downside of inheritance blocking in Section 5.1, as well as the less studied—but arguably more severe—issue of neglecting quantification in defeasible consequences in Section 5.2. We make a strong case against naive translations of aspects in propositional RC to the DL case, and argue for an explicit treatment of quantification in the light of defeasible knowledge. Equality of the different characterisations clearly shows that the discussed issues are impartial to the underlying characterisation. This discussion provides a natural bridge to the main contribution of our work, capturing expressive extensions of RC with new model-theoretic semantics in PART III.

4 MATERIALISATION

Semantics for defeasible Description Logics based on *materialisation* are reductions to classical reasoning. In the most basic case, we query knowledge bases, containing strict as well as defeasible statements, for defeasible subsumptions $C \sqsubsim D$. Intuitively, we are asking whether the most typical instances of C also satisfy the concept D. Loosely speaking, in materialisation-based approaches, elements of the query concept C are considered more typical, the more DCIs from the given KB they satisfy.[1] At the same time, elements of a query subject such as $\mathit{Cat} \sqcap \neg \mathit{Smart}$ may not be able to satisfy all DCIs from the KB, e.g. $\mathit{Cat} \sqsubsim \mathit{Smart}$. Materialisation-based approaches therefore proceed in two steps.

In the first step, it has to be determined which subset of the given DBox can be satisfied by (elements of) the query subject. While this subset should intuitively be as large as possible, its definition is decisive for the strength of the resulting semantics. In the second step, the classical variant of the respective inference is determined, while imposing the set of DCIs obtained in the first step on the query subject. To understand this imposition, consider the query subject Cat with the defeasible property from above. The elements of Cat that also satisfy $\mathit{Cat} \sqsubsim \mathit{Smart}$ in any interpretation can be syntactically characterised by $\mathit{Cat} \sqcap (\neg \mathit{Cat} \sqcup \mathit{Smart})$, using the material implication of the defeasible statement. Likewise, when arguing about the specific individual daisy, this information can be asserted with $(\neg \mathit{Cat} \sqcup \mathit{Smart})(\mathit{daisy})$, extending the ABox. This technique does not come without its difficulties.

If intended to maintain the complexity of classical reasoning in the underlying DL, the materialisations we illustrated before are not appropriate for $\mathcal{EL}_\bot$ queries. When introducing concepts with full negation and disjunction, reasoning does not trivially remain tractable. A different form of materialisation that is not relying on these constructors is discussed in Section 4.4. Furthermore, determining subsets of the DBox that are consistent with individuals is not always a local operation. When considering role assertions, such as $\mathit{friend}(\mathit{molli}, \mathit{lilly})$, asserting defeasible information on lilly might influence the defeasible information that molli is consistent with. Deciding defeasible subsumption and instance checks is inherently different. Specific definitions are covered separately in Section 4.2 and 4.3.

Nevertheless, the commonalities of materialisation-based approaches outweigh their differences, allowing us to present the core of these procedures

1 In general, two domain elements in an interpretation may be of incomparable typicality w.r.t. the subset order for the sets of DCIs they satisfy. However, in rational reasoning, the satisfaction of subsets of the given DBox is only considered for a predetermined totally ordered chain of subsets, ergo, the notion of being more or less typical is applicable in the rational case.

in terms of an abstract framework. Instantiations of this framework are essentially parametric on a DL and a semantic strength. In Section 4.1 we will introduce the general notions that are central to this framework, as well as some of the basic instantiations that are required for deciding defeasible subsumption entailment under rational materialisation-based semantics $(\mathtt{rat}, \mathtt{mat})$.

4.1 A FRAMEWORK FOR MATERIALISATION

The framework capable of producing materialisation-based reductions for deciding defeasible entailments with different semantic strength and relying on different DL concept constructors, is made up of various functions. These functions are introduced in a generic way, to manifest their duty within the framework, while their specific definitions—hereinafter called *instantiations*—are parametric on a semantic strength $\mathbf{s}$, a DL $\mathcal{L}$, or both. To be more precise, consider the following notion of materialisation functions, generalising the reduction of defeasible queries to classical queries.

Definition 4.1 (Materialisation Function)**.** For a semantic strength $\mathbf{s}$ and a DL $\mathcal{L}$, $\mathtt{Mat}^{\mathcal{L}}_{\mathbf{s}}()$ and $\mathtt{Mat}^{\mathcal{L}}()$ are *materialisation functions*. Both types take as input any $\mathcal{L}$ inference problem $(\mathcal{K}, \alpha)$ and return an $\mathcal{L}$ inference problem $(\mathcal{K}', \alpha')$ where both $\mathcal{K}'$ and α' are strict. $\mathtt{Mat}^{\mathcal{L}}()$ is called *simple*.

The function $\mathtt{Mat}^{\mathcal{L}}_{\mathbf{s}}()$ is finally used to characterise defeasible entailment under $(\mathbf{s}, \mathtt{mat})$-semantics, when instantiated with a specific semantic strength $\mathbf{s}$. $\mathtt{Mat}^{\mathcal{L}}()$ does not rely on a semantic strength, because its purpose is to fix the transformation of DCIs to strict DL statements that are at the core of any materialisation approach. We have already encountered its instantiation for the DL $\mathcal{ALC}$, relying on the usual material implication $\neg C \sqcup D$ for a DCI $C \mathrel{\underset{\sim}{\sqsubset}} D$ [CS'10]. Clearly, when presented with an $\mathcal{EL}_\bot$ inference problem $(\mathcal{K}, \alpha)$, this transformation introduces concept constructors that exceed the $\mathcal{EL}_\bot$ profile, hence the parametrisation of (simple) materialisation functions by a DL $\mathcal{L}$.

Definition 4.2 ($\mathcal{ALC}$ Simple Materialisation)**.** The $\mathcal{ALC}$ material implication of a DCI is $\overline{E \mathrel{\underset{\sim}{\sqsubset}} F} = \neg E \sqcup F$ and for a set of DCIs $\mathcal{E}$, $\overline{\mathcal{E}} = \bigsqcap_{E \mathrel{\underset{\sim}{\sqsubset}} F \in \mathcal{E}} \overline{E \mathrel{\underset{\sim}{\sqsubset}} F}$. The *simple materialisation function* $\mathtt{Mat}^{\mathcal{ALC}}()$ is defined for a knowledge base $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ and any query α as

$$\mathtt{Mat}^{\mathcal{ALC}}(\mathcal{K}, \alpha) = \begin{cases} (\mathcal{K}_{\mathtt{strict}}, \alpha) & \text{, if } \alpha \text{ is classical} \\ (\mathcal{K}_{\mathtt{strict}}, \overline{\mathcal{D}} \sqcap C \sqsubseteq D) & \text{, if } \alpha = C \mathrel{\underset{\sim}{\sqsubset}} D \\ (((\mathcal{A} \cup \{\overline{\mathcal{D}}(a)\}, \mathcal{T}), C(a)) & \text{, if } \alpha = C\{a\} \end{cases}$$

The function $\mathtt{Mat}^{\mathcal{ALC}}()$ can be seen as a direct transformation from a defeasible inference problem to a classical one, stubbornly "enriching" the inference with all defeasible knowledge in $\mathcal{D}$.

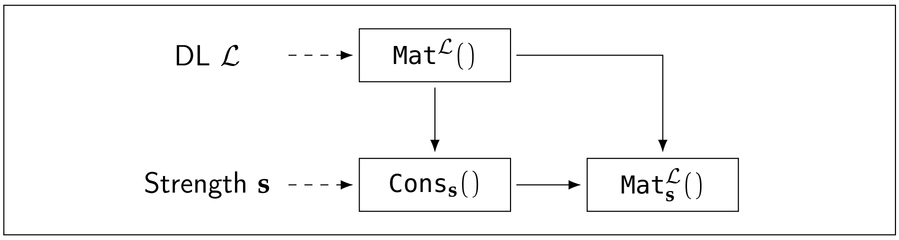

Figure 4.1: Overview and dependencies of the materialisation framework.

Simple materialisation can be considered as the core of this framework, while its companion, equipped with a parameter for semantic strength, is used to define defeasible entailment on the highest level. Formally, defeasible entailment (through materialisation) can be defined in a generic way, relying on $\mathtt{Mat}_{\mathbf{s}}^{\mathcal{L}}()$. Specific types of semantics are then identified through instantiations of $\mathtt{Mat}_{\mathbf{s}}^{\mathcal{L}}()$ on the DL $\mathcal{L}$ and strength $\mathbf{s}$, without needing to adjust the following definition in any way.

Definition 4.3 (Defeasible Entailment under $(\mathbf{s}, \mathtt{mat})$-Semantics)**.** A defeasible inference problem $(\mathcal{K}, \alpha)$ in the DL $\mathcal{L}$ is true under $(\mathbf{s}, \mathtt{mat})$-semantics *iff* $\mathtt{Mat}_{\mathbf{s}}^{\mathcal{L}}(\mathcal{K}, \alpha)$ is true under classical semantics. Formally,

$$\mathcal{K} \models^{(\mathbf{s},\mathtt{mat})} \alpha \textit{ iff } \mathcal{K}' \models \alpha',$$

for $\mathtt{Mat}_{\mathbf{s}}^{\mathcal{L}}(\mathcal{K}, \alpha) = (\mathcal{K}', \alpha')$.

As both functions $\mathtt{Mat}^{\mathcal{L}}()$ and $\mathtt{Mat}_{\mathbf{s}}^{\mathcal{L}}()$ return classical inference problems, we often use the shorthand "$\mathtt{Mat}^{\mathcal{L}}(\mathcal{K}, \alpha)/\mathtt{Mat}_{\mathbf{s}}^{\mathcal{L}}(\mathcal{K}, \alpha)$ is true/false", referring to the entailment of $(\mathcal{K}', \alpha')$ under classical semantics. Clearly, simple materialisation puts us nowhere near rational consequences, as many concepts $\overline{\mathcal{D}} \sqcap C$ (materialising the full DBox of the input KB $\mathcal{K}$) might be unsatisfiable with $\mathcal{K}_{\mathtt{strict}}$. Clearly, enriching a query subject in such a way that it becomes unsatisfiable allows to derive any consequence about this (enriched) subject, rendering such entailments worthless. Determining an appropriate subset of the DBox, as per the *first part* of any materialisation approach, is the responsibility of a *consistent-selection function* $\mathtt{Cons}_{\mathbf{s}}()$, that is instantiated on a semantic strength $\mathbf{s}$.

The dependency graph depicted in Figure 4.1 shows how the three main components of this framework interact. The DL $\mathcal{L}$ and the strength $\mathbf{s}$ are parameters determining the concrete definitions. As seen in Definition 4.2 the return value of a function somewhat depends on its input (e.g. $(\mathcal{K}, \alpha)$) as well as its parameters (e.g. $\mathcal{L}$). Nevertheless, instantiations of functions are purposefully parametrised only by semantic strength and DL as these are considered to be the distinguishing features determining the resulting semantics. To cover the first part of materialisation-based approaches, $\mathtt{Cons}_{\mathbf{s}}()$ will be instantiated with a semantic strength $\mathbf{s}$ and rely on a simple materialisation function to determine consistency of DBox subsets.[2] A materialisation function $\texttt{Mat}^{\mathcal{L}}_{\mathbf{s}}()$, finally determining entailments under materialisation-based semantics, relies in turn on $\texttt{Cons}_{\mathbf{s}}()$ and the corresponding simple materialisation function $\texttt{Mat}^{\mathcal{L}}()$ of the same DL. The generality of these notions helps to emphasize the commonalities among several semantics such as Relevant or Lexicographic Closure. In Chapter 6 we will illustrate how these more expressive semantics can be achieved by different instantiations of only $\texttt{Cons}_{\mathbf{s}}()$, leaving the materialisation functions untouched. In this chapter we focus on instantiations to produce entailments of rational strength, to initially illustrate what instantiations of this framework look like.

As for materialisation functions, we capture in a very generic way, which properties are to be expected of the input and output of a consistent-selection function. This characterisation relies on a notion of exceptionality that is essential in all materialisation-based procedures. It refers to a specific kind of inconsistency, that is introduced subsequently.

Definition 4.4. (Selecting Consistent DCIs) Let $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ be an $\mathcal{L}$ KB, $\chi \in \mathfrak{C}(\mathcal{L}) \cup \mathsf{N_I}$ a concept or individual and $\mathbf{s}$ a semantic strength. $\texttt{Cons}_{\mathbf{s}}(\mathcal{K}, \chi)$ is called a *consistent-selection function iff* $\texttt{Cons}_{\mathbf{s}}(\mathcal{K}, \chi) \subseteq \mathcal{D}$ s.t. χ is *not exceptional* w.r.t. $\texttt{Cons}_{\mathbf{s}}(\mathcal{K}, \chi)$ and $\mathcal{K}_{\texttt{strict}}$. If χ is exceptional already w.r.t. $\emptyset$ and $\mathcal{K}_{\texttt{strict}}$, then by convention $\texttt{Cons}_{\mathbf{s}}(\mathcal{K}, \chi) = \emptyset$.

The core of consistent-selection functions is clearly the *exceptionality* of a concept or individual w.r.t. a set of DCIs and an underlying knowledge base. Intuitively, a concept C is exceptional w.r.t. a set of DCIs $\mathcal{E}$ and a KB $\mathcal{K}$, if no model of $\mathcal{K}$ contains an element satisfying C and $\mathcal{E}$ simultaneously. Satisfiability of DCIs for elements in extensions of C can be reduced to classical entailment using simple materialisation.

Definition 4.5 (Exceptionality). Let $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ be a (not necessarily defeasible) $\mathcal{L}$ KB and $\mathcal{E}$ a set of DCIs.[3]

1. A concept C is *exceptional* w.r.t. $\mathcal{E}$ and $\mathcal{K}$ *iff*

$$\texttt{Mat}^{\mathcal{L}}((\mathcal{A}, \mathcal{T}, \mathcal{E}), \mathrm{C} \mathrel{\underset{\sim}{\sqsubset}} \bot)$$

   is true.

2. A DCI $\mathrm{C} \mathrel{\underset{\sim}{\sqsubset}} \mathrm{D}$ is exceptional w.r.t. $\mathcal{E}$ and $\mathcal{K}$ *iff* C is exceptional w.r.t. $\mathcal{E}$ and $\mathcal{K}$.

3. An individual $a \in \mathsf{N_I}$ is exceptional w.r.t. $\mathcal{E}$ and $\mathcal{K}$ *iff*

$$\texttt{Mat}^{\mathcal{L}}((\mathcal{A}, \mathcal{T}, \mathcal{E}), \bot\{a\})$$

   is true.

2 Technically, $\texttt{Cons}_{\mathbf{s}}()$ relying on $\texttt{Mat}^{\mathcal{L}}()$ makes it rely indirectly on $\mathcal{L}$ as well. However, its instantiations by $\mathbf{s}$ will be shared for any DL $\mathcal{L}$. Therefore, $\mathcal{L}$ does not appear in the notation of $\texttt{Cons}_{\mathbf{s}}()$ to not suggest an immediate dependency on $\mathcal{L}$.

3 We carry the DBox $\mathcal{D}$ in the input KB, to keep this definition as general as possible. For the same reason, exceptionality is determined w.r.t. a set of DCIs $\mathcal{E}$ that might stand in no relation to the DBox $\mathcal{D}$ or just as well be a subset of $\mathcal{D}$ or $\mathcal{D}$ itself.

The set of exceptional DCIs within a KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, is defined as

$$\mathtt{Exc}(\mathcal{K}) = \{C \mathrel{\tilde{\sqsubseteq}} D \in \mathcal{D} \mid C \text{ is exceptional w.r.t. } \mathcal{D} \text{ and } \mathcal{K}\}$$

Using this notation, we have an understanding for the purpose of consistent-selection functions. Ideally, $\mathtt{Cons_s}()$ is applied to the query subject (a concept or individual) and selects a consistent subset of the input DBox that is as large as possible in some sense. In practice, the amount of defeasible information that is selected by these functions directly determines the semantic strength of the reasoning mechanism. To obtain consistent subsets of a DBox in semantics of rational strength requires to predetermine sets of DCIs that are eligible for checking consistency. This procedure follows the original algorithm by Lehmann and Magidor [LM'92] and iteratively produces a totally ordered, chain of decreasing subsets of the given DBox. The benefit of this method is that it requires only a linear number of classical entailment checks. Its drawback on the other hand, is that, in general, it does not come close to any notion of "as large as possible". This issue is predominantly recognised as inheritance blocking (cf. Section 5.1). Attempts to alleviate it, effectively propose alternative instantiations of $\mathtt{Cons_s}()$ (cf. Section 6.1, 6.2), even though this modularisation has never been formalised prior to our work.

Definition 4.6 (Rational Chain). For an $\mathcal{L}$ KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, the *rational chain* of $\mathcal{K}$, is the following set of subsets of the DBox:

$$\mathtt{chain}(\mathcal{K}) = \langle \mathcal{D}_0, \ldots, \mathcal{D}_n, \mathcal{D}_\infty \rangle$$

such that $\mathcal{D}_0 = \mathcal{D}$, $\mathcal{D}_i = \mathtt{Exc}((\mathcal{A}, \mathcal{T}, \mathcal{D}_{i-1}))$ for $i > 0$. The index n is the smallest integer s.t. $\mathtt{Exc}((\mathcal{A}, \mathcal{T}, \mathcal{D}_{n+1})) = \mathcal{D}_{n+1}$ and $\mathcal{D}_\infty = \mathcal{D}_{n+1}$.

Clearly, $\mathtt{chain}(\mathcal{K})$ is finite, if $\mathcal{D}$ is finite. As for $\mathcal{D}_\infty$, there are two possibilities, either $|\mathcal{D}_\infty| > 0$ or it is empty. In case $\mathcal{D}_\infty \neq \emptyset$, all antecedents C of DCIs in $\mathcal{D}_\infty$ are exceptional w.r.t. $\mathcal{D}_\infty$ and $\mathcal{K}$. From the original and derived definition of Rational Closure in [BMV'11b] it becomes clear that the antecedents in $\mathcal{D}_\infty$ are required to support any consequence, in order to produce entailments under RC. More intuitively, there is no rational preference among the DCIs in $\mathcal{D}_\infty$ that would allow removal of some subset of $\mathcal{D}_\infty$ to provide non-exceptionality for the remaining DCIs. As a consequence, the DCIs in $\mathcal{D}_\infty$, and in particular their antecedents, are considered *strictly* unsatisfiable. From a semantic point of view, for every $C \mathrel{\tilde{\sqsubseteq}} D \in \mathcal{D}_\infty$, we must conclude $\mathcal{K}_{\mathtt{strict}} \models C \sqsubseteq \bot$. As argued in [CMMN'14] from an algorithmic point of view, $\mathcal{D}_\infty$ provides strict information that should not be part of the DBox. Formally, a KB $\mathcal{K}$ is considered *well-separated iff* $\mathcal{D}_\infty = \emptyset$ and Britz et al. [BCM+'13] show that every KB $\mathcal{K}$ can be transformed to an equivalent well-separated KB $\mathcal{K}$', simply by replacing every $C \mathrel{\tilde{\sqsubseteq}} D \in \mathcal{D}_\infty$ with $C \sqsubseteq \bot \in \mathcal{T}$.

Remark 4.7. For the remainder of this thesis, we assume all knowledge bases to be *well-separated*, unless explicitly stated otherwise. Effectively, this allows for the obvious assumption that $\mathtt{chain}(\mathcal{K}) = \langle \mathcal{D}_0, \cdots, \mathcal{D}_n \rangle$ ends with $\mathcal{D}_n = \emptyset$ in theorems, definitions and proofs.

For an intuition behind the rational chain and its role in the enrichment of query subjects, consider the following example.

Example 4.8. Let $\mathcal{K} = (\mathcal{T}, \mathcal{D})$ with

$$\begin{aligned}\mathcal{T} &= \{\mathsf{Evil} \sqcap \mathsf{Friendly} \sqsubseteq \bot\}\\ \mathcal{D} &= \{\mathsf{Cat} \mathrel{\sqsubset\!\!\!\sim} \mathsf{Smart},\\ &\qquad \mathsf{Cat} \sqcap \exists \mathsf{friend}.\mathsf{Dog} \mathrel{\sqsubset\!\!\!\sim} \neg\mathsf{Smart} \sqcap \mathsf{Friendly},\\ &\qquad \mathsf{Cat} \sqcap \exists \mathsf{friend}.(\mathsf{Dog} \sqcap \mathsf{Gullible}) \mathrel{\sqsubset\!\!\!\sim} \mathsf{Evil}\}\end{aligned}$$

In this knowledge base, there are essentially two "reasons" for an element (or concept) to be exceptional. No element can ever satisfy $\mathsf{Friendly}$ at the same time as Evil, and obviously $\mathsf{Smart} \sqcap \neg\mathsf{Smart}$ must also be extended to the empty set in any interpretation. Answering the question of what to *rationally* conclude for typical instances of specific classes, such as $\mathsf{Gray\ Cats}$ or $\mathsf{Stupid\ Cats}$, requires consultation of the rational chain to determine a consistent subset of $\mathcal{D}$. We unravel Definition 4.6 step by step, starting with $\mathcal{D}_0 = \mathcal{D}$. Note the strict subsumption hierarchy between the antecedents of our input DCIs:

$$\mathsf{Cat} \sqcap \exists \mathsf{friend}.(\mathsf{Dog} \sqcap \mathsf{Gullible}) \sqsubseteq \mathsf{Cat} \sqcap \exists \mathsf{friend}.\mathsf{Dog} \sqsubseteq \mathsf{Cat} \quad (4.1)$$

This hierarchy is often referred to as an order of specificity, considering concepts lower down in the hierarchy ($\mathsf{Cat} \sqcap \exists \mathsf{friend}.\mathsf{Dog}$) to be more specific than their subsumers (Cat). This specificity ordering implicitly influences the construction of the rational chain, inducing a natural precedence among DCIs.

$$\begin{aligned}\mathcal{D}_1 = \mathsf{Exc}(\mathcal{D}_0) = \{&\mathsf{Cat} \sqcap \exists \mathsf{friend}.\mathsf{Dog} \mathrel{\sqsubset\!\!\!\sim} \neg\mathsf{Smart} \sqcap \mathsf{Friendly},\\ &\mathsf{Cat} \sqcap \exists \mathsf{friend}.(\mathsf{Dog} \sqcap \mathsf{Gullible}) \mathrel{\sqsubset\!\!\!\sim} \mathsf{Evil}\}\end{aligned}$$

Intuitively, more specific concepts are influenced by more DCIs, increasing the chance for exceptionality. For example (for brevity abbreviating the antecedents of DCIs by A_1, A_2, A_3 in the order they are displayed above), in order for an element to satisfy the concept

$$A_2 \sqcap (\neg A_1 \sqcup \mathsf{Smart}) \sqcap (\neg A_2 \sqcup (\neg\mathsf{Smart} \sqcap \mathsf{Friendly}))$$

it is forced to satisfy Smart and $\neg\mathsf{Smart}$ simultaneously (because $A_2 \sqsubseteq A_1$), which is impossible. On the other hand, elements in

$$A_1 \sqcap (\neg A_1 \sqcup \mathsf{Smart}) \sqcap (\neg A_2 \sqcup (\neg\mathsf{Smart} \sqcap \mathsf{Friendly}))$$

are not required to satisfy $\neg\mathsf{Smart} \sqcap \mathsf{Friendly}$. Intuitively, the DCI $A_2 \mathrel{\sqsubset\!\!\!\sim} \neg\mathsf{Smart} \sqcap \mathsf{Friendly}$ does not influence elements of the more general class Cat (A_1).

From the next set of exceptional DCIs, we can also see how the application of defeasible knowledge can be staggered in arbitrary (but finitely deep) levels, according to the hierarchy of antecedents.

$$\mathcal{D}_2 = \mathsf{Exc}(\mathcal{D}_1) = \{\mathsf{Cat} \sqcap \exists \mathsf{friend}.(\mathsf{Dog} \sqcap \mathsf{Gullible}) \mathrel{\sqsubset\!\!\!\sim} \mathsf{Evil}\}$$

As $\mathrm{Cat} \sqcap \exists \mathrm{friend}.(\mathrm{Dog} \sqcap \mathrm{Gullible})$ remains not exceptional with $\mathcal{D}_2$, we must conclude

$$\mathtt{chain}(\mathcal{K}) = \langle \mathcal{D}, \mathcal{D}_1, \mathcal{D}_2, \emptyset \rangle.$$

Being confined to the rational chain for the selection of consistent subsets of the input DBox, leaves really only one candidate from $\mathtt{chain}(\mathcal{K})$ when aspiring for "as much as possible" defeasible information.

Definition 4.9 (Rationally Consistent DCIs). For an $\mathcal{L}$ KB $\mathcal{K}$ with $\mathtt{chain}(\mathcal{K}) = \langle \mathcal{D}_0, \ldots, \mathcal{D}_n \rangle$ and a concept or individual $\chi \in \mathfrak{C}(\mathcal{L}) \cup \mathsf{N}_\mathsf{I}$, the *rational* consistent-selection function $\mathtt{Cons}_{\mathsf{rat}}(\mathcal{K}, \chi)$ is defined as

$$\mathtt{Cons}_{\mathsf{rat}}(\mathcal{K}, \chi) = \mathcal{D}_i,$$

where i is the smallest integer in $\{0, 1, \ldots, n\}$ s.t. χ is not exceptional w.r.t. $\mathcal{D}_i$ and $\mathcal{K}$.

Note that $\mathtt{Cons}_{\mathsf{rat}}()$ does not carry $\mathcal{L}$ in its notation. When used in the definition of e.g. $\mathtt{Mat}^{\mathcal{ALC}}_{\mathsf{rat}}()$, it is only sensible to assume the simple materialisation function $\mathtt{Mat}^{\mathcal{ALC}}()$ to be used within $\mathtt{Cons}_{\mathsf{rat}}()$. Likewise for other DLs $\mathcal{L}$.

Remark 4.10. For a well-separated KB $\mathcal{K}$, i.e. with $\mathcal{D}_n = \emptyset$, concepts or individuals that are only not exceptional w.r.t. $\mathcal{D}_n$ are considered as atypical as possible, their defeasible consequences essentially coincide with their strict consequences. Recall that the assertion of consistent subsets of a DBox to an individual a in the ABox depends on previous such assertions to other individuals in the relational neighbourhood of a. $\mathtt{Cons}_{\mathsf{rat}}()$ (and other instantiations), does not account for these interactions itself. This task is referred to instantiations of $\mathtt{Mat}^{\mathcal{L}}_{\mathsf{s}}()$, which will rely iteratively on "local" consistent-selections of $\mathtt{Cons}_{\mathsf{s}}()$ (cf. Definition 4.16), to fully enrich the ABox with the appropriate materialised assertions.

We have now introduced all necessary ingredients to define materialisation functions $\mathtt{Mat}^{\mathcal{L}}_{\mathsf{s}}()$, including a universal characterisation of defeasible entailment relying on such functions in Definition 4.3. As described in Remark 4.10, the treatment of defeasible subsumption queries and defeasible instance checks differs notably within $\mathtt{Mat}^{\mathcal{L}}_{\mathsf{s}}()$. Therefore, we cover both reasoning problems separately in Section 4.2 and 4.3, leading up to the final instantiations of $\mathtt{Mat}^{\mathcal{ALC}}_{\mathsf{rat}}()$.

4.2 DEFEASIBLE SUBSUMPTION

Instance checking and subsumption are inherently different and unless allowing for the construction of nominal concepts ([Bon'19; BV'17a; BV'17b; CMS'18]), need to be treated individually also in their defeasible variant. For defeasible subsumption inferences α, the definition of the materialisation function(s) $\mathtt{Mat}^{\mathcal{L}}_{\mathsf{s}}(\mathcal{K}, \alpha)$ is straightforward and shared for any semantic strength s and DL $\mathcal{L}$.

Definition 4.11 (Materialisation of Defeasible Subsumption). For an inference problem $(\mathcal{K}, C \mathrel{\tilde{\sqsubseteq}} D)$ in a DL $\mathcal{L}$ and a semantic strength $\mathbf{s}$,

$$\mathtt{Mat}^{\mathcal{L}}_{\mathbf{s}}(\mathcal{K}, C \mathrel{\tilde{\sqsubseteq}} D) = \mathtt{Mat}^{\mathcal{L}}(\mathcal{A}, \mathcal{T}, \mathtt{Cons}_{\mathbf{s}}(\mathcal{K}, C)), C \mathrel{\tilde{\sqsubseteq}} D).$$

$\mathtt{Cons}_{\mathbf{s}}(\mathcal{K}, C)$ delivers a consistent subset of the input DBox, depending on the strength $\mathbf{s}$, that is then transformed by simple materialisation to enrich the resulting classical subsumption query. While direct approaches to materialisation-based reductions ([CMMN'14; CMM+'15; CMMV'13; CMVM'13; CS'10; CS'12; CS'13; GGOP'15]) seem slightly less convoluted, the framework approach shows almost trivially the commonalities among different semantics and lets us switch out or adapt only what is necessary to adapt. Consider the following observations:

1. The simple materialisation function $\mathtt{Mat}^{\mathcal{ALC}}()$ (Def. 4.2), together with the consistent-selection function $\mathtt{Cons}_{\mathtt{rat}}()$ (Def. 4.9) and the above definition of materialisation for defeasible subsumption (Def. 4.11) presents a full instantiation of the materialisation framework. Thus, the universal definition for defeasible entailment (Def. 4.3) immediately provides the procedure to decide defeasible subsumption under $(\mathtt{rat}, \mathtt{mat})$-semantics.

2. $\mathcal{ALC}$ material implications (Def. 4.2) are clearly concepts in $\mathfrak{C}(\mathcal{ALC})$. It is also not hard to see (for the full details, see Chapter 8) that the number of calls to $\mathtt{Mat}^{\mathcal{L}}()$ that are required to compute $\mathtt{chain}(\mathcal{K})$ are linear in the size of the input DBox. This shows that defeasible reasoning under the materialisation-based reduction resides in the same complexity as classical reasoning in $\mathcal{ALC}$ (i.e. ExpTime). The same is not obvious for all DLs, in particular the sub-boolean $\mathcal{EL}_{\bot}$. While $\mathcal{EL}_{\bot}$ queries are technically also $\mathcal{ALC}$ queries, relying on $\mathcal{ALC}$ material implications in the reduction, does not obviously guarantee the reasoning complexity to remain polynomial. It gives us the opportunity to further showcase the simple modular exchange of single aspects of this framework, to achieve different goals. To tailor this reduction specifically to $\mathcal{EL}_{\bot}$, it is only necessary to define an appropriate simple materialisation function $\mathtt{Mat}^{\mathcal{EL}_{\bot}}()$, while Def. 4.11 and others remain untouched.

It remains to place the present entailment of defeasible subsumptions under $(\mathtt{rat}, \mathtt{mat})$-semantics into the context of Rational Closure through concept ranks. As a result of [BCM+'13] this has been successfully accomplished. Over the years, many formalisations of materialisation have appeared in fundamentally different ways.[4] We rely on [BCM+'13] for the most sophisticated version of the following result.

4 Earlier approaches formalised materialisation as a transformation of DCIs to GCIs [BMV'11b], and again others utilised material implications of GCIs alongside DCIs [CMMV'13; CMVM'13; CS'12; CS'13], which clearly results in semantics that are inferentially weaker than classical reasoning. The appropriate transformation is of course the present one, which is also adopted in [BCM+'13; CMMN'14; CMM+'15; CMMV'13; CS'10; CS'13].

An intermediate discovery by Britz et al. [BMV'11b] shows that Rational Closure, as syntactically translated from Lehmann and Magidor [LM'92] in terms of a preference over rational entailment relations, can be equivalently characterised by a ranking function for concepts.

Definition 4.12 (Concept Rank). The *concept rank* of an $\mathcal{ALC}$ concept C in a (not necessarily well-separated) KB $\mathcal{K} = (\mathcal{T}, \mathcal{D})$ with $\mathtt{chain}(\mathcal{K}) = \langle \mathcal{D}_0, \ldots, \mathcal{D}_n, \mathcal{D}_\infty \rangle$ is $r_{\mathcal{K}}(C) = i$, if i is the smallest integer in $\{0, \ldots, n\}$ such that C is not exceptional w.r.t. $\mathcal{D}_i$ and $\mathcal{K}$. If no such i exists, then $r_{\mathcal{K}}(C) = \infty$.[5]

The following characterisation of defeasible subsumption entailments under RC in DLs has been adopted as the definition of RC in the related literature, except for [BMV'11b], where it is presented as a consequence of the more faithful, preference-based definition of RC à la KLM.

Definition 4.13. The *Rational Closure* of a KB $\mathcal{K} = (\mathcal{T}, \mathcal{D})$ is the set $\mathsf{RC}(\mathcal{K}) = \{C \sqsubsetsim D \mid r_{\mathcal{K}}(C) < r_{\mathcal{K}}(C \sqcap \neg D) \text{ or } r_{\mathcal{K}}(C) = \infty\}$.

Most of the literature utilising materialisation [BCM+'13; CMMN'14; CMS'18] proves (or at least claims) that the entailments obtained with the reduction through $\mathtt{Mat}^{\mathcal{ALC}}_{\mathtt{rat}}()$ coincide with the entailments characterised by concept rank (Def. 4.13).

Theorem 4.14 ([BCM+'13]). *For two concepts* C, D *and a KB* $\mathcal{K} = (\mathcal{T}, \mathcal{D})$, $r_{\mathcal{K}}(C) < r_{\mathcal{K}}(C \sqcap \neg D)$ *or* $r_{\mathcal{K}}(C) = \infty$ *iff* $\mathtt{Mat}^{\mathcal{ALC}}_{\mathtt{rat}}(\mathcal{K}, C \sqsubsetsim D)$ *is true.*

The proof of this theorem is relatively straightforward, as concept ranks are determined by the same rational chain as the consistent subsets produced by $\mathtt{Cons}_{\mathtt{rat}}()$. For the formal details, consult Theorem 5 in [BCM+'13] p. 17. Note that in [BCM+'13] RC is defined through concept ranks, formalised with an entailment relation $\models^{\leqslant}_{R}$, and materialisation-based entailment is characterised with the relation $\vdash_{rat}$.

4.3 DEFEASIBLE INSTANCE CHECKING

Defeasible instance checking has been approached with materialisation in [CMVM'13; CS'10; CS'12; CS'13] for different semantic strengths. The fundamental idea of materialisation is transferable from subsumption to instance checking. A query subject (here, an individual) should be assigned some predetermined consistent set of defeasible statements, to derive consequences based on the defeasible part of the KB. In a naive way, this would mean to employ $\mathtt{Cons}_{\mathtt{s}}()$ to determine such consistent knowledge and $\mathtt{Mat}^{\mathcal{L}}()$, extending the ABox appropriately, to reduce the problem to classical instance checking. However, the defeasible instance checking variant of the KLM postulates (e.g. Fig. 3.2) as introduced in [CS'10] expect a sort of coherence among instance checks over distinct individuals. Take

5 Note that the ranking function in [BMV'11b] is defined in a "reversed" manner to Definition 4.12, e.g. where the rank in [BMV'11b] is 0, it will be ∞ here.

Cautious Monotonicity for example (Fig. 3.2) and suppose $C\{a\}$ and $C\{b\}$ are derivable from $(\mathcal{A},\mathcal{T},\mathcal{D})$ while a and b are related in the ABox (say $r(a,b) \in \mathcal{A}$). Extending $\mathcal{A}$ with $D(b)$ may not preserve the consistent subset of $\mathcal{D}$ that is selected for a by $\mathtt{Cons_s}()$ and therefore, not all consequences about (defeasible) membership of a might be preserved. To illustrate such interactions more explicitly, consider the following example.

Example 4.15. For $\mathcal{K} = (\mathcal{A},\mathcal{D})$, let

$$\begin{aligned}\mathcal{A} &= \{A(a), B(b), r(a,b)\}\\ \mathcal{D} &= \{A \mathrel{\sqsubseteq\!\!\!\sim} \forall r.(\neg X),\\ &\qquad B \mathrel{\sqsubseteq\!\!\!\sim} X\}\end{aligned}$$

It is not hard to verify that $\mathtt{Cons_{rat}}(a) = \mathtt{Cons_{rat}}(b) = \mathcal{D}$. Through the naive method, we would conclude $\forall r.(\neg X)\{a\}$ as well as $X\{b\}$, after (separately) extending the ABox with $(\neg A \sqcup \forall r.(\neg X)) \sqcap (\neg B \sqcup X)(a)$, for the former consequence, and with the same assertion for b, to obtain the latter consequence. Applying (CM) as follows

$$\frac{(\mathcal{A},\mathcal{T},\mathcal{D}) \mathrel{|\!\!\approx} \forall r.(\neg X)\{a\} \quad (\mathcal{A},\mathcal{T},\mathcal{D}) \mathrel{|\!\!\approx} X\{b\}}{(\mathcal{A} \cup \{X(b)\},\mathcal{T},\mathcal{D}) \mathrel{|\!\!\approx} \forall r.(\neg X)\{a\}}$$

would require any rational entailment relation to support contradicting conclusions and thus violate the principle of defeasible reasoning.

Several of the adopted postulates refer to instance checks over potentially distinct individuals and can therefore cause problems when determining consistent sets of DCIs for individuals only "locally". To ensure some sort of coherence among defeasible instance checks for different individuals, their consistent sets of DCIs should *all* be considered in the reduction algorithm at the same time. To avoid extending the ABox with contradictory concept assertions, the prominent approach taken in [CMVM'13; CS'10; CS'12; CS'13] is to consider different ABox extensions, much like the default assumption extensions studied by Reiter [Rei'80]. As a result of [CMVM'13], such extensions are uniquely determined by processing individuals in a given order. The distinct entailment relations relying on different ABox extensions are therefore identified by this order as an *additional parameter*. Giordano et al. [GGOP'15] consider entailments of the ABox as those that can be derived from all ABox extensions, as in *sceptical* semantics for Reiter's default logic. For a more thorough discussion of the different options, including the sceptical approach, consult Casini et al. [CMVM'13]. They prove that the sceptical entailment relation does not satisfy the rationality postulates in general, motivating us to adopt the parametrised approach of classifying a multitude of rational entailment relations for defeasible instance checking.

As the consistent-selection function $\mathtt{Cons_{rat}}()$ operates locally, we refer the iterative extension of the input ABox to the definition of $\mathtt{Mat}_{\mathtt{s}}^{\mathcal{L}}()$, calling $\mathtt{Cons_{rat}}()$ multiple times in the process. As opposed to the sequence of individuals in the given ABox, used for parametrisation of entailment relations by [CS'10], we free this notation from being syntactically linked to the input KB, by assuming a total preference order $\prec$ over all individuals N_I. The order $\prec$ over N_I then induces a sequence $\mathtt{seq}_\prec(\mathbb{O}) = \langle a_1, \dots, a_n \rangle$ for every finite subset (e.g. $\mathbb{O} = sig_I(\mathcal{A})$) of N_I in the obvious way. Additionally, we utilise our uniform notation for different semantics to capture defeasible instance checking entailment under the given order $\prec$, by parametrising the strength identifier as $\mathbf{s}_\prec$. The definition of $\mathtt{Mat}^{\mathcal{L}}_{\mathbf{s}_\prec}()$ is extended to accept defeasible instance queries, but remains generic on the DL $\mathcal{L}$ and the strength $\mathbf{s}$. To allow the following definition to be reused for stronger instantiations (Sec. 6.1, 6.2) and different DLs, we use an initialisation function $\mathtt{init}^{\mathcal{L}}_{\mathbf{s}_\prec}()$. It provides the initial ABox, TBox and DBox to start the iterative processing of individuals from (cf. Rem. 4.17).

Definition 4.16 (Materialisation of Defeasible Instance Checking)**.** For an $\mathcal{L}$ KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, a semantic strength $\mathbf{s}$, a total order $\prec$ over N_I such that $\mathtt{seq}_\prec(sig_I(\mathcal{A})) = \langle a_1, a_2, \dots, a_n \rangle$ and an initialisation function $\mathtt{init}^{\mathcal{L}}_{\mathbf{s}_\prec}()$, let

$$(\mathcal{A}_0, \mathcal{T}_0, \mathcal{D}_0) = \mathtt{init}^{\mathcal{L}}_{\mathbf{s}_\prec}(\mathcal{K}), \tag{4.2}$$

$$\mathcal{D}_{a_i} = \mathtt{Cons}_{\mathbf{s}}((\mathcal{A}_{i-1}, \mathcal{T}_0, \mathcal{D}_0), a_i), \text{ and} \tag{4.3}$$

$$(\mathcal{A}_i, \mathcal{X}) = \mathtt{KB}(\mathtt{Mat}^{\mathcal{L}}((\mathcal{A}_{i-1}, \mathcal{T}_0, \mathcal{D}_{a_i}), \top\{a_i\})) \tag{4.4}$$

for $1 \leqslant i \leqslant n$. The final ABox $\mathcal{A}_n$ will be denoted with $\mathcal{A}_{\mathbf{s}_\prec}$, likewise, $\mathcal{K}_{\mathbf{s}_\prec} = (\mathcal{A}_{\mathbf{s}_\prec}, \mathcal{T}_0)$. For simple access to each of the DBox subsets selected for the individuals $a_1, \dots, a_n$, let

$$\mathtt{ext}(\mathcal{K}_{\mathbf{s}_\prec}, a_i) = \mathcal{D}_{a_i}.$$

For a defeasible instance check $(\mathcal{K}, C\{a\})$, its materialisation w.r.t. $\mathcal{L}$ and $\mathbf{s}_\prec$ is

$$\mathtt{Mat}^{\mathcal{L}}_{\mathbf{s}_\prec}(\mathcal{K}, C\{a\}) = (\mathcal{K}_{\mathbf{s}_\prec}, C(a))$$

For rational strength and the DL $\mathcal{ALC}$, the iteration in Def. 4.16 is simply initialised with $\mathtt{init}^{\mathcal{ALC}}_{\mathtt{rat}_\prec}(\mathcal{K}) = \mathcal{K}$.

Remark 4.17. Note the following remarks on Definition 4.16, in particular explaining the chained functions applied in (4.4).

- The initialisation function $\mathtt{init}^{\mathcal{L}}_{\mathbf{s}_\prec}()$ is a formality and allows later on (Sec. 6.1, 6.2) to build more refined ABox extensions, that are initially based on e.g. the rational ABox extension. In anticipation of such inferentially stronger semantics, consider the initialisation $\mathtt{init}^{\mathcal{ALC}}_{\mathtt{rel}_\prec}((\mathcal{A}, \mathcal{T}, \mathcal{D})) = (\mathcal{A}_{\mathtt{rat}_\prec}, \mathcal{T}, \mathcal{D})$ for ABox extensions of relevant strength. The technical motivation for this initialisation is given in Section 6.1.

- In (4.3), it is important to use the initial input TBox $\mathcal{T}_0$ and DBox $\mathcal{D}_0$ in $\mathtt{Cons_s}()$, because the rational chain needs to be established always w.r.t. the initial KB (it does not depend on the ABox). In practice, the rational chain is of course computed only once, prior to this iteration.

- The sole purpose of (4.4) is to iteratively produce $\mathcal{A}_i$ from the previous ABox $\mathcal{A}_{i-1}$. First of all, materialisation functions produce inference problems, hence the projection ($\mathtt{KB}()$) of such a pair to its KB component. On the same note, the inference $\top\{a_i\}$ in (4.4) can use any arbitrary concept (and individual for that matter), as the query does not influence the extended ABox in this type of simple materialisation, as you can see from Def. 4.2. We rely on the simple materialisation function because the instantiation of $\mathtt{Mat}^{\mathcal{L}}()$ determines the concept constructors used in the assertions extending $\mathcal{A}_{i-1}$.

- Later, we are considering relevant contexts $\mathbb{C}, \mathbb{O}$ (Def. 2.13) in general (with $sig_I(\mathcal{A}) \subseteq \mathbb{O}$), possibly including individuals $a \notin sig_I(\mathcal{A})$. Def. 4.16 extends easily to this case, by considering the sequence of individuals $\mathtt{seq}_{\prec}(\mathbb{O})$. As long as $\mathbb{O}$ is finite, the resulting ABox extension remains finite as well. It is not hard to see that individuals in $\mathbb{O} \setminus sig_I(\mathcal{A})$ are consistent with $\mathtt{Cons_s}(\mathcal{K}, \top)$, because $\mathcal{K}$ holds no explicit information about them. Nevertheless, when DCIs of the form $\top \sqsubsetsim \cdots$ are part of the DBox, typical information might still be inferred for those individuals.

For an illustration of this iterative ABox extension, consider the following (less naive) approach to reason over the ABox of Example 4.15.

Example 4.18. Let $\mathcal{K} = (\mathcal{A}, \mathcal{D})$ as in Example 4.15 and recall that if a is a typical member of A then b is not allowed to belong to X and vice versa. First, for $a \prec b$ clearly $\mathcal{A}_0 = \mathcal{A}$ and $\mathcal{A}_1 = \mathcal{A} \cup \{\overline{\mathcal{D}}(a)\}$, because a is processed first and we know $\mathtt{Cons_{rat}}(\mathcal{K}, a) = \mathcal{D}$ already from Example 4.15. Now, when extending $\mathcal{A}_1$ by processing b, the outcome of $\mathtt{Cons_{rat}}((\mathcal{A}_1, \mathcal{D}), b)$ is different, because $(\neg B \sqcup X)(b)$ is no longer consistent with $\mathcal{A}_1$. In particular, $\mathcal{A}_1 \models (\neg X)(b)$. From $\mathtt{chain}(\mathcal{K}) = \langle \mathcal{D}, \emptyset \rangle$ it follows that $\mathcal{A}_2 = \mathcal{A}_1$ and thus $\mathcal{K}_{\mathtt{rat}_\prec} = \mathcal{A} \cup \{\overline{\mathcal{D}}(a)\}$. Thus, instantiating Definition 4.3 with $\mathtt{Mat}^{\mathcal{ALC}}_{\mathtt{rat}_\prec}()$ allows to conclude $(\neg X)(b)$ under $(\mathtt{rat}_\prec, \mathtt{mat})$-semantics. In contrast, let $\prec'$ be a different order over N_I, with $b \prec' a$. Using the same arguments and the same conflict between the locally consistent extensions for a and b, it is easy to see that $\mathcal{K}_{\mathtt{rat}_{\prec'}} = \mathcal{A} \cup \{\overline{\mathcal{D}}(b)\}$. Thus, we derive $X(b)$ under $(\mathtt{rat}_{\prec'}, \mathtt{mat})$-semantics.

Example 4.18 shows clearly how entailments diverge based on the given order over individuals. This total order over N_I, can be seen to classify a range of entailment relations that are rational if they comply with any such order, providing something very similar to a representation result. From a theoretical point of view, this classification is entirely valid and the satisfaction of ABox postulates (see [CMVM'13]) provides some kind of guarantee for the behaviour of the individual entailment relations that are characterised. However, from a practical point of view, it is not immediately clear how to select appropriate orders. A number of different approaches (some aligned with Reiter's treatment of default extensions) are applicable and should be argued for on a case by case basis, depending on the application domain. As this is a theoretical contribution, we refer the interested reader to [CMVM'13] for further investigations, such as the sceptical approach, the analytical approach—specifying ABoxes for which all orders over individuals produce the same entailments—and ways to determine a minimal set of pairs of individuals in the ABox that have to be explicitly ordered by the knowledge engineer.

We close this section by unifying the definition of Rational Closure to include defeasible subsumption and instance checks simultaneously. Theorem 4.14 shows that RC (as in Def. 4.13) might as well be defined through materialisation, i.e. relying on $(\mathtt{rat}, \mathtt{mat})$-semantics directly. The difference between defeasible subsumption and instance checking is that different orders over individuals create different Rational Closures of the knowledge base. At the same time, deciding entailment of defeasible subsumption (Def. 4.11) is impartial to such an order, allowing the following definition.

Definition 4.19 (Materialisation-Based Rational Closure)**.** For an $\mathcal{L}$ KB $\mathcal{K}$, a defeasible inference α and a total order $\prec$ on N_I, the Rational Closure of $\mathcal{K}$ contains $\alpha \in \mathsf{RC}_\prec(\mathcal{K})$ *iff* $\mathcal{K} \models^{(\mathtt{rat}_\prec, \mathtt{mat})} \alpha$ *iff* $\mathtt{Mat}^{\mathcal{L}}_{\mathtt{rat}_\prec}(\mathcal{K}, \alpha)$ is true under classical semantics.

As an immediate corollary of Theorem 4.14 it holds that a defeasible subsumption is contained in the RC of a KB as per Definition 4.13 *iff* it belongs to $\mathsf{RC}_\prec(\mathcal{K})$ for any order $\prec$ over N_I, i.e. $\mathsf{RC}(\mathcal{K}) \subseteq \mathsf{RC}_\prec(\mathcal{K})$. We include here a formal characterisation of Rational Closure as an extension of the input knowledge base mainly for the sake of completeness. In the remainder we shall rely on $(\mathtt{rat}_\prec, \mathtt{mat})$-semantics ($\models^{(\mathtt{rat}_\prec, \mathtt{mat})}$) as formally defined with Definition 4.3, using the appropriate instantiations that have been presented in this and previous sections. Note that when informally discussing semantics characterised as a pair $(\mathbf{s}, \mathbf{c})$, we will often omit the parameter $\prec$, without discriminating against defeasible instance checking in the argument.

4.4 MATERIALISATION IN $\mathcal{EL}_\bot$

Formalising material implications in $\mathcal{EL}_\bot$ without exceeding its tractable reasoning complexity is not as straightforward as in $\mathcal{ALC}$, because we cannot rely on disjunction and negation. However, examining what is essentially the difference of a concept C and a concept $(\neg E \sqcup F) \sqcap C$, reveals that there is an equivalent variant of materialisation, relying only on $\mathcal{EL}_\bot$ concept constructors. In an interpretation $\mathcal{I}$, $((\neg E \sqcup F) \sqcap C)^{\mathcal{I}}$ refers to the subset of $C^{\mathcal{I}}$, such that every $d \in ((\neg E \sqcup F) \sqcap C)^{\mathcal{I}}$

$$\text{either satisfies } F \text{ or does not satisfy } E. \tag{4.5}$$

Adding $E \sqsubseteq F$ as a GCI to the KB would have a similar effect, where surely all elements in $C^{\mathcal{I}}$ (in models of the extended KB) comply with Property 4.5. As a matter of fact, this version of materialisation has been adopted in an earlier approach by Britz et al. [BMV'11b]. However, including a GCI $E \sqsubseteq F$ forces *all* domain elements in a model of the KB to satisfy (4.5), instead of only typical members of the concept in question. For a simple demonstration that this kind of materialisation is too strong, consider the $\mathcal{EL}_\perp$ GCIs $A \sqsubseteq \exists r.A$, $\exists r.B \sqsubseteq \perp$ and the DCI $A \mathrel{\underset{\sim}{\sqsubset}} B$. If this DCI is transformed to a strict GCI, extending the KB, we would have to conclude (strictly) $A \sqsubseteq \perp$. However, typical elements of A (those that should satisfy B) need not be related via r to other *typical* elements of A. The distinction between top-level concepts and quantified concepts is lost with this type of materialisation.

To restrict the elements on which to apply such a new GCI, we introduce a *fresh concept name* $A_{E \mathrel{\underset{\sim}{\sqsubset}} F}$ and use it within the GCI $A_{E \mathrel{\underset{\sim}{\sqsubset}} F} \sqcap E \sqsubseteq F$, to individually restrict the elements satisfying its left-hand side. Now, in all models of the (new) KB, only (but all) elements in $A_{E \mathrel{\underset{\sim}{\sqsubset}} F}$ will adhere to (4.5). Thus, to refer to the subclass of C for which (4.5) is true, we consider the query concept $A_{E \mathrel{\underset{\sim}{\sqsubset}} F} \sqcap C$. If E, F, C and the original KB are in $\mathcal{EL}_\perp$, so is the modified KB and the modified query concept. For individuals, the same argument is applied, with the difference of extending the ABox with assertions of the form $A_{E \mathrel{\underset{\sim}{\sqsubset}} F}(a)$, rather than $(\neg E \sqcup F)(a)$.

Definition 4.20 ($\mathcal{EL}_\perp$ Simple Materialisation)**.** Let $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ be a KB and $N_C^{aux} \subseteq N_C$ a set of auxiliary concept names s.t. $sig(\mathcal{K}) \cap N_C^{aux} = \emptyset$. Let

- $\widehat{E \mathrel{\underset{\sim}{\sqsubset}} F} = A_{E \mathrel{\underset{\sim}{\sqsubset}} F}$ with $A_{E \mathrel{\underset{\sim}{\sqsubset}} F} \in N_C^{aux}$ ($\mathcal{EL}$ material implication),
- $\widehat{\mathcal{D}} = \bigsqcap_{E \mathrel{\underset{\sim}{\sqsubset}} F \in \mathcal{D}} \widehat{E \mathrel{\underset{\sim}{\sqsubset}} F}$, and
- $\mathcal{T}^{\mathcal{D}} = \mathcal{T} \cup \{\widehat{E \mathrel{\underset{\sim}{\sqsubset}} F} \sqcap E \sqsubseteq F \mid E \mathrel{\underset{\sim}{\sqsubset}} F \in \mathcal{D}\}$ (TBox extension).

The simple materialisation function $\mathtt{Mat}^{\mathcal{EL}_\perp}()$ is defined for an $\mathcal{EL}_\perp$ KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ and an $\mathcal{EL}_\perp$ query α as follows

$$\mathtt{Mat}^{\mathcal{EL}_\perp}(\mathcal{K}, \alpha) = \begin{cases} (\mathcal{K}_{\mathtt{strict}}, \alpha) & \text{, if } \alpha \text{ is classical} \\ ((\mathcal{A}, \mathcal{T}^{\mathcal{D}}), \widehat{\mathcal{D}} \sqcap C \sqsubseteq D) & \text{, if } \alpha = C \mathrel{\underset{\sim}{\sqsubset}} D \\ ((\mathcal{A} \cup \{\widehat{\mathcal{D}}(a)\}, \mathcal{T}^{\mathcal{D}}), C(a)) & \text{, if } \alpha = C\{a\} \end{cases}$$

Clearly, for an $\mathcal{EL}_\perp$ KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, $\mathcal{T}^{\mathcal{D}}$ is an $\mathcal{EL}_\perp$ TBox. Intuitively, because $\mathcal{T}^{\mathcal{D}}$ contains fresh concept names only on the left-hand side of GCIs, none of the new GCIs will influence the consequences of a concept with $\mathrm{sig}_C(C) \cap N_C^{aux} = \emptyset$. Formally, $(\mathcal{A}, \mathcal{T}^{\mathcal{D}})$ is a *conservative extension* of $(\mathcal{A}, \mathcal{T})$ (for an ABox $\mathcal{A}$, TBox $\mathcal{T}$, DBox $\mathcal{D}$ in $\mathcal{EL}_\bot$), which means the two KBs are inseparable in terms of their entailments over the original signature [LW'10].

Proposition 4.21. *For an $\mathcal{EL}_\bot$ KB $(\mathcal{A}, \mathcal{T})$, a set of DCIs $\mathcal{E}$ and a classical subsumption or instance check α with $\mathrm{sig}_C(\alpha) \cap N_C^{aux} = \emptyset$, $(\mathcal{A}, \mathcal{T}^{\mathcal{E}}) \models \alpha$ iff $(\mathcal{A}, \mathcal{T}) \models \alpha$.*

Proof. The *if*-direction is trivial, because $\mathtt{Mod}((\mathcal{A}, \mathcal{T}^{\mathcal{E}})) \subseteq \mathtt{Mod}((\mathcal{A}, \mathcal{T}))$ and for the *only-if*-direction, it suffices to show that from a counterexample for $(\mathcal{A}, \mathcal{T}) \models \alpha$, a counterexample for $(\mathcal{A}, \mathcal{T}^{\mathcal{E}}) \models \alpha$ can be derived. Such a model of $(\mathcal{A}, \mathcal{T}^{\mathcal{E}})$ can be trivially obtained by extending all $A_{E \mathrel{\underset{\sim}{\sqsubset}} F} \in N_C^{aux}$ with the empty set, trivially satisfying $(\mathcal{A}, \mathcal{T}^{\mathcal{E}})$, without changing any entailments w.r.t. the original signature. □

Thanks to the generality of the framework established in Section 4.1, we almost immediately obtain a full instantiation of $\mathtt{Mat}_{\mathtt{rat}}^{\mathcal{EL}_\bot}()$. For treating defeasible instance checks, we specify the initial KB for Def. 4.16 as $\mathtt{init}_{\mathtt{rat}}^{\mathcal{EL}_\bot}((\mathcal{A}, \mathcal{T}, \mathcal{D})) = (\mathcal{A}, \mathcal{T}^{\mathcal{D}}, \mathcal{D})$. To show that the $\mathcal{EL}_\bot$ materialisation is working as intended, we need to prove that the result of $\mathtt{Mat}_{\mathtt{rat}}^{\mathcal{ALC}}()$ is true *if and only if* the inference problem produced by $\mathtt{Mat}_{\mathtt{rat}}^{\mathcal{EL}_\bot}()$ is true (both under classical semantics) for all $\mathcal{EL}_\bot$ inference problems. Because most of the involved definitions (Def. 4.3, 4.5, 4.6, 4.9, 4.11, 4.16 and 4.19) are identical for both $\mathcal{ALC}$ and $\mathcal{EL}_\bot$ (modulo using the respective simple materialisation function), it boils down to showing that the functions $\mathtt{Mat}^{\mathcal{ALC}}()$ and $\mathtt{Mat}^{\mathcal{EL}_\bot}()$ produce equivalent inference queries in terms of classical semantics, when presented with an $\mathcal{EL}_\bot$ defeasible inference problem. We prove this separately for subsumption and instance checking.

Theorem 4.22. *For an $\mathcal{EL}_\bot$ KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ and an $\mathcal{EL}_\bot$ defeasible subsumption $C \mathrel{\underset{\sim}{\sqsubset}} D$, it holds that*

$$\mathtt{Mat}^{\mathcal{ALC}}(\mathcal{K}, C \mathrel{\underset{\sim}{\sqsubset}} D) \text{ iff } \mathtt{Mat}^{\mathcal{EL}_\bot}(\mathcal{K}, C \mathrel{\underset{\sim}{\sqsubset}} D)$$

under classical semantics.

Proof. Assuming $(\mathcal{A}, \mathcal{T})$ to be a consistent classical knowledge base (the alternative proves this theorem trivially), any subsumption $X \sqsubseteq Y$ is classically entailed by $(\mathcal{A}, \mathcal{T})$ *iff* it is classically entailed by $\mathcal{T}$. Therefore, we assume w.l.o.g. $\mathcal{A} = \emptyset$, i.e. $\mathcal{K} = (\mathcal{T}, \mathcal{D})$. Looking at the inference queries obtained from $\mathtt{Mat}^{\mathcal{ALC}}(\mathcal{K}, C \mathrel{\underset{\sim}{\sqsubset}} D)$ and $\mathtt{Mat}^{\mathcal{EL}_\bot}(\mathcal{K}, C \mathrel{\underset{\sim}{\sqsubset}} D)$ (Def. 4.2 and 4.20), we effectively need to prove

$$\mathcal{T} \models \overline{\mathcal{D}} \sqcap C \sqsubseteq D \textit{ iff } \mathcal{T}^{\mathcal{D}} \models \widehat{\mathcal{D}} \sqcap C \sqsubseteq D. \tag{4.6}$$

We prove both direction by contraposition.

Recall for $\mathcal{D} = \{E_1 \mathrel{\underset{\sim}{\sqsubset}} F_1, \ldots, E_n \mathrel{\underset{\sim}{\sqsubset}} F_n\}$ that $\widehat{\mathcal{D}} = \bigsqcap_{1 \leqslant i \leqslant n} A_{E_i \mathrel{\underset{\sim}{\sqsubset}} F_i}$ ($A_{E_i \mathrel{\underset{\sim}{\sqsubset}} F_i} \in N_C^{aux}$) and $\overline{\mathcal{D}} = \bigsqcap_{1 \leqslant i \leqslant n} (\neg E_i \sqcup F_i)$. First of all, note that $\mathcal{I} \models \mathcal{T}^{\mathcal{D}}$ implies

1. $\mathcal{I} \models \mathcal{T}$, because $\mathcal{T} \subseteq \mathcal{T}^{\mathcal{D}}$ and
2. $\widehat{\mathcal{D}}^{\mathcal{I}} \subseteq \overline{\mathcal{D}}^{\mathcal{I}}$.

The latter holds, because $d \in A^{\mathcal{I}}_{E_i \sqsubsetsim F_i}$ and $\mathcal{I} \models A_{E_i \sqsubsetsim F_i} \sqcap E_i \sqsubseteq F_i$ imply $d \in (\neg E_i \sqcup F_i)^{\mathcal{I}}$ (for $1 \leqslant i \leqslant n$). These facts immediately prove the contraposition of $\mathcal{T}^{\mathcal{D}} \not\models \widehat{\mathcal{D}} \sqcap C \sqsubseteq D$ implying $\mathcal{T} \not\models \overline{\mathcal{D}} \sqcap C \sqsubseteq D$, because a counterexample (model of $\mathcal{T}^{\mathcal{D}}$) for $\widehat{\mathcal{D}} \sqcap C \sqsubseteq D$ is also a counterexample (model of $\mathcal{T}$) for $\overline{\mathcal{D}} \sqcap C \sqsubseteq D$.

For the contraposition of the other direction, assume there is a model $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ of $\mathcal{T}$, s.t. $\exists d \in (\overline{\mathcal{D}} \sqcap C)^{\mathcal{I}} \setminus D^{\mathcal{I}}$. Let $\mathcal{J} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{J}})$ with $\chi^{\mathcal{J}} = \chi^{\mathcal{I}}$ for all $\chi \in sig(\mathcal{K}, C, D)$ and $A^{\mathcal{J}}_{E_i \sqsubsetsim F_i} = (\neg E_i \sqcup F_i)^{\mathcal{I}}$ for all $1 \leqslant i \leqslant n$. It holds that

- $X^{\mathcal{J}} = X^{\mathcal{I}}$ (for $X \in \mathfrak{C}(\mathcal{ALC})$ with $sig(X) \subseteq sig(\mathcal{K}, C, D)$), i.e. $\mathcal{J} \models \mathcal{T}$,
- $\widehat{\mathcal{D}}^{\mathcal{J}} = \overline{\mathcal{D}}^{\mathcal{I}}$ (by definition of $\mathcal{J}$), and
- $\mathcal{J} \models A_{E_i \sqsubsetsim F_i} \sqcap E_i \sqsubseteq F_i$ for all $1 \leqslant i \leqslant n$.

Therefore, $\mathcal{J}$ is a counterexample for $\mathcal{T}^{\mathcal{D}} \models \widehat{\mathcal{D}} \sqcap C \sqsubseteq D$. □

Theorem 4.22 directly implies that for an $\mathcal{EL}_\bot$ KB $\mathcal{K}$, query subject χ and set of DCIs $\mathcal{E}$, χ is exceptional w.r.t. $\mathcal{E}$ and $\mathcal{K}$ (Def. 4.5) using $\mathtt{Mat}^{\mathcal{EL}_\bot}()$ *iff* it is when using $\mathtt{Mat}^{\mathcal{ALC}}()$. Therefore, the rational chain as well as $\mathtt{Cons}_\mathbf{s}()$ are identical for both simple materialisation functions.[6] We continue with a very similar proof for the case of defeasible instance checking, before concluding this section with the final result as a simple consequence of the preceding theorems. The case of defeasible instance checking is covered more generically, because $\mathtt{Mat}^{\mathcal{L}}_{\mathbf{s}_\prec}(\mathcal{K}, C\{a\})$ involves the addition of multiple $\mathcal{L}$ material implication concept assertions for different individuals in the ABox. The idea is to prove the equivalence of entailments for two versions of an arbitrary ABox extension, relying on the two types of material implications, respectively.

Theorem 4.23. *For an $\mathcal{EL}_\bot$ KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, an $\mathcal{EL}_\bot$ defeasible instance check $C\{a\}$, sets of $\mathcal{EL}_\bot$ DCIs $\mathcal{D}_1, \ldots, \mathcal{D}_k \subseteq \mathcal{D}$ and $a_1, \ldots, a_k \in sig_I(\mathcal{A})$, the following are equivalent*

1. $(\mathcal{A} \cup \{\widehat{\mathcal{D}_j}(a_j) \mid 1 \leqslant j \leqslant k\}, \mathcal{T}^{\mathcal{D}}) \models C(a)$
2. $(\mathcal{A} \cup \{\overline{\mathcal{D}_j}(a_j) \mid 1 \leqslant j \leqslant k\}, \mathcal{T}) \models C(a)$.

Proof. We assume $(\mathcal{A}, \mathcal{T})$ to be consistent, as the alternative trivially satisfies the claim. Let $\mathcal{D} = \{E_1 \sqsubsetsim F_1, \cdots, E_m \sqsubsetsim F_m\}$ and assume w.l.o.g. that all $\mathcal{D}_j$ are non-empty, hence $|\mathcal{D}_j| = n_j$ ($1 \leqslant n_j \leqslant m$) for all $1 \leqslant j \leqslant k$. Then, the most general presentation of all $\mathcal{D}_j$ ($1 \leqslant j \leqslant k$) is

$$\mathcal{D}_j = \{E_{j_1} \sqsubsetsim F_{j_1}, \cdots, E_{j_{n_j}} \sqsubsetsim F_{j_{n_j}}\},$$

with $1 \leqslant j_i \leqslant m$ for all $1 \leqslant i \leqslant n_j$. For a shorthand, let

6 This further supports the notation of $\mathtt{Cons}_\mathbf{s}()$ to be independent of the DL $\mathcal{L}$.

- $\widehat{\mathcal{A}} = \{\widehat{\mathcal{D}_j}(a_j) \mid 1 \leqslant j \leqslant k\}$ and
- $\overline{\mathcal{A}} = \{\overline{\mathcal{D}_j}(a_j) \mid 1 \leqslant j \leqslant k\}$.

We show both directions of the claim by contraposition.

[$1 \Longleftarrow 2$] Assume $\mathcal{I} \models (\mathcal{A} \cup \widehat{\mathcal{A}}, \mathcal{T}^{\mathcal{D}})$ with $a^{\mathcal{I}} \notin C^{\mathcal{I}}$ and recall $\mathcal{T} \subseteq \mathcal{T}^{\mathcal{D}}$. $a_j^{\mathcal{I}} \in A^{\mathcal{I}}_{E_{j_i} \lesssim F_{j_i}}$ and $\mathcal{I} \models A_{E_{j_i} \lesssim F_{j_i}} \sqcap E_{j_i} \sqsubseteq F_{j_i}$ directly implies $a_j^{\mathcal{I}} \in (\neg E_{j_i} \sqcup F_{j_i})^{\mathcal{I}}$ for all $1 \leqslant j \leqslant k$ and $1 \leqslant i \leqslant n_j$. Therefore $\mathcal{I} \models (\mathcal{A} \cup \overline{\mathcal{A}}, \mathcal{T})$ and provides a counterexample for $(\mathcal{A} \cup \overline{\mathcal{A}}, \mathcal{T}) \models C(a)$.

[$1 \Longrightarrow 2$] Assume $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ is a model of $(\mathcal{A} \cup \overline{\mathcal{A}}, \mathcal{T})$ with $a^{\mathcal{I}} \notin C^{\mathcal{I}}$. Construct $\mathcal{J} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{J}})$ with $\chi^{\mathcal{J}} = \chi^{\mathcal{I}}$ for all $\chi \in sig(\mathcal{K}, C\{a\})$ (including individuals) and for all $E \lesssim F \in \mathcal{D}$, let $A^{\mathcal{J}}_{E \lesssim F} = \{a_j^{\mathcal{I}} \mid E \lesssim F \in \mathcal{D}_j\}$. It follows that

- $X^{\mathcal{J}} = X^{\mathcal{I}}$ (for $X \in \mathfrak{C}(\mathcal{ALC})$ with $sig(X) \subseteq sig(\mathcal{K}, C\{a\})$), i.e. $\mathcal{J} \models (\mathcal{A}, \mathcal{T})$,
- $\mathcal{J} \models \widehat{\mathcal{A}}$ (by definition of $\mathcal{J}$), and
- because $a_j^{\mathcal{J}} = a_j^{\mathcal{I}} \in (\neg E \sqcup F)^{\mathcal{I}} = (\neg E \sqcup F)^{\mathcal{J}}$ for all $E \lesssim F \in \mathcal{D}_j$ and $1 \leqslant j \leqslant k$, $\mathcal{J}$ also satisfies all $A_{E \lesssim F} \sqcap E \sqsubseteq F$ for $E \lesssim F \in \mathcal{D}$.

Thus $\mathcal{J} \models (\mathcal{A} \cup \widehat{\mathcal{A}}, \mathcal{T}^{\mathcal{D}})$ and provides a counterexample for $(\mathcal{A} \cup \widehat{\mathcal{A}}, \mathcal{T}^{\mathcal{D}}) \models C(a)$. □

Note that in Theorem 4.23, the query subject a can easily be one of the a_j and if $k = 1$, the result implies equivalence of simple materialisation $\mathtt{Mat}^{\mathcal{ALC}}(\mathcal{K}, C\{a\})$ *iff* $\mathtt{Mat}^{\mathcal{EL}_\bot}(\mathcal{K}, C\{a\})$. Consequently, in the initial iteration of Definition 4.16, $\mathtt{Cons_s}()$ produces the same consistent DBox (for $\mathcal{ALC}$ and $\mathcal{EL}_\bot$ materialisation) for the first processed individual. The exceptionality of this first individual is decided with simple materialisation-based on the original ABox, an equivalence that is covered by Theorem 4.23 with $k = 1$. For any subsequent iteration, the ABox $\mathcal{A}_{i-1}$ contains exactly the same concept assertions, modulo the syntax of material implications. Thus, $\mathtt{Cons_s}()$ will again determine the same set of consistent DCIs for the currently processed individual, based on the equivalent ABox extensions at this point. Consequently, each step of the iteration in Definition 4.16 is equivalent in terms of the reduction to classical reasoning w.r.t. both $\mathcal{ALC}$ and $\mathcal{EL}_\bot$ (regardless of $\mathbf{s}$), when the original inference problem is in $\mathcal{EL}_\bot$. All of the above, including Thm. 4.22 and 4.23, is summed up by the following corollary.

Corollary 4.24. *For an $\mathcal{EL}_\bot$ inference problem $(\mathcal{K}, \alpha)$ and a total order $\prec$ on N_I, it holds that*

$$\mathtt{Mat}^{\mathcal{ALC}}_{\mathtt{rat}_\prec}(\mathcal{K}, \alpha) \text{ is true iff } \mathtt{Mat}^{\mathcal{EL}_\bot}_{\mathtt{rat}_\prec}(\mathcal{K}, \alpha) \text{ is true}$$

under classical semantics.

In the next chapter we discuss issues that arise from this framework for materialisation-based procedures. Some problems are local to rational semantics and can be somewhat alleviated by other instantiations of the framework, i.e. varying on the semantic strength $\mathbf{s}$ (cf. Chapter 6). Other drawbacks are inherent for the technique of materialisation in general, and therefore persist for all instantiations of the framework, regardless of semantic strength. These issues motivate the concept of *coverage* for non-monotonic semantics and the need to revaluate materialisation-based approaches.

5 DISCUSSION

Rational Closure is often considered as a very stable foundation for non-monotonic reasoning in modern KR systems. It is formally well-behaved in terms of a set of rationality postulates and it allows for an efficient reduction to classical reasoning in the underlying formalism. Adaptations of the original RC in conditional (propositional) logic have been investigated for first order logic [LM'90], modal logic [BMV'11a] and of course Description Logics [Bon'19; BCM+'13; GGOP'15; PT'18]. This is evidence that RC maintains a strong foothold in non-monotonic KR.

Nevertheless, Rational Closure still suffers from several fatal drawbacks. As a matter of fact, much of its attention is directed precisely at such shortcomings [BV'17a; CMMN'14; CS'12; CS'13; GG'18; PT'17a; PT'17b; PT'18]. On the other hand, it is worth noting that, in early stages of their research, Giordano et al. [GGOP'10a; GGOP'10b] have argued that rational entailment is actually too strong, as it allows to draw unintuitive conclusions. However, the entailment they discuss does not correspond to Rational Closure.

We recognise the most important issues as the following and continue readily to discuss the first two in depth.

INHERITANCE BLOCKING. Defeasible properties of a super-concept are inherited in an all-or-nothing [GD'16] manner, regardless of their interactions or conflicts with more specific defeasible knowledge. This is the most widely known downside of RC that has been identified and approached on numerous occasions [BFPS'15; BS'17; CS'12; CS'13; PT'17b]. It occurs already for RC in conditional logic [LM'92] and it is only natural that direct translations thereof inherit its drawbacks. (Section 5.1)

QUANTIFICATION NEGLECT. Quantified concepts are oblivious to defeasible information. This issue has been largely overlooked in the literature, with only few mentions in [Bon'19; BFPS'15; KLM'90], excluding our contributions originating in [PT'17a]. We argue that *truly meaningful* rational consequences, in any logic that supports forms of quantification, must not suffer from such neglect. (Section 5.2)

DEPENDENCE ON THE DMUP. Semantic characterisations of Rational Closure [BCM+'13; GGOP'15] are not well-defined for Description Logics that do not enjoy the DMUP. Extensions of RC to more expressive logics than $\mathcal{ALC}$ ([BV'17a; BV'17b; GD'18; GGO'18]), relied on individual solutions for semantic characterisations, if non-satisfaction of the DMUP ([CMS'18; GD'18]) prohibited a simple translation of the approach for $\mathcal{ALC}$. Bonatti [Bon'19] provides a uniform solution by generalising the notion of concept ranks and ranked models, effectively introducing a semantic characterisation of *Rational Closure for all Description Logics*.

Due to the proven equivalence of various characterisations (such as in Thm. 4.14), any formalisation of (standard) RC suffers those shortcomings. We rely on the materialisation framework to illustrate and discuss the first two properties, because defeasible consequences, and thereby the causes of problematic cases, are easily understood and portrayed in terms of material implications. Going forward, we will often utilise $\mathcal{ALC}$ material implications $\overline{C \sqsubsetsim D}$, simply because they are handled much more easily than TBox extensions, as required when remaining with $\mathcal{EL}_\bot$. Corollary 4.24 ensures that any drawbacks or other illustrations are shared for both versions of material implications. We refer the interested reader to [Bon'19] for an in-depth study of the third issue. Interestingly enough, Bonatti [Bon'19] closes his contribution with an indisputable realisation, proclaiming the urgency of the aforementioned problems.[1]

> "All of these limitations need to be addressed before Rational Closure can be applied in practice."
>
> *[Bon'19, p. 215]*

5.1 INHERITANCE-BLOCKING

To understand inheritance blocking, consider DCIs in the rational chain of a DBox as layered information. Their interaction with query subjects is governed by the subsumption hierarchy of their antecedents, as indicated by the construction of the rational chain (Def. 4.6). Recall the DCIs $\mathsf{Cat} \sqsubsetsim \mathsf{Smart}$ and $\mathsf{Cat} \sqcap \exists \mathsf{friend}.\mathsf{Dog} \sqsubsetsim \neg\mathsf{Smart} \sqcap \mathsf{Friendly}$ from Example 4.8 (Page 48) and consider the additional property $\mathsf{Cat} \sqsubsetsim \mathsf{Lazy}$ for generic Cats. Example 4.8 already illustrated how concepts further down in the subsumption hierarchy (here: $\mathsf{Cat} \sqcap \exists \mathsf{friend}.\mathsf{Dog}$), are subject to *more* defeasible information. Similar to Example 4.8, the rational chain for these three DCIs (even with an empty TBox) is

$$\begin{aligned}\mathsf{chain}(\mathcal{K}) = \langle \mathcal{D}_0 &= \mathcal{D},\\ \mathcal{D}_1 &= \{\mathsf{Cat} \sqcap \exists \mathsf{friend}.\mathsf{Dog} \sqsubsetsim \neg\mathsf{Smart} \sqcap \mathsf{Friendly}\},\\ \mathcal{D}_2 &= \emptyset\rangle.\end{aligned}$$

When presented with a query subject such as the individual molli, satisfying $\mathsf{Cat} \sqcap \exists \mathsf{friend}.\mathsf{Dog}$, rationality dictates that due to the conflict $\mathcal{A} \cup \{\overline{\mathcal{D}}(\mathsf{molli})\} \models \bot(\mathsf{molli})$, more specific information, i.e. $\mathcal{D}_1$, shall prevail. Intuitively speaking, the property Smart, associated with typical instances

1 To be precise, Bonatti specifically refers to the first two problems and lists a third issue that is also addressed in [BFPS'15]. The urgency for resolving the dependence on the DMUP is of course given by the paper [Bon'19] itself.

of Cat, is not *inherited* to typical instances of $\mathsf{Cat} \sqcap \exists \mathsf{friend.Dog}$ (and there is a good reason for that).

At a closer look, the conflict that keeps molli from satisfying $\mathcal{D}_0$ is caused by the contradictory consequences of $\mathsf{Cat} \sqsubsetsim \mathsf{Smart}$ and $\mathsf{Cat} \sqcap \exists \mathsf{friend.Dog} \sqsubsetsim \neg \mathsf{Smart} \sqcap \mathsf{Friendly}$. Unfortunately, being bound to the rational chain, $\mathsf{Cat} \sqsubsetsim \mathsf{Lazy}$ is removed at the same time as $\mathsf{Cat} \sqsubsetsim \mathsf{Smart}$, because on the subsumption hierarchy their antecedents are indistinguishable. In a nutshell, the removal of irrelevant defeasible properties within a conflicting set of properties is the essence of *inheritance blocking*. One would expect that if no conflict or reason for the removal of a defeasible property exists, it *should* be inherited to a more specific class/instance.

The problem causing this all-or-nothing [GD'16] effect, is that the layering of DCIs according to the subsumption hierarchy is too coarse. For instance, it is easy to see that *all* antecedents of DCIs that are topmost on their subsumption hierarchy must be consistent with $\mathcal{D}_0$ in *any* rational chain. Consider the additional DCI $\mathsf{Small} \sqsubsetsim \mathsf{Cute}$ about the typical properties of a concept that, on its own, is completely unrelated to the concept Cat. Without further information, Small must be topmost in the subsumption hierarchy, hence belonging to $\mathcal{D}_0$. Now, any Cat suffering from the conflict between Smart and $\neg\mathsf{Smart}$ (concepts and individuals alike) are also *blocked* from inheriting the typical property of Small, because they are incapable of satisfying everything in $\mathcal{D}_0$. Naturally, this effect does not only rely on inheritance of properties that are consistent in $\mathcal{D}_0$. Deeper layers can exhibit the same behaviour, albeit requiring more involved exemplification.

This issue of Rational Closure is known for a long time, being rooted already in the conditional logic introduced by Lehmann and Magidor [LM'92]. Numerous approaches attempting to resolve inheritance blocking as best as possible have appeared in the literature ([CMMN'14; CS'12; CS'13; GG'18; Leh'95]). These attempts typically propose a stronger entailment relation that maintains consequences under RC, but is able to derive the missing but expected consequences as well. The quality of such extensions is usually appraised in terms of computational complexity and a self-reflecting analysis for satisfaction of KLM postulates. Among the most prominent, we recognise Relevant Closure(s) (Sec. 6.1), Lexicographic Closure (Sec. 6.2) and defeasible inheritance-based Description Logics [CS'11; CS'13]. The first two will be covered in terms of the materialisation framework in Chapter 6.

5.2 QUANTIFICATION NEGLECT

Another severe limitation that most approaches towards RC in DLs or extensions thereof suffer from, is the neglect of defeasible information in quantified concepts. This behaviour is inherent with all instantiations of the materialisation framework, rather than a single aspect of it (as it is the case for inheritance blocking with the rational chain). Quantification neglect is rooted in the very idea of utilising materialised forms of defeasible information to enrich desired consequences or to determine ranks of concepts. Its essence lies in the following observation.

Materialisation cannot infer defeasible properties for quantified concepts that are satisfied by the query subject.

Illustration and comprehension of this issue requires nothing more than a single DCI, e.g. $\mathrm{Dog} \mathrel{\underset{\sim}{\sqsubset}} \mathrm{Nice}$. Effectively, from the typical instances of Dog satisfying the defeasible property Nice, we are unable to conclude for any instance that is related to an instance of Dog via a role, that it is also related to an instance of Nice, regardless of typicality or conflict of the predecessor instance. For defeasible subsumption, it is very easy to pinpoint the cause of this behaviour to materialisation:

$$\not\models (\neg\mathrm{Dog} \sqcup \mathrm{Nice}) \sqcap \exists\mathrm{friend}.\mathrm{Dog} \sqsubseteq \exists\mathrm{friend}.\mathrm{Nice}\,.$$

For defeasible instance checks, quantification neglect is not necessarily as universal as it is for the query subjects of defeasible subsumptions, but it persists nonetheless. For a defeasible instance check $\mathrm{Cat}\{\mathrm{molli}\}$, clearly the query concept Cat does not syntactically contain quantified concepts. However, the query subject can still satisfy quantifications, even based on defeasible information, e.g. $\mathcal{K} \models^{(\mathsf{rat}_\prec,\mathsf{mat})} \exists\mathrm{friend}.\mathrm{Dog}\{\mathrm{molli}\}$. Quantification neglect occurs for individuals only when quantifications are derived from *anonymous* individuals. Specifically, it matters whether $\exists\mathrm{friend}.\mathrm{Dog}\{\mathrm{molli}\}$ is derived from $\mathrm{Cat} \mathrel{\underset{\sim}{\sqsubset}} \exists\mathrm{friend}.\mathrm{Dog}$ (anonymous) or from $\mathrm{friend}(\mathrm{molli}, \mathrm{lilly})$ and $\mathrm{Dog}(\mathrm{lilly})$ (named individual). In the latter case, the extended knowledge base $\mathcal{K}_\prec$ (Def. 4.16) may already assert defeasible information for lilly, allowing to derive typical consequences for the quantified concept Dog. However in the former case, no DCI, materialised as an assertion for molli (or any other individual) is able to "affect" the quantified concept Dog, aside of the normal capabilities of universal quantification.

We argue that defeasible information should influence top-level instances, as well as *any* existentially nested instances, while remaining conflict free in the rational sense. Consider an additional DCI $\mathrm{Dog} \mathrel{\underset{\sim}{\sqsubset}} \mathrm{Angry}$ and the simple restriction $\exists\mathrm{friend}.\mathrm{Angry} \sqsubseteq \bot$. To obtain defeasible consequences of the concept $\exists\mathrm{friend}.\mathrm{Dog}$, we would have to determine a set of DCIs that all instances of this nested Dog can satisfy, while maintaining consistency of the original concept $\exists\mathrm{friend}.\mathrm{Dog}$. For semantics overcoming quantification neglect, such a selection of consistent DCIs would ideally align with a well-chosen foundation, such as the rational chain. Then again, the issue of inheritance blocking would transfer to the defeasible consequences of quantified concepts as well. The most fine-grained solution here, considering all of the above DCIs, would of course be to derive $\exists\mathrm{friend}.\mathrm{Nice}$ but not $\exists\mathrm{friend}.\mathrm{Angry}$ as defeasible subsumers of $\exists\mathrm{friend}.\mathrm{Dog}$.

While materialisation-based approaches to defeasible reasoning have been subject to quantification neglect since their first appearance in the DL setting in [CS'10], this issue has (to the best of our knowledge) not been addressed in terms of rational reasoning until 2017 [PT'17a]. The method of *overriding* introduced by Bonatti et al. [BFPS'15] allows to individually refer to the *normal* instances of a class syntactically, e.g. $\exists \mathtt{friend}.\mathsf{N}\mathtt{Dog}$. This is different from our proposition, that defeasible information should influence consequences on any level of existentially nested elements and align with the rational (or relevant) selection of consistent DCIs.

Even though inheritance blocking has received a lot more attention to this day, it is hard to compare severity or urgency of the two problems. Inheritance blocking is somewhat logic-independent, as it persists for Rational Closure in DL and conditional logic, while quantification neglect requires some form of quantification to be expressible in the underlying logic. On the other hand, the main advantage of (many) Description Logics over the propositional calculus is the ability to express binary relations and impose knowledge on the relational neighbourhood. Any non-monotonic extensions of DLs not respecting this feature could be regarded as mere compositions of propositional rational reasoning and classical DL reasoning. That being said, both insufficiencies are orthogonal in the sense that one does not entail the other. Therefore, individual solutions as well as combinations thereof are not only a vital part of the present thesis, but should play a major role in the area of Rational Closure in Description Logics in the long term.

6 EXTENSIONS

Most of the extensions of Rational Closure that have been investigated in the literature [CMMN'14; CS'12; CS'13], are motivated to resolve the issue of inheritance blocking, while retaining computability and as many formal properties as possible. Another concern with the proposal of stronger semantics is not to stray too far from the KLM foundation. In particular, attempting to adhere to their thesis that any sensible non-monotonic entailment relation should be a rational *extension* of RC (Thesis 5.25 in [LM'92]). As noteworthy contributions for a strengthening of RC we recognise the following.

RELEVANT CLOSURE. Consistent subsets of DCIs are determined by means of consequence justification [Hor'11], also referred to as axiom pinpointing [Peñ'09]. Relevant Closure has only been studied in the context of DLs by Casini et al. [CMMN'14]. (see Section 6.1)

LEXICOGRAPHIC CLOSURE. For the selection of consistent sets of DCIs, all subsets of the DBox are ranked with tuples of natural numbers, measuring how many DCIs from each layer in the rational chain they contain. These tuples are then used to define a lexicographic order over the subset of the DBox and select the most preferred, (with the query subject) consistent set. Versions of this closure operator have been studied by Casini and Straccia [CS'12] and [GG'18]. (see Section 6.2)

BOOLEAN INHERITANCE NETWORKS. *Boolean inheritance networks* (BIN) ([Hor'94]) can be utilised to determine interactions between DCIs and build a dependency graph with different types of edges. Upon determining a consistent subset of the DBox for the query subject, this network is then queried to provide a more fine grained selection of appropriate DCIs than allowed by the rational chain. The resulting entailment operation is still referred to as Rational Closure, relying on the preprocessing of a BIN. It was originally introduced for DLs by Casini and Straccia [CS'11] with more refined results in [CS'13].

CONTEXT-BASED DEFEASIBLE SUBSUMPTION. In standard RC, defeasible subsumptions are considered under one single context. In model-theoretic characterisations, this is contingent on classical interpretations being extended with a single preference relation over domain elements. Britz and Varzinczak [BV'17a; BV'18b] study the effects of extending classical models with *several* preference relations, one for every context. Defeasible subsumptions (as well as DCIs) can then be subscripted by a specific context and conflicts are effectively localised within a context.

An important addition that does not really fit into the preceding list, is the refinement of concept-rank-based RC as studied by Bonatti [Bon'19]. It is a strengthening of RC on the level of the underlying logic. Rather than supporting more inferences, its goal is to support more expressive Description Logics.

It is interesting to see that three of the four approaches above, essentially investigate different strategies to determine consistent subsets of the DBox. The algorithm that is used to ultimately determine consequences under these stronger closures is structurally identical to the computation of entailments under RC. In terms of the materialisation framework we constructed in Section 4.1, they coincide on the (simple) materialisation functions and differ only on their instantiations of $\texttt{Cons}_{\mathbf{s}}()$.[1] To highlight the versatility of the materialisation framework, we present instantiations of $\texttt{Cons}_{\mathbf{s}}()$ for Relevant and Lexicographic Closure in the following (Sec. 6.1 and 6.2). As a matter of fact, we also adopt the method of relevant reasoning in PART III of this thesis, as testament for a semantics that neither suffers from inheritance blocking nor quantification neglect. At this time, we conjecture that Lexicographic Closure and inheritance-based RC can be similarly merged with our new semantics.

6.1 RELEVANT CLOSURE

In Section 5.1 it was illustrated why Rational Closure is too coarse and how disregarding DCIs in bulk can lead to the (undesired) removal of conflict-free defeasible knowledge. Another point of critique is that the candidates for consistent sets of DCIs are determined in a preprocessing step, without a specific query subject in mind. Relevant Closure (hereinafter also referred to as relevant reasoning or semantics of relevant strength) aims to determine consistent sets of DCIs that are tailored specifically to a given query subject with the finest possible granularity. An immediate consequence of this premise is that the produced sets of DCIs are supersets of those selected from the rational chain (for any given query subject). The name-giving main intuition is to "associate relevance with the subsumptions responsible for making the antecedent of a query exceptional" *Sec. 4.1 in [CMMN'14]*. Formally, this type of relevance is captured in terms of (minimal) justifications ([Hor'11; HPS'09]). Recall the notion of exceptionality for concepts and individuals in Definition 4.5 (Page 46).

Definition 6.1 (Justification [CMMN'14]). For an $\mathcal{L}$ KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, a set of DCIs $\mathcal{J} \subseteq \mathcal{D}$ and a concept or individual $\chi \in \mathfrak{C}(\mathcal{L}) \cup \mathsf{N}_\mathrm{I}$, $\mathcal{J}$ is called a χ-*Justification iff* χ is exceptional for $\mathcal{J}$ and there is no $\mathcal{J}' \subsetneq \mathcal{J}$ s.t. χ

1 As a technicality, $\texttt{init}^{\mathcal{L}}_{\mathbf{s}}()$ needs individual instantiation as well. For more details on its purpose consider Section 6.1 and 6.2.

is exceptional w.r.t. $\mathcal{J}'$. The set of all χ-Justifications w.r.t. the KB $\mathcal{K}$ is $\texttt{justifications}(\mathcal{K},\chi)$.

This particular definition of justifications is a special instance within the area that, in general, studies minimal justifications for any type of inference. Justification is heavily investigated in different flavours within ([BP'10a; BP'10b; Hor'11; HPS'09; Peñ'09]) and outside ([CD'91; GHN+'04; HO'48; LS'05]) of Description Logics. In DLs, justifications can be effectively determined by means of automata [BP'10a] as well as extensions of general tableaux [BP'10b]. We shall not introduce these methods in detail and rather accept them as (decidable) black-box procedures, appropriately saturating the set $\texttt{justifications}(\mathcal{K},\chi)$. In terms of computational complexity, the additional step of determining justifications turns out to be non-negligible for $\mathcal{EL}_\bot$. As the detailed complexity analysis of materialisation-based entailment is referred to Chapter 8, so are the relevant results from the area of justification.

Example 6.2. Consider the KB $\mathcal{K} = (\mathcal{T},\mathcal{D})$[2]

$$\begin{aligned}
\mathcal{T} = \{&\mathit{FamilyCat} \sqsubseteq \mathit{Cat}\}\\
\mathcal{D} = \{&\mathit{Smart} \mathrel{\tilde{\sqsubseteq}} \neg\mathit{Lazy},\\
&\mathit{Cat} \mathrel{\tilde{\sqsubseteq}} \mathit{Smart},\\
&\mathit{Cat} \mathrel{\tilde{\sqsubseteq}} \mathit{Lazy},\\
&\mathit{Cat} \mathrel{\tilde{\sqsubseteq}} \mathit{Loner},\\
&\mathit{FamilyCat} \mathrel{\tilde{\sqsubseteq}} \neg\mathit{Loner}\}
\end{aligned}$$

First of all, this simple KB showcases that the rational chain is not determined solely through the strict subsumption hierarchy induced by $\mathcal{T}$, but also somewhat through the defeasible subsumption hierarchy. Clearly,

$$(\neg\mathit{Cat} \sqcup \mathit{Smart}) \sqcap (\neg\mathit{Cat} \sqcup \mathit{Lazy}) \sqcap (\neg\mathit{Smart} \sqcup \neg\mathit{Lazy})$$

cannot be satisfied by elements of Cat. This shows that Cat cannot satisfy the entire DBox. The defeasible information about $\mathit{FamilyCats}$ conflicting with one of the assumptions about generic Cats, shows that $\mathit{FamilyCat}$ must be exceptional w.r.t. any subset of the DBox containing both $\mathit{FamilyCat} \mathrel{\tilde{\sqsubseteq}} \neg\mathit{Loner}$ and $\mathit{Cat} \mathrel{\tilde{\sqsubseteq}} \mathit{Loner}$. Consequently, the rational chain of the KB is as follows:

$$\begin{aligned}
\texttt{chain}(\mathcal{K}) = \langle \mathcal{D}_0 = \mathcal{D},\\
\mathcal{D}_1 &= \{\mathit{Cat} \mathrel{\tilde{\sqsubseteq}} \mathit{Smart}, \mathit{Cat} \mathrel{\tilde{\sqsubseteq}} \mathit{Lazy}, \mathit{Cat} \mathrel{\tilde{\sqsubseteq}} \mathit{Loner},\\
&\qquad \mathit{FamilyCat} \mathrel{\tilde{\sqsubseteq}} \neg\mathit{Loner}\},\\
\mathcal{D}_2 &= \{\mathit{FamilyCat} \mathrel{\tilde{\sqsubseteq}} \neg\mathit{Loner}\},\\
\mathcal{D}_3 &= \emptyset\rangle
\end{aligned}$$

2 This KB is of propositional nature, that is, it does not utilise any form of DL quantification. Rest assured that the ideas presented in this example remain valid with more expressive DL consequences. We opt for simplicity, to convey an understandable intuition. Furthermore, the forms of negation used in this KB are easily transformed into GCIs such as $\mathit{NotLazy} \sqcap \mathit{Lazy} \sqsubseteq \bot$, allowing to transfer the arguments in this example to $\mathcal{EL}_\bot$ KBs.

Clearly, this is another good example for inheritance blocking, as any rational consequence about $\mathrm{FamilyCats}$ is constrained by the rational chain to disregard e.g. $\mathrm{Cat} \sqsubsetsim \mathrm{Smart}$, even though this is not in conflict with any (defeasible) information known for $\mathrm{FamilyCat}$. Resolving this particular conflict in a more fine-grained manner, requires to disregard only one of the involved DCIs. On the other hand, the conflict between $\mathrm{Cat} \sqsubsetsim \mathrm{Smart}$, $\mathrm{Cat} \sqsubsetsim \mathrm{Lazy}$ and $\mathrm{Smart} \sqsubsetsim \neg\mathrm{Lazy}$ is also inherited to $\mathrm{FamilyCat}$, showing that exceptionality might be witnessed by several conflicts.

A major aspect in relevant reasoning is the strategy of how to resolve (possibly multiple) identified conflicts in a reasonable way. This question relates quite elegantly to Reiter's hitting set duality [Rei'87]. Minimality of the justifications for a conflict implies that a consistent subset of the DBox can be found by removing at least one statement of every justification. Casini et al. [CMMN'14] propose two methods for disregarding DCIs that appear in the justifications of a query subject, resulting in the characterisation of *Basic* and *Minimal* Relevant Closure. Both propose a specific removal of relevant DCIs, that is closely aligned with the rational chain. In the effort to obtain a semantics that *extends* Rational Closure, it is intuitively preferred to remove more general DCIs, i.e. those that belong to bigger members of the rational chain.

BASIC RELEVANT CLOSURE produces consistent subsets of the DBox by iteratively removing all DCIs along the rational chain that are part of some exceptionality justification. This iteration begins by removing *all* relevant DCIs that are *exclusive* to $\mathcal{D}_0$ and continues with increasing indices until the current query subject is consistent with the remaining defeasible knowledge. In Example 6.2, Basic Relevant Closure produces the same consequences for $\mathrm{FamilyCat}$ as Rational Closure, because all DCIs that are *exclusive* to $\mathcal{D}_0$ and $\mathcal{D}_1$ belong to some justification for $\mathrm{FamilyCat}$.

MINIMAL RELEVANT CLOSURE is a further refinement, reducing the amount of removed DCIs. Only the *rank-minimal* DCIs of every justification will be removed, allowing to inherit the properties Smart and Lazy for $\mathrm{FamilyCats}$, because the rank-minimal DCI in this particular conflict is $\mathrm{Smart} \sqsubsetsim \neg\mathrm{Lazy}$. (cf. Exm. 6.4)

In [CMMN'14] it was shown that Minimal Relevant Closure extends both Rational and Basic Relevant Closure. Therefore, we skip the intermediate step and formally define only Minimal Relevant Closure, also omitting the keyword *Minimal* in the following. The semantics of relevant strength presented in PART III are also based on the foundation of Minimal Relevant Closure. Furthermore, note that in [CMMN'14] only defeasible subsumption is considered. The following instantiation enables defeasible instance checking, a novel and unparalleled contribution that resolves inheritance blocking for this type of inference.

A formalisation of Relevant Closure in terms of the materialisation framework requires little more than an appropriate instantiation of the consistent-selection function $\mathtt{Cons}_{\mathtt{rel}}()$.

Definition 6.3 (Relevant Consistent DCIs). For an $\mathcal{L}$ KB $\mathcal{K}$ with $\mathtt{chain}(\mathcal{K}) = \langle \mathcal{D}_0, \cdots, \mathcal{D}_n \rangle$, a concept or individual $\chi \in \mathfrak{C}(\mathcal{L}) \cup \mathsf{N}_\mathsf{I}$ and $\mathtt{justifications}(\mathcal{K}, \chi) = \langle \mathcal{J}_0, \cdots, \mathcal{J}_m \rangle$, let

$$\mathcal{J}_i^{min} = \{C \sqsubsetsim D \in \mathcal{J}_i \mid \forall E \sqsubsetsim F \in \mathcal{J}_i . r_{\mathcal{K}}(C) \leqslant r_{\mathcal{K}}(E)\}$$

for $0 \leqslant i \leqslant m$. The *relevant* consistent-selection function $\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, \chi)$ is defined as follows

$$\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, \chi) = \mathcal{D} \setminus \bigcup_{0 \leqslant i \leqslant m} \mathcal{J}_i^{min}.$$

Continuing with Example 6.2 allows to show superiority of $\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, \chi)$ over $\mathtt{Cons}_{\mathtt{rat}}(\mathcal{K}, \chi)$.

Example 6.4. The two justifications for $\mathit{FamilyCat}$ are clearly

$$\mathcal{J}_0 = \{\mathit{Smart} \sqsubsetsim \neg\mathit{Lazy}, \mathit{Cat} \sqsubsetsim \mathit{Smart}, \mathit{Cat} \sqsubsetsim \mathit{Lazy}\}$$
$$\mathcal{J}_1 = \{\mathit{Cat} \sqsubsetsim \mathit{Loner}, \mathit{FamilyCat} \sqsubsetsim \neg\mathit{Loner}\}$$

with

$$\mathcal{J}_0^{min} = \{\mathit{Smart} \sqsubsetsim \neg\mathit{Lazy}\}$$
$$\mathcal{J}_1^{min} = \{\mathit{Cat} \sqsubsetsim \mathit{Loner}\}.$$

Using $\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, \mathit{FamilyCat}) = \mathcal{D} \setminus (\mathcal{J}_0^{min} \cup \mathcal{J}_1^{min})$, we can clearly derive

$$\mathcal{T} \models \overline{\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, \mathit{FamilyCat})} \sqcap \mathit{FamilyCat} \sqsubseteq \mathit{Smart} \sqcap \mathit{Lazy},$$

something that was impossible under Rational Closure.

The definition of $\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, \chi)$ already suffices to determine defeasible subsumption inferences under $(\mathtt{rel}, \mathtt{mat})$-semantics using Definition 4.11. The clear separation of functionality in the materialisation framework make this type of instantiation particularly simple.

To determine entailment of defeasible instance checks under $(\mathtt{rel}_\prec, \mathtt{mat})$-semantics requires only to define a starting point $\mathtt{init}^{\mathcal{L}}_{\mathtt{rel}}(\mathcal{K})$ for the ABox extension defined generically in Def. 4.16. Now, we can finally illustrate why this initialisation is parametric on a semantic strength, rather than starting from the original ABox in any case.

Example 6.5. Consider an extension of Example 6.2 by an ABox and an additional GCI as $\mathcal{K} = (\mathcal{A}, \mathcal{T} \cup \mathcal{T}', \mathcal{D})$, with

$$\begin{aligned}\mathcal{A} = \{&\mathit{FamilyCat}(\mathit{molli}),\\ &\mathit{Cat}(\mathit{daisy}),\\ &\mathit{friend}(\mathit{daisy}, \mathit{molli})\}\\ \mathcal{T}' = \{&\mathit{Smart} \sqsubseteq \forall \mathit{friend}.(\neg\mathit{Smart})\}\end{aligned}$$

and $\mathcal{T},\mathcal{D}$ as in Example 6.2. Clearly the rational chain is

$$\mathtt{chain}(\mathcal{K}) = \langle \mathcal{D}_0, \mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3 \rangle,$$

as before. We assume the order $\prec$ over $\mathsf{N_I}$ to assign $\mathfrak{molli} \prec \mathfrak{daisy}$, to showcase a problem that arises when initialising the algorithm in Definition 4.16 with $\mathtt{init}^{\mathcal{ALC}}_{\mathtt{rel}}(\mathcal{K}) = \mathcal{K}$. Because $\mathfrak{molli}$ is considered before $\mathfrak{daisy}$, the consistent subset $\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, \mathfrak{molli})$ is determined. It is not hard to confirm that $\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, \mathfrak{molli}) = \mathcal{D} \setminus \{\mathsf{Smart} \mathrel{\tilde{\sqsubseteq}} \neg\mathsf{Lazy}, \mathsf{Cat} \mathrel{\tilde{\sqsubseteq}} \mathsf{Loner}\}$. Thus, the initial ABox is extended with $(\neg\mathsf{Cat} \sqcup \mathsf{Smart})(\mathfrak{molli})$, among others.[3] This allows to draw the conclusion $\mathsf{Smart}(\mathfrak{molli})$ from $(\mathcal{A}_1, \mathcal{T}, \mathcal{D})$. Upon determining $\mathtt{Cons}_{\mathtt{rel}}((\mathcal{A}_1, \mathcal{T}, \mathcal{D}), \mathfrak{daisy})$, it is clear from the GCI in $\mathcal{T}'$, the newly derivable fact $\mathsf{Smart}(\mathfrak{molli})$, and $\mathsf{friend}(\mathfrak{daisy}, \mathfrak{molli})$, that we will not be able to derive $\mathsf{Smart}(\mathfrak{daisy})$ after the ABox extension is complete. According to our knowledge base, *all* Smart elements only befriend elements that are not Smart.

So far so good, but how do the consequences under the relevant ABox extension relate to those under the rational ABox extension? With the same order, $\mathfrak{molli}$ would only be (rationally) consistent with $\mathcal{D}_2$, hence we would not be able to conclude $\mathsf{Smart}(\mathfrak{molli})$ after the first iteration of the extension algorithm. This in turn allows $\mathfrak{daisy}$ to be consistent with the DCI $\mathsf{Cat} \mathrel{\tilde{\sqsubseteq}} \mathsf{Smart}$ and thus, with $\mathcal{D}_1$, because there is no conflict derivable through the relation $\mathsf{friend}(\mathfrak{daisy}, \mathfrak{molli})$. Initialising the relevant ABox extension with the original KB resulted in the *loss* of a consequence that is derivable under RC. As we are motivated to strictly extend the consequences obtainable through Rational Closure, we need to initialise the relevant ABox extension appropriately.

Formally, for the initial KB in Definition 4.16 we simply use

$$\mathtt{init}^{\mathcal{ALC}}_{\mathtt{rel}_\prec}(\mathcal{K}) = (\mathcal{A}_{\mathsf{rat}_\prec}, \mathcal{T}, \mathcal{D}),$$

and

$$\mathtt{init}^{\mathcal{EL}_\bot}_{\mathtt{rel}_\prec}(\mathcal{K}) = (\mathcal{A}_{\mathsf{rat}_\prec}, \mathcal{T}^{\mathcal{D}}, \mathcal{D}),$$

for an $\mathcal{ALC}$ or $\mathcal{EL}_\bot$ KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, respectively. This ensures the following result for both types of defeasible inference, obtained for the framework instantiation $\mathtt{Mat}^{\mathcal{L}}_{\mathtt{rel}_\prec}(\mathcal{K}, \alpha)$. Superiority of defeasible subsumption under Relevant Closure was shown in [CMMN'14], while superiority of defeasible instance checking as well as its definition altogether, are our contributions in [PT'18].

Theorem 6.6. *Relevant materialisation-based reasoning is strictly stronger than rational materialisation-based reasoning. In particular, for $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ and a defeasible subsumption or instance check α,*

$$\mathcal{K} \models^{(\mathsf{rat}_\prec, \mathsf{mat})} \alpha \implies \mathcal{K} \models^{(\mathsf{rel}_\prec, \mathsf{mat})} \alpha \tag{6.1}$$

$$\mathcal{K} \models^{(\mathsf{rat}_\prec, \mathsf{mat})} \alpha \not\Longleftarrow \mathcal{K} \models^{(\mathsf{rel}_\prec, \mathsf{mat})} \alpha \tag{6.2}$$

3 Concept assertions with conjunction on the top level are clearly equivalent to separate concept assertions for each conjunct.

Proof. For α being a defeasible subsumption, this result appeared as Proposition 2 in [CMMN'14].

Let $\alpha = C\{a\}$ and suppose $\mathcal{K} \models^{(\mathtt{rat}_{\prec},\mathtt{mat})} \alpha$. Theorem 4.22 and 4.23 allow us to rely on $\mathcal{ALC}$ materialisation here and transfer the result to $\mathcal{EL}_{\bot}$. The assumption that α is entailed by $\mathcal{K}$ under rational strength implies the classical entailment $(\mathcal{A}_{\mathtt{rat}_{\prec}}, \mathcal{T}) \models C(a)$. From Definition 4.16 and $\mathtt{init}^{\mathcal{ALC}}_{\mathtt{rel}_{\prec}}(\mathcal{K}) = (\mathcal{A}_{\mathtt{rat}_{\prec}}, \mathcal{T}, \mathcal{D})$ it follows that $\mathcal{A}_{\mathtt{rat}_{\prec}} \subseteq \mathcal{A}_{\mathtt{rel}_{\prec}}$. Due to monotonicity of classical reasoning, this implies that $(\mathcal{A}_{\mathtt{rel}_{\prec}}, \mathcal{T}) \models C(a)$, hence $\mathcal{K} \models^{(\mathtt{rel}_{\prec},\mathtt{mat})} \alpha$.

Example 6.4 (with the ABox from Exm. 6.5) shows that the converse is not always true. With the appropriate initialisation, we are able to derive $Lazy\{molli\}$ from the relevant but not from the rational ABox extension. □

6.2 LEXICOGRAPHIC CLOSURE

Determining consistent sets of DCIs in a more fine-grained manner than provided by the rational chain goes back to the *Lexicographic Closure* in [BCD+'93; Leh'95], for the propositional case. To the best of our knowledge, justification-based reasoning was not considered for propositional defeasible reasoning in this context. Casini and Straccia [CS'12] have lifted this more expressive closure by a syntactic translation from the propositional to the DL case, much like they did for Rational Closure in the first place. The fundamental idea for Lexicographic Closure is similar to Relevant Closure, in that *all* subsets of the DBox are considered to determine $\mathtt{Cons}_{\mathtt{lex}}()$. However, the approaches differ in the technique to select such a consistent subset. Lexicographic Closure is based on a lexicographic ranking over all subsets of the DBox. Due to its similarity in DBox-granularity with the Relevant Closure, we can reuse Example 6.2 and be very brief with this instantiation of the materialisation framework.

For the mathematical foundation, recall the lexicographic ordering on strings of natural numbers of length k, to be defined as

$$\langle n_0, \cdots, n_k \rangle <_{\mathtt{lex}} \langle m_0, \cdots, m_k \rangle$$
iff
$$\exists i \in \{0, \cdots, k\}. \big(\forall j < i. n_j = m_j\big) \wedge n_i < m_i$$

The intuition behind a lexicographic ordering of sets of DCIs, is that every subset of the DBox $\mathcal{D}$ is measured by the number of DCIs it contains, component-wise separated by their antecedents concept ranks. Recall concept ranks from Def. 4.12 (Page 51), and note that $r_{\mathcal{K}}(C) = i$ *iff* $C \sqsubseteq\!\!\!\!\sim D \in \mathcal{D}_i \setminus \mathcal{D}_{i+1}$ is an equivalent characterisation of the concept ranks for antecedents C in $\mathcal{D}$.

Definition 6.7 (Lexicographic Rank). For a KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, let $\mathtt{chain}(\mathcal{K}) = \langle \mathcal{D}_0, \cdots, \mathcal{D}_{k+1} \rangle$. For each subset $\mathcal{U} \subseteq \mathcal{D}$, the tuple $\langle n_0, \cdots, n_k \rangle_{\mathcal{U}}$ is defined with

$$n_i = |\{C \sqsubsim D \in \mathcal{U} \mid r_{\mathcal{K}}(C) = k - i\}|$$

for $0 \leqslant i \leqslant k$.

The lexicographic rank of $\mathcal{U}$ is simply counting the number of contained DCIs for every antecedent rank, in reversed order. This ordering gives an implicit preference on more specific (i.e. higher antecedent rank) DCIs, when maximising the lexicographic rank. Recall $\mathcal{K}$ and $\mathtt{chain}(\mathcal{K})$ from Example 6.2. In rational strength of reasoning, we were forced to use $\mathcal{D}_2 = \{\mathit{FamilyCat} \sqsubsim \neg \mathit{Loner}\}$ as the consistent set of DCIs for $\mathit{FamilyCat}$. The lexicographic rank of $\mathcal{D}_2$ is $\langle 1, 0, 0 \rangle_{\mathcal{D}_2}$ and there are clearly more preferable (in terms of a higher lexicographic rank) subsets of $\mathcal{D}$, that are consistent with $\mathit{FamilyCat}$. Take for instance $\mathcal{U} = \mathtt{Cons_{rel}}(\mathcal{K}, \mathit{FamilyCat})$, as in Example 6.4. It is easy to see that

$$\langle 1, 0, 0 \rangle_{\mathcal{D}_2} <_{\mathtt{lex}} \langle 1, 2, 0 \rangle_{\mathcal{U}}\,.$$

Ideally, the lexicographic ordering provides a most preferable set of consistent DCIs, maximising the number of more specific DCIs in descending order. However, there might be multiple consistent sets of DCIs with the same lexicographic rank. This case is treated *cautiously*, by considering the intersection of those rank-maximal candidates.

For a slight abuse of notation, suppose we have a universal instance of the $\mathtt{Cons}()$ function, providing the set of all consistent subsets of the DBox w.r.t. a query subject. Explicitly, for $\alpha \in \mathfrak{C}(\mathcal{L}) \cup \mathrm{N_I}$ and a KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, let

$$\mathtt{Cons}(\mathcal{K}, \alpha) = \{\mathcal{U} \subseteq \mathcal{D} \mid \alpha \text{ is not exceptional w.r.t. } \mathcal{U} \text{ and } \mathcal{K}\}$$

Definition 6.8 (Lexicographic Consistent DCIs). For an $\mathcal{L}$ KB $\mathcal{K}$ with $\mathtt{chain}(\mathcal{K}) = \langle \mathcal{D}_0, \cdots, \mathcal{D}_{k+1} \rangle$, and a concept or individual $\chi \in \mathfrak{C}(\mathcal{L}) \cup \mathrm{N_I}$, the *lexicographic* consistent-selection function $\mathtt{Cons_{lex}}(\mathcal{K}, \chi)$ is defined as

$$\mathtt{Cons_{lex}}(\mathcal{K}, \chi) = \bigcap \mathit{max}_{<_{\mathtt{lex}}}(\mathtt{Cons}(\mathcal{K}, \alpha))$$

for $\mathit{max}_{<_{\mathtt{lex}}}(\mathtt{Cons}(\mathcal{K}, \alpha))$ selecting the lexicographic-rank-maximal elements from $\mathtt{Cons}(\mathcal{K}, \alpha)$.

As opposed to Relevant Closure, Casini and Straccia [CS'12] consider defeasible instance checking under Lexicographic Closure. However, in [CS'12] the ABox extensions are not initialised with the rational ABox extension. Consequently, the effect illustrated in Example 6.5 could easily take hold. Nevertheless, they show that the closure resulting from their ABox extension is rational, in terms of translated KLM-postulates for instance checks. To align the Lexicographic Closure that we present, with the ABox extension approach taken for Relevant Closure, we define $\mathtt{init}^{\mathcal{L}}_{\mathtt{lex}_\prec}(\mathcal{K}) = \mathtt{init}^{\mathcal{L}}_{\mathtt{rel}_\prec}(\mathcal{K})$ to be based on the rational ABox extension as well. While this does not guarantee that the satisfaction of KLM postulates carries over from [CS'12], it ensures the second part of Lehmann and Magidor's thesis, by extending RC.

The following result is proven analogous to Theorem 6.6. Example 6.2 can still be used to show superiority, and the simple argument that no consistent set of DCIs $\mathcal{U} \subsetneq \mathtt{Cons}_{\mathtt{rat}}(\mathcal{K}, \chi)$ can be lexicographically preferable to $\mathtt{Cons}_{\mathtt{rat}}(\mathcal{K}, \chi)$, allows to conclude that the Lexicographic Closure contains the Rational Closure.

Theorem 6.9. *Lexicographic materialisation-based reasoning is strictly stronger than rational materialisation-based reasoning. In particular, for $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ and a defeasible subsumption or instance check α,*

$$\mathcal{K} \models^{(\mathtt{rat}_{\prec}, \mathtt{mat})} \alpha \implies \mathcal{K} \models^{(\mathtt{lex}_{\prec}, \mathtt{mat})} \alpha \tag{6.3}$$

$$\mathcal{K} \models^{(\mathtt{rat}_{\prec}, \mathtt{mat})} \alpha \not\Longleftarrow \mathcal{K} \models^{(\mathtt{lex}_{\prec}, \mathtt{mat})} \alpha \tag{6.4}$$

Part III

NESTED RATIONAL CLOSURE AND BEYOND

INTRODUCTION AND OVERVIEW

In the preceding parts we have introduced the foundations for reasoning in Description Logics and the paradigm of non-monotonic reasoning. A promising approach to combine these two foundations into *defeasible reasoning*, is a reduction algorithm relying on materialisation to compute entailments under Rational Closure. By abstracting the central notions of this reduction, we were able to present a framework capable of deriving entailments under several materialisation-based semantics. This branch of research evidently strives for continued improvement, be it the incorporation of more powerful tools to apply defeasibility [BCMV'13; BV'16; BV'17a; BV'18b], encourage more expressive DL formalisms [Bon'19; BV'17a; BV'17b; GGO'18], or to alleviate crucial drawbacks [CMMN'14; CS'12; CS'13; GG'18; Leh'95]. Surprisingly, until 2017 ([PT'17a; PT'17b]), studies towards the latter aspect have focused almost exclusively on resolving the well-known problem of inheritance blocking (Sec. 5.1), in favour of supporting *more intuitive* consequences. However, we argue that the issue of *quantification neglect* (Sec. 5.2) is much more concerning, when the underlying logic distinguishes itself from less expressive formalisms mainly by the ability to express forms of quantification. The defeasible consequences that are produced by any formalism suffering from quantification neglect are at best of propositional coverage, as defeasible information is only inferred on the top level (conjunction/disjunction) of concepts. This dulls the benefit of representing defeasible knowledge in Description Logics, as the resulting consequences are hardly *more meaningful* than their corresponding representation in propositional logic (modulo the expressivity of classical DL consequences).

In this part we will present new model-theoretic semantics that are capable of deriving defeasible consequences for arbitrarily *nested* concepts, both for defeasible subsumption and defeasible instance checks. Our two dimensional characteristic for semantics of specific strength and coverage, will be extended by *nested coverage* with this model-theoretic characterisation. We showcase the capabilities of nested semantics in both, rational and relevant strength. The latter provides a solution for both of the aforementioned drawbacks, introducing a very powerful form of defeasible entailment. The main difficulty in defining nested semantics is that a formal description of how defeasible information in DLs should be inferred for quantified concepts is not easily derived from the KLM foundations. Unfortunately, the postulates about first-order quantification that are discussed in [LM'90], do not transfer to the restricted form of quantification in DLs. The only occurrence of DL quantification in a discussion on the entailment of formal properties appears in [BCM+'13]. However, the postulates Britz et al. present are derived from the basic set of KLM postulates for rational entailment relations. While they provide some insight on the behaviour of quantified concepts w.r.t. the propositional postulates, they obviously do not impose additional restrictions on the entailment relation, such that quantification neglect might be resolved. At this time there is no formal recognition of properties, such as new postulates, describing the expected behaviour of nested semantics in DLs. No doubt, answering this question requires extensive studies, marking an important path for future research.

To maintain a strong foothold within the area of rational reasoning, our nested semantics are designed in such a way that consequences about quantified concepts align with the strength of the respective materialisation-based (propositional) coverage. Consider for instance a query subject $\exists r.A$. We propose nested *rational* semantics such that the consequences derivable for the quantified occurrence of A never exceed the consequences that would be derivable if A itself was the query subject, i.e. the consequences derivable for A in the materialisation framework. For relevant strength, the fine-graded nature of the consistent subset selection of the DBox increases the complexity of our semantics. Effectively all subsets of the DBox (exponentially many) have to be considered for quantified occurrences of A, rather than a predetermined (linear) number of subsets.

The presentation of our nested semantics is separated into two technical chapters. In Chapter 7, the novel formalism of *typicality models*, is defined and qualitatively analysed, studying the behaviour and inferable entailments for such models. This chapter concludes by formally showing *superiority* (preserving and extending entailments under RC) of nested over materialisation-based semantics, in particular, showcasing the ability to derive the missing consequences. An extensive study of the computational complexity for deciding entailment under the new semantics follows in Chapter 8. It includes algorithmic characterisations of entailments under nested rational (and relevant) semantics, relying on classical DL reasoning. The reduction from a known SAT problem is used to prove hardness of deciding nested rational entailment, whereas hardness results for relevant strength remain an open problem.

7 TYPICALITY MODELS

Quantification neglect can be illustrated with a very simple scenario. For a solitary DCI $A \sqsubsetsim B$ ($\mathcal{A} = \mathcal{T} = \emptyset$), the conclusion $\exists r.A \sqsubsetsim \exists r.B$ cannot be drawn with the materialisation framework, because the produced concept $(\neg A \sqcup B) \sqcap \exists r.A$ can be satisfied by an interpretation $\mathcal{I}$ with an element $d \in ((\neg A \sqcup B) \sqcap \exists r.A)^{\mathcal{I}}$, such that no element $e \in B^{\mathcal{I}}$ is an r-successor of d in $\mathcal{I}$. The way all materialisation reductions are defined (regardless of strength), such models are always allowed and will inevitably enable *quantification neglect*.

We carefully design a semantics that—instead of considering *all* models of a KB—disregards models which do not satisfy the *expected* DCIs at role-successor elements. This selection of "appropriate" models is formalised with a so-called *typicality preference relation* on sets of models. Reasoning only with the most preferred sets of models effectively allows to derive inferences including defeasible knowledge on quantified concepts. Because it is highly non-trivial to achieve this for general models and classical semantics, we make the first step in the right direction for a slightly restricted scenario.

In Section 2.4 we showed that for classical $\mathcal{EL}_{\bot}$ query entailment, it suffices to consider representative models over a unique domain $\Delta^{\mathbb{C},\mathbb{O}}$ for a certain (finite) context $\mathbb{C}$, $\mathbb{O}$. Because comparing (sets of) models becomes much easier when they share a common structure, we adopt this feature of representability for our semantics. Representative interpretations will be extended to incorporate defeasible information. More specifically, it is necessary to

1. capture concept representatives that are *typical*, in terms of satisfied DCIs, and
2. require individual representatives to satisfy appropriate sets of DCIs.

The former is required to decide entailments $C \sqsubsetsim D$ by checking if the *most typical* C representative—the one that also satisfies $\mathtt{Cons}_{\mathbf{s}}(\mathcal{K}, C)$ (for instantiations of $\mathbf{s}$)—satisfies D in the selected set of models. Likewise, the latter is required to decide entailments $C\{a\}$ while requiring the a-representative to satisfy the defeasible information extending the original ABox $\mathcal{A}$, i.e. $\mathcal{A}_{\mathbf{s}_{\prec}}$. Both of these entailment conditions are aligned with the result in Theorem 2.19, showing that it is sufficient to inspect a single representative for deciding entailments of the KB.

Formally, we introduce *typicality interpretations* based on the appropriate extensions of representative domains, called *typicality domains*. In such domains, every element is associated with a set of DCIs in addition to the original association with a concept or individual from the given context. Our exposition is then separated into two parts.

In the beginning (Sec. 7.1, 7.2), we follow the outline for representative models closely (recall Sec. 2.3), lifting notions such as the standard and model property to typicality interpretations. This provides a foundation for the second part and shows that the restriction to a unique domain is a natural one. The first part (much like for representative models), concludes in a semantics characterised in terms of all typicality models. We show that entailments under these semantics coincide with the entailments produced by the materialisation framework, and thus consider them to be of *propositional coverage* (prop). In the second part (Sec. 7.3), we strengthen the propositional semantics by defining the discussed preference relation over sets of typicality models and achieve two types of *nested coverage* (nest).

For the remainder of Chapter 7 and 8, we assume all inference problems, KBs, and concepts to be in $\mathcal{EL}_\perp$, unless specified otherwise. Note that Theorem 4.22 and 4.23 justify using $\mathcal{ALC}$ material implications, to allow for more intuitive illustrations and simpler proofs.

7.1 INTRODUCING TYPICALITY INTERPRETATIONS

The central notion in our extension of representative models for reasoning over defeasible knowledge is the *typicality domain*. The main idea is to extend classical representative domains by attaching to every concept- and individual-representative a set of DCIs. The meaning of this attached defeasible knowledge is that the respective domain element shall satisfy the associated DCIs, something that can be enforced by extending the classical model-property (Def. 2.7). Entailments are, as in the classical case (Thm. 2.19), read off of the representative domain element that is associated with the appropriate DCIs. For example to determine $\mathcal{K} \models^{(\mathbf{s},\mathbf{c})} C \sqsubsetsim D$, we must inspect the C-representative that is associated with $\mathtt{Cons}_\mathbf{s}(\mathcal{K}, C)$. In that sense, the sets of DCIs that representatives are associated with in the typicality domain, determine the strength of the resulting semantics. Thus, to obtain reasoning semantics of a certain strength $\mathbf{s}$, the typicality domain must be specifically constructed to align the typicality of representatives with $\mathbf{s}$. As this construction relies on $\mathtt{Cons}_\mathbf{s}(\mathcal{K}, \chi)$, specific typicality domains depend on an input KB $\mathcal{K}$. Similar to the generality of the materialisation framework, in terms of different instantiations on the strength $\mathbf{s}$, most of the results in this part are impartial to the specific structure of the underlying typicality domain, giving plausibility for a generic presentation.

Definition 7.1 (Typicality Domain). For a finite set of DCIs $\mathcal{E}$, and a context $\mathbb{C} \subseteq \mathfrak{C}(\mathcal{EL}_\perp)$, $\mathbb{O} \subseteq \mathsf{N_I}$, $\Delta^{\mathbb{C},\mathbb{O}}$ is called a *typicality domain iff*

1. $\Delta^{\mathbb{C},\mathbb{O}} \subseteq (\mathbb{C} \cup \mathbb{O}) \times \mathcal{P}(\mathcal{E})$,
2. $\forall C \in \mathbb{C}.(C, \emptyset) \in \Delta^{\mathbb{C},\mathbb{O}}$,
3. $\forall a \in \mathbb{O}.|\{(\chi, \mathcal{U}) \in \Delta^{\mathbb{C},\mathbb{O}} \mid \chi = a\}| = 1$

A domain element $(\chi, \mathcal{U}) \in \Delta^{\mathbb{C},\mathbb{O}}$ is called

- a *concept representative*, if $\chi \in \mathbb{C}$, and
- an *individual representative*, if $\chi \in \mathbb{O}$.

Intuitively, an element $(\chi, \mathcal{U})$ represents either the subclass of elements in C that satisfy the DCIs in $\mathcal{U} \subseteq \mathcal{E}$, or an individual that is required to satisfy the DCIs in $\mathcal{U}$. For generic domain elements of a typicality domain $\Delta^{\mathbb{C},\mathbb{O}}$, we may still use the usual descriptors $d, e \in \Delta^{\mathbb{C},\mathbb{O}}$.

Remark 7.2. Note that general typicality domains will not carry the set of DCIs $\mathcal{E}$ in their notation, because any (finite) set of DCIs can be considered. The definition of specific typicality domains will eventually rely on a given defeasible KB $\mathcal{K}$ with its finite DBox $\mathcal{D}$. Every element in $\mathbb{C} \cup \mathbb{O}$ is represented in $\Delta^{\mathbb{C},\mathbb{O}}$ at least once and individuals are also represented at most once. Therefore, an important, yet simple consequence of Def. 7.1 is that $\Delta^{\mathbb{C},\mathbb{O}}$ is *finite iff* the context $\mathbb{C}, \mathbb{O}$ is finite. We frequently introduce $\Delta^{\mathbb{C},\mathbb{O}}$ as finite, simultaneously ensuring finiteness of $\mathbb{C}$ and $\mathbb{O}$.

We continue to lift basic notions from representative domains (Sec. 2.4) to typicality domains.

Definition 7.3 (Typicality Interpretation)**.** For a context $\mathbb{C}, \mathbb{O}$, an interpretation $\mathcal{I} = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\mathcal{I}})$ is called a *typicality interpretation iff* $\Delta^{\mathbb{C},\mathbb{O}}$ is a typicality domain over $\mathbb{C}, \mathbb{O}$.

In Theorem 2.19, it was shown that in models over a representative domain, it suffices to check entailment for a single representative, rather than for all instances of the query subject. This motivates that satisfaction of defeasible inferences in typicality interpretations is also captured in terms of the concepts satisfied by the appropriate representatives. Defeasible instance checking has the same definition as its counterpart in classical semantics, whereas for defeasible subsumption, we investigate the *most typical* representatives of a concept in general. Intuitively, the typicality of two representatives for the same concept are compared in terms of the subset relation between associated DCIs.

Definition 7.4 (Satisfaction in Typicality Interpretations)**.** For a context $\mathbb{C}, \mathbb{O}$, a typicality interpretation $\mathcal{I} = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\mathcal{I}})$, $C \in \mathbb{C}$ and $a \in \mathbb{O}$, it holds that

1. $\mathcal{I} \models C \sqsubsim D$ *iff* $(C, \mathcal{U}) \in D^{\mathcal{I}}$ for all $(C, \mathcal{U}) \in \Delta^{\mathbb{C},\mathbb{O}}$, such that $\neg\exists (C, \mathcal{U}') \in \Delta^{\mathbb{C},\mathbb{O}}$ with $\mathcal{U} \subsetneq \mathcal{U}'$, and
2. $\mathcal{I} \models C\{a\}$ *iff* $a^{\mathcal{I}} \in C^{\mathcal{I}}$.

For the specific semantic strengths rational and relevant, the typicality domain is defined in such a way, that there is only one $\subseteq$-maximal $\mathcal{U} \subseteq \mathcal{D}$ (for the DBox $\mathcal{D}$ of the input KB) associated to each represented concept. Specifically, for rational reasoning, this will be $\mathsf{Cons}_{\mathsf{rat}}(\mathcal{K}, C)$. The DCIs associated with individual representatives, will be the DCIs that contributed to the ABox extension based on the total order $\prec$ on N_I.

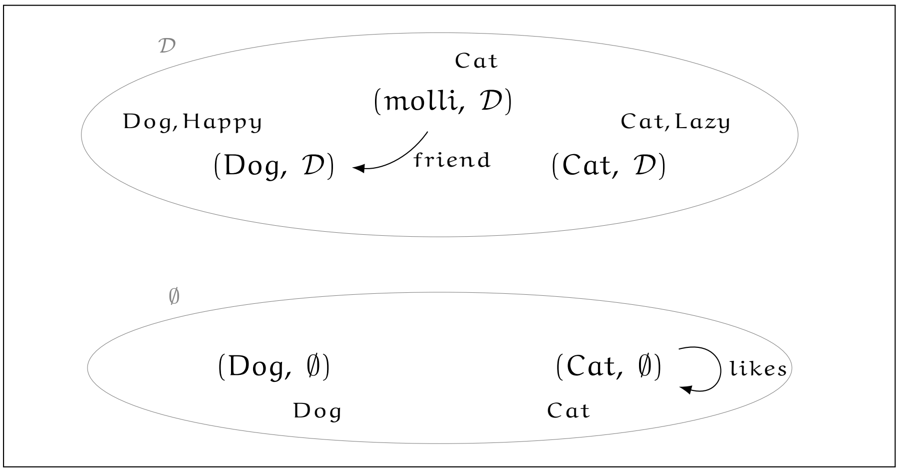

Figure 7.1: The (structured) labelled graph representation for the typicality interpretation $\mathcal{I} = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\mathcal{I}})$ in Example 7.5.

Example 7.5. To illustrate a generic typicality interpretation, consider a context with two concept names $\mathbb{C} = \{\mathsf{Cat}, \mathsf{Dog}\}$ and a single individual $\mathbb{O} = \{\mathsf{molli}\}$. Figure 7.1 depicts an interpretation $\mathcal{I}$ over the typicality domain

$$\Delta^{\mathbb{C},\mathbb{O}} = \{(\mathsf{Cat},\ \mathcal{D}), (\mathsf{Dog},\ \mathcal{D}), (\mathsf{molli},\ \mathcal{D}), (\mathsf{Cat},\ \emptyset), (\mathsf{Dog},\ \emptyset)\},$$

with $\mathsf{molli}^{\mathcal{I}} = (\mathsf{molli},\ \mathcal{D})$. Because we are not discussing any specific KB or the relation that $\mathcal{I}$ has to it, suppose $\mathcal{D}$ is an arbitrary finite set of DCIs. The following is a list of noteworthy observations and features that can be extracted from the labelled graph representation of $\mathcal{I}$.

1. Domain elements can be (visually) grouped according to their associated set of DCIs. For instance, in semantics of rational strength, all domain elements can be (visually) organised as a matrix (e.g. Fig. 7.3).

2. Note, there is only one representative for molli (due to Property 3 of Def. 7.1) but several for Cat and Dog, at least those associated with $\emptyset$ (cf. Property 2 of Def. 7.1). Other concepts such as Happy or Lazy can occur in the interpretation (signature) without requiring representation, as long as they do not belong to the context.

3. Despite its specific domain, $\mathcal{I}$ is still a (classical) DL interpretation and therefore, maintains their functionality. In particular, the satisfaction of classical queries/GCIs/assertions/KBs remains intact. Assuming $\mathsf{molli}^{\mathcal{I}} = (\mathsf{molli},\ \mathcal{D})$, the classical entailments $(\exists \mathsf{friend}.\mathsf{Dog})(\mathsf{molli})$ and $\exists \mathsf{likes}.\mathsf{Cat} \sqsubseteq \mathsf{Cat}$ are supported by $\mathcal{I}$.

4. Definition 7.4 covers the entailment of defeasible subsumptions and instance checks by single typicality interpretations. For example, $\mathcal{I}$ satisfies $\mathsf{Cat} \mathrel{\sqsubseteq\!\!\!\!\sim} \mathsf{Lazy}$, because the concept representative of Cat that is associated with a $\subseteq$-maximal set of DCIs also satisfies Lazy.

On the other hand, satisfaction of defeasible instance checks w.r.t. solitary typicality interpretations coincides with satisfaction of classical instance checks, i.e. $\mathcal{I} \models (\exists \mathtt{friend}.\mathtt{Happy})\{\mathtt{molli}\}$ holds as much as $\mathcal{I} \models (\exists \mathtt{friend}.\mathtt{Happy})(\mathtt{molli})$.

This list of observations is not exhaustive, but it should be enough to provide an intuition of how to read visualised typicality models in future examples. For clarity, we will often display the labelled graph representation of an interpretation, rather than its set-based definition. It is only important to note that, in those cases, the labelled graph depicts the *full* domain and interpretation mapping, unless stated otherwise.

7.1.1 *Properties of Typicality Interpretations*

We continue to lift the properties in Definition 2.15 and 2.16 to accommodate for the added association with DCIs, beginning with the characterisation of *standard* typicality interpretations.

Definition 7.6 (Standard Property)**.** For a context $\mathbb{C}, \mathbb{O}$, a typicality interpretation $\mathcal{I} = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\mathcal{I}})$ is *standard iff*

1. $(C, \mathcal{U}) \in C^{\mathcal{I}}$ for all $(C, \mathcal{U}) \in \Delta^{\mathbb{C},\mathbb{O}}$,
2. $a^{\mathcal{I}} = (a, \mathcal{U}) \in \Delta^{\mathbb{C},\mathbb{O}}$ for all $a \in \mathbb{O}$ and
3. $d \in (\exists r.F)^{\mathcal{I}} \implies (d, (F, \emptyset)) \in r^{\mathcal{I}}$ for all $d \in \Delta^{\mathbb{C},\mathbb{O}}$, $r \in N_R$ and $F \in \mathbb{C}$.

The standard property of typicality interpretations is analogous to that in the classical case (Def. 2.15), projecting the first two properties on the first component of a typicality domain element. The last property ensures again that existing role successors are witnessed by an appropriate representative (as in Def. 2.15). However, now the expected successor is additionally restricted to not be associated with any DCIs. For one, this is because only concept representatives associated with $\emptyset$ are guaranteed to exist in $\Delta^{\mathbb{C},\mathbb{O}}$. Also, we intuitively expect elements of a typicality domain to comply with the defeasible information they are associated with. Thus, assigning witnesses (for satisfied existential restrictions) that are associated with (some) DCIs would already influence the derivable entailments for that particular quantified concept. This intuitive compliance with associated DCIs is formally captured by extending the standard model property to check for domain elements satisfying their DCIs.

Definition 7.7 (Model Property)**.** For a KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ and a context $\mathbb{C}, \mathbb{O}$, a typicality interpretation $\mathcal{I} = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\mathcal{I}})$ is a *typicality model* of $\mathcal{K}$ *iff*

1. $\mathcal{I} \models (\mathcal{A}, \mathcal{T})$ and
2. $\forall (\chi, \mathcal{U}) \in \Delta^{\mathbb{C},\mathbb{O}}.\big(\forall E \mathrel{\underset{\sim}{\sqsubset}} F \in \mathcal{U}.((\chi, \mathcal{U}) \in E^{\mathcal{I}} \implies (\chi, \mathcal{U}) \in F^{\mathcal{I}})\big)$.

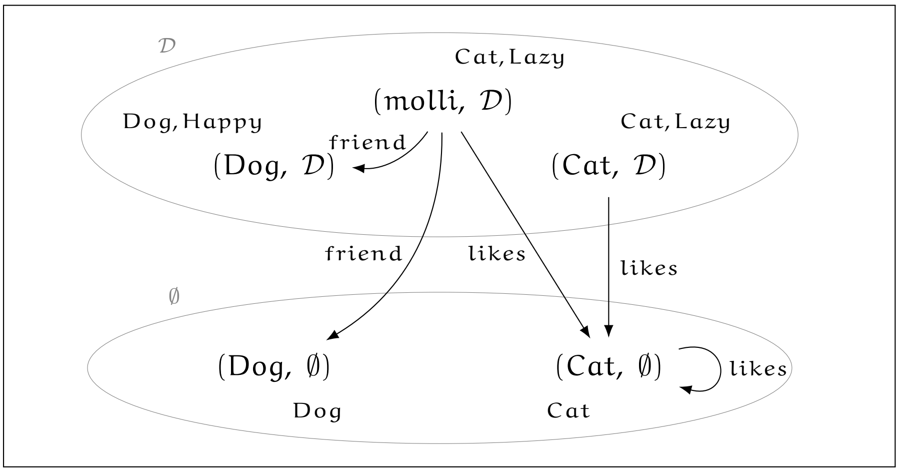

Figure 7.2: The (structured) labelled graph representation for the standard typicality model $\mathcal{J} = (\Delta^{\mathbf{C},\mathbf{O}}, \cdot^{\mathcal{J}})$ in Example 7.9.

$\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbf{C},\mathbf{O}})$ denotes the set of all *standard typicality models* of $\mathcal{K}$ over $\Delta^{\mathbf{C},\mathbf{O}}$. The typicality domain $\Delta^{\mathbf{C},\mathbf{O}}$ is *consistent* with the KB $\mathcal{K}$ *iff* $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbf{C},\mathbf{O}}) \neq \emptyset$.

Remark 7.8. A direct consequence of Def. 7.7 is, that $(\chi, \mathcal{U}) \in \overline{\mathcal{U}}^{\mathcal{I}}$ for all $\mathcal{I} \in \mathtt{TMod}(\mathcal{K}, \Delta^{\mathbf{C},\mathbf{O}})$ and $(\chi, \mathcal{U}) \in \Delta^{\mathbf{C},\mathbf{O}}$. Because all interpretations $\mathcal{I} \in \mathtt{TMod}(\mathcal{K}, \Delta^{\mathbf{C},\mathbf{O}})$ are standard (Def. 7.6), it holds for all $(C, \mathcal{U}), (a, \mathcal{U}') \in \Delta^{\mathbf{C},\mathbf{O}}$, that $(C, \mathcal{U}) \in C^{\mathcal{I}}$ and $(a, \mathcal{U}') = a^{\mathcal{I}}$. Together with the first observation, it follows that $(C, \mathcal{U}) \in (C \sqcap \overline{\mathcal{U}})^{\mathcal{I}}$ and $\mathcal{I} \models \overline{\mathcal{U}'}(a)$. Clearly, $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbf{C},\mathbf{O}}) \subseteq \mathtt{Mod}(\mathcal{K})$, because of 1 in Def. 7.7. Hence, $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbf{C},\mathbf{O}}) \neq \emptyset$ implies for all $(C, \mathcal{U}) \in \Delta^{\mathbf{C},\mathbf{O}}$ that $\overline{\mathcal{U}} \sqcap C$ is consistent with $\mathcal{K}$. Furthermore, from $\mathcal{I} \models \overline{\mathcal{U}'}(a)$ (for all standard typicality models $\mathcal{I}$ of $\mathcal{K}$ and $(a, \mathcal{U}') \in \Delta^{\mathbf{C},\mathbf{O}}$), it follows that

$$\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbf{C},\mathbf{O}}) \subseteq \mathtt{Mod}((\mathcal{A} \cup \{\overline{\mathcal{U}'}(a) \mid (a, \mathcal{U}') \in \Delta^{\mathbf{C},\mathbf{O}}\}, \mathcal{T})).$$

This assortment of immediate observations for standard typicality models is very handy in several of the upcoming proofs.

Example 7.9. Recall the interpretation $\mathcal{I}$ from Example 7.5 (see Fig. 7.1) and consider a specific KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ with

$$\mathcal{A} = \{(\mathrm{Cat} \sqcap \exists \mathrm{friend}.\mathrm{Dog})(\mathrm{molli})\}$$
$$\mathcal{T} = \{\mathrm{Cat} \sqsubseteq \exists \mathrm{likes}.\mathrm{Cat}\}$$
$$\mathcal{D} = \{\mathrm{Dog} \mathrel{\underset{\sim}{\sqsubset}} \mathrm{Happy}, \mathrm{Cat} \mathrel{\underset{\sim}{\sqsubset}} \mathrm{Lazy}\}$$

To begin with, $\mathcal{I}$ is not a standard model of $\mathcal{K}$. While it does satisfy $\mathrm{molli}^{\mathcal{I}} = (\mathrm{molli}, \mathcal{D})$ (Property 2) and Property 1 of Def. 7.6, it lacks the edge $\mathrm{friend}((\mathrm{molli}, \mathcal{D}), (\mathrm{Dog}, \emptyset))$, to satisfy the last aspect of the standard property. Extending $\mathcal{I}$ with this edge would render it a standard typicality interpretation, though it would still not satisfy $\mathcal{K}$. In $\mathcal{I}$, $(\mathrm{molli}, \mathcal{D})$ is a counterexample for both, the satisfaction of $\mathrm{Cat} \sqsubseteq \exists \mathrm{likes}.\mathrm{Cat}$ and

$Cat \mathrel{\tilde{\sqsubseteq}} Lazy$ (in terms of Def. 7.7). Figure 7.2 depicts another typicality interpretation $\mathcal{J}$ that does satisfy the standard model property as well as the KB $\mathcal{K}$ according to Definition 7.7. As it turns out, $\mathcal{J}$ is the smallest extension of $\mathcal{I}$ (i.e. $\mathcal{I} \subseteq \mathcal{J}$, cf. Definition 2.11) that belongs to $\texttt{TMod}(\mathcal{K}, \Delta^{\mathrm{C,O}})$. At the same time, $\mathcal{J}$ is not a $\subseteq$-smallest element in $\texttt{TMod}(\mathcal{K}, \Delta^{\mathrm{C,O}})$, because the edge $friend((molli, \mathcal{D}), (Dog, \mathcal{D}))$ is superfluous for $\mathcal{J}$ to satisfy $\mathcal{K}$ and Definition 7.6.

Typicality models will be used to determine entailments under semantics of different strength and coverage. For variations on the strength $\mathbf{s} \in \{\texttt{rat}, \texttt{rel}\}$, we explained how the construction of a specific typicality domain, formally characterised as $\Delta_{\mathbf{s}}^{\mathrm{C,O},\mathcal{K}}$, provides the appropriate association of represented elements with DCIs. To achieve different coverages $\mathbf{c}$, we intuitively described how the set of considered models will be reduced in terms of a typicality preference relation. Formally, this set of *considered models* is captured as a subscripted instance of $\texttt{TMod}()$, parametrised by the specific *coverage* $\mathbf{c}$ (e.g. $\texttt{prop}$, $\texttt{nest}$) that the resulting semantics will possess, i.e. $\texttt{TMod}_{\mathbf{c}}(\mathcal{K}, \Delta^{\mathrm{C,O}}) \subseteq \texttt{TMod}(\mathcal{K}, \Delta^{\mathrm{C,O}})$. Defeasible entailment under $(\mathbf{s}, \mathbf{c})$-Semantics, using typicality models, can be captured in general, similar to defeasible entailment in the materialisation framework (Def. 4.3).

Definition 7.10 (Defeasible Entailment under $(\mathbf{s}, \mathbf{c})$-Semantics)**.** A defeasible inference problem $(\mathcal{K}, \alpha)$ is entailed under $(\mathbf{s}, \mathbf{c})$-semantics *iff*

$$\forall \mathcal{I} \in \texttt{TMod}_{\mathbf{c}}(\mathcal{K}, \Delta_{\mathbf{s}}^{\mathrm{C,O},\mathcal{K}}).\mathcal{I} \models \alpha. \tag{7.1}$$

where $\Delta_{\mathbf{s}}^{\mathrm{C,O},\mathcal{K}}$ is the $\mathbf{s}$-typicality domain for $\mathcal{K}$.

As before, $\mathbf{s}$ includes a total order $\prec$ on $\mathrm{N_I}$ if individuals and defeasible instance checks are considered. Finally, it can be explicitly seen, how semantics for instantiations of $\mathbf{s}$ and $\mathbf{c}$ are captured by defining $\Delta_{\mathbf{s}}^{\mathrm{C,O},\mathcal{K}}$ as well as $\texttt{TMod}_{\mathbf{c}}()$ in terms of their parameters. At first, in Sec. 7.2, we will illustrate how the straight-forward approach, using *all* standard typicality models of a KB, defines semantics that still suffer from quantification neglect. Such semantics are considered to be of propositional coverage, explicitly

$$\texttt{TMod}_{\texttt{prop}}(\mathcal{K}, \Delta_{\mathbf{s}}^{\mathrm{C,O}}) = \texttt{TMod}(\mathcal{K}, \Delta_{\mathbf{s}}^{\mathrm{C,O}}).$$

However, before investigating propositional coverage of rational and relevant strength explicitly, we establish some important properties for standard typicality models that are true for any underlying typicality domain. In particular, the second observation in the end of Section 2.4 can be adopted to the defeasible case, and (unlike in Sec. 2.4) we now make this result explicit. Specifically, we show that standard typicality models are closed under intersection, identifying a $\subseteq$-smallest member in their midst, providing a powerful tool for algorithmic characterisations of entailments based on typicality models.

7.2 MINIMAL TYPICALITY MODELS

Using classical representative models over a KB, entailments can be decided if their context permits it. Specifically, $\mathcal{K} \models C \sqsubseteq D$ can only be decided using representative models over the context $\mathbb{C}$, $\mathbb{O}$, if it contains and is consistent with $\mathcal{K}$, is quantification closed, and $C \in \mathbb{C}$ (cf. Theorem 2.19). In theory, considering a context to encompass all concepts and individuals ($\mathbb{C} = \mathfrak{C}(\mathcal{EL}_\bot)$, $\mathbb{O} = N_I$) is the least restricted case, because all queries are covered by this context. One can argue that, in practice, it suffices to consider only finite contexts, that are still large enough w.r.t. $\mathcal{K}$ and the relevant query subjects. This is mostly motivated for computability reasons. In theoretical studies it is also a common approach to consider a restriction to the finite case before moving to the more general variant (e.g. [Bon'19]). For typicality interpretations, any typicality domain over a finite context is also finite (cf. Rem. 7.2). Because KBs are always assumed to be finite, they also occupy only a finite signature. This limits the number of distinct typicality interpretations over that signature and a finite typicality domain, to be finite as well. In turn, this provides some well-behavioural properties of finite typicality interpretations/models, beginning with the following intermediary result.

Lemma 7.11. *For a finite typicality domain $\Delta^{\mathbb{C},\mathbb{O}} \subseteq (\mathbb{C} \cup \mathbb{O}) \times \mathcal{P}(\mathcal{E})$ and two standard typicality interpretations $\mathcal{I} = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\mathcal{I}})$ and $\mathcal{J} = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\mathcal{J}})$ it holds that*

$$C^{\mathcal{I} \cap \mathcal{J}} = C^{\mathcal{I}} \cap C^{\mathcal{J}} \text{ for all } \mathcal{EL}_\bot \text{ concepts } C \text{ with } \mathrm{Qc}(C) \subseteq \mathbb{C}.$$

Proof. We prove this claim by induction on C. The cases $C = A \in N_C$ as well as $C = E \sqcap F$ are trivial (assuming the claim holds for E and F). Let $C = \exists r.E$ for $r \in N_R$ and hypothesise that $E^{\mathcal{I} \cap \mathcal{J}} = E^{\mathcal{I}} \cap E^{\mathcal{J}}$.

- $(\exists r.E)^{\mathcal{I} \cap \mathcal{J}} \subseteq (\exists r.E)^{\mathcal{I}} \cap (\exists r.E)^{\mathcal{J}}$ holds by Definition 2.11 and the induction hypothesis, because $r^{\mathcal{I} \cap \mathcal{J}} = r^{\mathcal{I}} \cap r^{\mathcal{J}}$ and $E^{\mathcal{I} \cap \mathcal{J}} = E^{\mathcal{I}} \cap E^{\mathcal{J}}$.
- For $(\exists r.E)^{\mathcal{I} \cap \mathcal{J}} \supseteq (\exists r.E)^{\mathcal{I}} \cap (\exists r.E)^{\mathcal{J}}$, $d \in (\exists r.E)^{\mathcal{I}} \cap (\exists r.E)^{\mathcal{J}}$ implies that $(d, (E, \emptyset)) \in r^{\mathcal{I}} \cap r^{\mathcal{J}}$, as well as $(E, \emptyset) \in E^{\mathcal{I}} \cap E^{\mathcal{J}}$, because $\mathcal{I}$ and $\mathcal{J}$ are standard (Def. 7.6). It then follows by definition of $\cap$ and the induction hypothesis, that $d \in (\exists r.E)^{\mathcal{I} \cap \mathcal{J}}$. □

The preceding lemma is used to show that the information shared between two standard typicality models (for the same $\mathcal{K}$ and $\Delta^{\mathbb{C},\mathbb{O}}$) suffices to satisfy $\mathcal{K}$ and the standard property. More specifically, their intersection interpretation is again a standard typicality model of $\mathcal{K}$ over $\Delta^{\mathbb{C},\mathbb{O}}$.

Proposition 7.12. *For a finite typicality domain $\Delta^{\mathbb{C},\mathbb{O}} \subseteq (\mathbb{C} \cup \mathbb{O}) \times \mathcal{P}(\mathcal{E})$ and a knowledge base $\mathcal{K}$, $\texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ is closed under intersection.*

Proof. For a finite $\mathcal{K}$ we only need to consider the finite signature $sig(\mathcal{K})$ in $\texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$. Thus, $\Delta^{\mathbb{C},\mathbb{O}}$ being finite implies $\texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ to be finite. Consequently, it is sufficient to show the claim for finite intersection, i.e.

$$\mathcal{I}, \mathcal{J} \in \texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \text{ implies } \mathcal{I} \cap \mathcal{J} \in \texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}).$$

We need to show that $\mathcal{I} \cap \mathcal{J}$ is a standard typicality interpretation (Def. 7.6) and a model of $\mathcal{K}$ (Def. 7.7). $\texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \neq \emptyset$ implies that the context contains the necessary individuals and quantified concepts in $\mathcal{K}$ to yield at least one standard typicality model of $\mathcal{K}$ over $\Delta^{\mathbb{C},\mathbb{O}}$. The converse trivially satisfies the claim.

Because $\mathcal{I}$ and $\mathcal{J}$ are standard, $\mathcal{I} \cap \mathcal{J}$ is well-defined and immediately satisfies 2 of Def. 7.6. Additionally, Lemma 7.11 implies that $\mathcal{I} \cap \mathcal{J}$ satisfies $(\mathsf{C}, \mathcal{U}) \in \mathsf{C}^{\mathcal{I}\cap\mathcal{J}}$ for all $(\mathsf{C}, \mathcal{U}) \in \Delta^{\mathbb{C},\mathbb{O}}$ (Property 1 of Def. 7.6). Towards the third property, for each $\mathsf{E} \in \mathbb{C}$, it holds that $d \in (\exists r.\mathsf{E})^{\mathcal{I}\cap\mathcal{J}}$ implies $d \in (\exists r.\mathsf{E})^{\mathcal{I}}$ and $d \in (\exists r.\mathsf{E})^{\mathcal{J}}$, again, by Lemma 7.11. The definition of $\cap$ and $\mathcal{I}, \mathcal{J}$ each satisfying 3 in Def. 7.6, imply $(d, (\mathsf{E}, \emptyset)) \in r^{\mathcal{I}\cap\mathcal{J}}$, making $\mathcal{I} \cap \mathcal{J}$ a standard typicality interpretation.

A simple consequence of Lemma 7.11 is

$$(\mathsf{E}^{\mathcal{I}} \subseteq \mathsf{F}^{\mathcal{I}}) \wedge (\mathsf{E}^{\mathcal{J}} \subseteq \mathsf{F}^{\mathcal{J}}) \implies \mathsf{E}^{\mathcal{I}\cap\mathcal{J}} \subseteq \mathsf{F}^{\mathcal{I}\cap\mathcal{J}}. \quad (\star)$$

Thus, assuming $\mathcal{I}$ and $\mathcal{J}$ to satisfy $\mathcal{K}$ according to Definition 7.7, Lemma 7.11 is sufficient to imply satisfaction of all

- GCIs in $\mathcal{T}$ (with $(\star)$),
- concept assertions in $\mathcal{A}$ (with Lem. 7.11),
- role assertions in $\mathcal{A}$ (with Def. 2.11), and
- DCIs associated with domain elements (with $(\star)$),

by $\mathcal{I} \cap \mathcal{J}$. Hence, $\mathcal{I} \cap \mathcal{J} \in \texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$. □

The implication of Proposition 7.12 is (analogous to the classical case) the existence of a $\subseteq$-smallest typicality model in $\texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$.

Definition 7.13 (Minimal Typicality Model). For a finite typicality domain $\Delta^{\mathbb{C},\mathbb{O}}$, and a KB $\mathcal{K}$ the *minimal typicality model* in $\texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ is defined as

$$\texttt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) = \bigcap_{\mathcal{J} \in \texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})} \mathcal{J} \quad (7.2)$$

if $\texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \neq \emptyset$.

Remark 7.14. $\texttt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ is called the *minimal* typicality model, rather than least, smallest or canonical, because the only role edges (besides edges between individual representatives) that are guaranteed to be shared among all standard typicality models of $\mathcal{K}$ over $\Delta^{\mathbb{C},\mathbb{O}}$, have *exclusively* elements $(\mathsf{C}, \emptyset)$ as successors. For an illustration, recall Example 7.9 with $\mathcal{J} \in \texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ depicted in Figure 7.2. All standard interpretations must contain a role successor to a concept representative associated with $\emptyset$, for all domain elements satisfying existential restrictions[1]. However, as

1 For all elements in existential restrictions with fillers contained in the context.

noted in Exm. 7.9, the edge $\mathtt{friend}((\mathtt{molli},\ \mathcal{D}),(\mathtt{Dog},\ \mathcal{D}))$ is not required to uphold the model or standard property of $\mathcal{J}$.

This behaviour is desired for standard interpretations, to guarantee closure under intersection and thus, existence of the minimal typicality model. Therefore, its name is derived from the *minimal typicality* that the ranges of its extensions of role names are subjected to.

It follows trivially from Definition 7.13 that $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \subseteq \mathcal{J}$ for all standard typicality models $\mathcal{J}$ of $\mathcal{K}$ over $\Delta^{\mathbb{C},\mathbb{O}}$. Thus, the previously discussed straightforward semantics, considering *all* models in $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$, can be captured by reasoning only over the minimal typicality model. In the following, entailments under such semantics (of propositional coverage) are compared to the entailments captured by the materialisation framework, on a technical level. This link appropriately aligns our new formalism with the established fundamental results.

7.2.1 *Propositional Coverage vs. Materialisation*

From Definition 7.13 it is obvious, that if $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \neq \emptyset$ then $\mathcal{I} \models \alpha$ for all $\mathcal{I} \in \mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ *iff* $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \models \alpha$ for defeasible queries α over the context $\mathbb{C}$, $\mathbb{O}$. We show in this section, how minimal typicality models can be used, to bridge from typicality model semantics of *propositional* coverage to materialisation and vice versa. Initially, this correspondence is shown for general typicality domains $\Delta^{\mathbb{C},\mathbb{O}}$. Those results can then be applied to the specific typicality domains that ultimately produce rational or relevant semantics.

In the classical case of reasoning with representative interpretations, a powerful tool was provided in Lemma 2.17. It showed that arbitrary interpretations can be somewhat normalised into a given representative domain, while maintaining their entailments in terms of classical queries. This transformation is the essence of why the restriction to a fixed domain is prosperous. We lift the transformation in Lemma 2.17 to construct typicality interpretations from arbitrary interpretations, in terms of their materialisation-based entailments.

Definition 7.15. For an interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ and a typicality domain $\Delta^{\mathbb{C},\mathbb{O}}$, let $\tau(\mathcal{I}, \Delta^{\mathbb{C},\mathbb{O}}) = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\tau(\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})})$ with

$$
\begin{aligned}
A^{\tau(\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})} &= \{(C,\ \mathcal{U}) \in \Delta^{\mathbb{C},\mathbb{O}} \mid \mathcal{I} \models \overline{\mathcal{U}} \sqcap C \sqsubseteq A\} \\
&\quad \cup \{(a,\ \mathcal{U}) \in \Delta^{\mathbb{C},\mathbb{O}} \mid \mathcal{I} \models A(a)\} \\
r^{\tau(\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})} &= \{((C,\ \mathcal{U}), (D,\ \emptyset)) \in (\Delta^{\mathbb{C},\mathbb{O}})^2 \mid \mathcal{I} \models \overline{\mathcal{U}} \sqcap C \sqsubseteq \exists r.D\} \\
&\quad \cup \{((a,\ \mathcal{U}), (D,\ \emptyset)) \in (\Delta^{\mathbb{C},\mathbb{O}})^2 \mid \mathcal{I} \models (\exists r.D)(a)\} \\
&\quad \cup \{((a,\ \mathcal{U}), (b,\ \mathcal{U}')) \in (\Delta^{\mathbb{C},\mathbb{O}})^2 \mid (a^{\mathcal{I}}, b^{\mathcal{I}}) \in r^{\mathcal{I}}\} \\
a^{\tau(\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})} &= (a,\ \mathcal{U})
\end{aligned}
$$

for all $A \in N_C$, $r \in N_R$ and $a \in \mathbb{O}$. $\tau(\mathcal{I}, \Delta^{\mathbb{C},\mathbb{O}})$ is called the *typicality interpretation transformation* of $\mathcal{I}$ into $\Delta^{\mathbb{C},\mathbb{O}}$.

Intuitively, the domain elements in $\tau(\mathcal{I}, \Delta^{\mathbb{C},\mathbb{O}})$ represent the information that $\mathcal{I}$ contains about each represented concept/individual, also regarding the DCIs associated with concept representatives. As far as this tool is concerned, it is enough to consider defeasible information only for concept representatives. Satisfaction of defeasible information for individuals is implicit when considering this transformation for the model of an *extended ABox*. The following intermediary result, verifying the construction in Definition 7.15, is central to link entailments based on typicality interpretations with those based on materialisation. It is analogous to the result for the classical case in Lemma 2.17.

Lemma 7.16. *For a typicality domain $\Delta^{\mathbb{C},\mathbb{O}}$ over a quantification closed context $\mathbb{C}$, $\mathbb{O}$ and an interpretation $\mathcal{I}$ with*

$$\forall (C, \mathcal{U}) \in \Delta^{\mathbb{C},\mathbb{O}}.(\overline{\mathcal{U}} \sqcap C)^{\mathcal{I}} \neq \emptyset, \tag{7.3}$$

it holds that

1. $(C, \mathcal{U}) \in (\overline{\mathcal{U}} \sqcap C)^{\tau(\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})}$ *(for all* $(C, \mathcal{U}) \in \Delta^{\mathbb{C},\mathbb{O}}$*),*
2. $(C, \mathcal{U}) \in D^{\tau(\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})}$ iff $\mathcal{I} \models \overline{\mathcal{U}} \sqcap C \sqsubseteq D$ *(for all* $(C, \mathcal{U}) \in \Delta^{\mathbb{C},\mathbb{O}}$ *and* $Qc(D) \subseteq \mathbb{C}$*), and*
3. $(a, \mathcal{U}) \in D^{\tau(\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})}$ iff $\mathcal{I} \models D(a)$ *(for all* $(a, \mathcal{U}) \in \Delta^{\mathbb{C},\mathbb{O}}$ *and* $Qc(D) \subseteq \mathbb{C}$*).*

Proof. We prove the claims one by one, with the exception of Claim 1. At first, we show a special case of Claim 1, where $\mathcal{U} = \emptyset$. This special case suffices to prove Claim 2, which in turn can be used to prove the general case of Claim 1. In the corresponding classical case, Claim 2 would immediately imply Claim 1. This is not the case in the present lemma, because Claim 1 is about materialised concepts and in Claim 2, D is an $\mathcal{EL}_\bot$ concept.

CLAIM 1 ($\mathcal{U} = \emptyset$). We begin by proving the special case of 1 where $\mathcal{U} = \emptyset$. Explicitly, for every $(C, \emptyset) \in \Delta^{\mathbb{C},\mathbb{O}}$ it holds that

$$(C, \emptyset) \in C^{\tau(\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})}. \tag{$\star$}$$

This proof relies on a slightly different induction than usual, covering conjunction implicitly.[2] For the induction start, assume $C = A_1 \sqcap \cdots \sqcap A_n$ is a conjunction of concept names $A_i \in N_C$. From Def. 7.15 it directly follows that $(C, \emptyset) \in A_i^{\tau(\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})}$ for $1 \leqslant i \leqslant n$. For the induction step, let $C = \bigsqcap_{1 \leqslant i \leqslant n} A_i \sqcap \bigsqcap_{1 \leqslant j \leqslant m} \exists r_j.E_j$.[3] Because of $C \in \mathbb{C}$ and $\mathbb{C}$ being quantification closed, it holds that $E_j \in \mathbb{C}$ and thus $(E_j, \emptyset) \in \Delta^{\mathbb{C},\mathbb{O}}$ for $1 \leqslant j \leqslant m$. Therefore, it is justified for the induction hypothesis, to assume

2 For $C = E \sqcap F$, the induction hypothesis that the claim holds for E and F, would only be justified if $E, F \in \mathbb{C}$, something that is not a requirement for $\mathbb{C}$ in this lemma.
3 This is the most general representation of any $\mathcal{EL}$ concept. We need not consider $C = \bot$, because of assumption (7.3).

that $(\star)$ holds for $E_1, \ldots, E_m$. It trivially holds that C is subsumed by each of its conjuncts individually, in any interpretation. In particular, $\mathcal{I} \models C \sqsubseteq A_i$ as well as $\mathcal{I} \models C \sqsubseteq \exists r_j.E_j$. Therefore, it follows from Definition 7.15 that $(C,\ \emptyset) \in A_i^{\tau(\mathcal{I},\Delta^{C,O})}$ as well as $((C,\ \emptyset),(E_j,\ \emptyset)) \in r_j^{\tau(\mathcal{I},\Delta^{C,O})}$ for $1 \leqslant i \leqslant n$ and $1 \leqslant j \leqslant m$. Using the induction hypothesis ($(E_j,\ \emptyset) \in E_j^{\tau(\mathcal{I},\Delta^{C,O})}$), we conclude that $(C,\ \emptyset) \in C^{\tau(\mathcal{I},\Delta^{C,O})}$.

CLAIM 2. We prove Claim 2 by (the usual) induction on D, relying only on the special case $(\star)$ of Claim 1 The induction start, for $D = A \in N_C$ holds by Def. 7.15. Assume the claim holds for two concepts E, F and consider $D = E \sqcap F$. The following are equivalent:

- $(C,\ \mathcal{U}) \in (E \sqcap F)^{\tau(\mathcal{I},\Delta^{C,O})}$,
- $(C,\ \mathcal{U}) \in E^{\tau(\mathcal{I},\Delta^{C,O})}$ and $(C,\ \mathcal{U}) \in F^{\tau(\mathcal{I},\Delta^{C,O})}$,
- $\mathcal{I} \models \overline{\mathcal{U}} \sqcap C \sqsubseteq E$ and $\mathcal{I} \models \overline{\mathcal{U}} \sqcap C \sqsubseteq F$ (by induction hypothesis),
- $\mathcal{I} \models \overline{\mathcal{U}} \sqcap C \sqsubseteq E \sqcap F$.

For the most involved case, consider $D = \exists r.E$ ($r \in N_R$). By assumption, $Qc(D) \subseteq \mathbb{C}$ which means $(E,\ \emptyset)$ is contained in $\Delta^{C,O}$.

[$\Longrightarrow$] $(C,\ \mathcal{U}) \in (\exists r.E)^{\tau(\mathcal{I},\Delta^{C,O})}$ implies the existence of an edge $((C,\ \mathcal{U}),(X,\ \emptyset)) \in r^{\tau(\mathcal{I},\Delta^{C,O})}$ such that $(X,\ \emptyset) \in E^{\tau(\mathcal{I},\Delta^{C,O})}$. That the successor element $(X,\ \emptyset)$ has to be a concept representative, associated with $\emptyset$, can be seen from Definition 7.15. By the induction hypothesis, $(X,\ \emptyset) \in E^{\tau(\mathcal{I},\Delta^{C,O})}$ implies $\mathcal{I} \models X \sqsubseteq E$. From Def. 7.15 and $((C,\ \mathcal{U}),(X,\ \emptyset)) \in r^{\tau(\mathcal{I},\Delta^{C,O})}$ it follows that $\mathcal{I} \models \overline{\mathcal{U}} \sqcap C \sqsubseteq \exists r.X$, which in turn implies $\mathcal{I} \models \overline{\mathcal{U}} \sqcap C \sqsubseteq \exists r.E$.

[$\Longleftarrow$] For the other direction, $\mathcal{I} \models \overline{\mathcal{U}} \sqcap C \sqsubseteq \exists r.E$ directly implies $((C,\ \mathcal{U}),(E,\ \emptyset)) \in r^{\tau(\mathcal{I},\Delta^{C,O})}$ by Def. 7.15. It then follows from $(\star)$ (i.e. $(E,\ \emptyset) \in E^{\tau(\mathcal{I},\Delta^{C,O})}$) that $(C,\ \mathcal{U}) \in (\exists r.E)^{\tau(\mathcal{I},\Delta^{C,O})}$.

CLAIM 1 (CONTINUED). Finally, we use Claim 2 to show Claim 1 for any set of DCIs $\mathcal{U}$ (for $(C,\ \mathcal{U}) \in \Delta^{C,O}$). The induction used to prove $(\star)$ also proves $(C,\ \mathcal{U}) \in C^{\tau(\mathcal{I},\Delta^{C,O})}$. Assume for a contradiction that there is $E \precsim F \in \mathcal{U}$ s.t. $(C,\ \mathcal{U}) \in E^{\tau(\mathcal{I},\Delta^{C,O})} \setminus F^{\tau(\mathcal{I},\Delta^{C,O})}$. $(C,\ \mathcal{U}) \in E^{\tau(\mathcal{I},\Delta^{C,O})}$ is equivalent to $\mathcal{I} \models \overline{\mathcal{U}} \sqcap C \sqsubseteq E$ (by Claim 2). Thus, $\neg E \sqcup F$ being a conjunct in $\overline{\mathcal{U}}$ implies $\mathcal{I} \models \overline{\mathcal{U}} \sqcap C \sqsubseteq F$, contradicting $(C,\ \mathcal{U}) \notin F^{\tau(\mathcal{I},\Delta^{C,O})}$, by Claim 2.

CLAIM 3. For Claim 3, the direction

$$\mathcal{I} \models D(a) \implies (a,\ \mathcal{U}) \in D^{\tau(\mathcal{I},\Delta^{C,O})} \tag{7.4}$$

is proven for general $\mathcal{EL}$ concepts $D = \bigsqcap_{1 \leqslant i \leqslant n} A_i \sqcap \bigsqcap_{1 \leqslant j \leqslant m} \exists r_j.E_j$. By the assumption $Qc(D) \subseteq \mathbb{C}$, it holds for $1 \leqslant j \leqslant m$ that $(E_j,\ \emptyset) \in \Delta^{C,O}$.

Therefore, Claim 1 implies $(E_j,\ \emptyset) \in E_j^{\tau(\mathcal{I},\Delta^{C,O})}$ for $1 \leqslant j \leqslant m$. $(a,\ \mathcal{U}) \in (\bigsqcap_{1\leqslant i\leqslant n} A_i)^{\tau(\mathcal{I},\Delta^{C,O})}$ follows again by Definition 7.15. For $1 \leqslant j \leqslant m$, $\mathcal{I} \models (\exists r_j.E_j)(a)$ implies $((a,\ \mathcal{U}),(E_j,\ \emptyset)) \in r_j^{\tau(\mathcal{I},\Delta^{C,O})}$, which, using Claim 1, allows to conclude $(a,\ \mathcal{U}) \in D^{\tau(\mathcal{I},\Delta^{C,O})}$.

We perform the same induction as for $(\star)$ to prove the other direction. For $D = A_1 \sqcap \cdots \sqcap A_n$, Def. 7.15 directly implies $\mathcal{I} \models D(a)$. For the induction hypothesis, assume for m concepts $E_1, \ldots, E_m$ that for any domain element $(b,\ \mathcal{U}') \in \Delta^{C,O}$, $\mathcal{I} \models E_j(b)$ *iff* $(b,\ \mathcal{U}') \in E_j^{\tau(\mathcal{I},\Delta^{C,O})}$ holds. For the induction step, let $D = \bigsqcap_{1\leqslant i\leqslant n} A_i \sqcap \bigsqcap_{1\leqslant j\leqslant m} \exists r_j.E_j$. As before, $\mathcal{I} \models A_i(a)$ $(1 \leqslant i \leqslant n)$ holds by Def. 7.15. For $1 \leqslant j \leqslant m$, $(a,\ \mathcal{U}) \in (\exists r_j.E_j)^{\tau(\mathcal{I},\Delta^{C,O})}$ implies the existence of $((a,\ \mathcal{U}), d) \in r_j^{\tau(\mathcal{I},\Delta^{C,O})}$ and $d \in E_j^{\tau(\mathcal{I},\Delta^{C,O})}$. From Def. 7.15, it is clear that d can either be a concept representative associated with $\emptyset$, or an individual representative.

CASE $d = (F,\ \emptyset)$. By Claim 2, $(F,\ \emptyset) \in E_j^{\tau(\mathcal{I},\Delta^{C,O})}$ holds *iff* $\mathcal{I} \models F \sqsubseteq E_j$. By Def. 7.15, it also holds that $\mathcal{I} \models (\exists r_j.F)(a)$, thus implying $\mathcal{I} \models (\exists r_j.E_j)(a)$.

CASE $d = (b,\ \mathcal{U}')$. For the induction hypothesis, we assumed $\mathcal{I} \models E_j(b)$ *iff* $(b,\ \mathcal{U}') \in E_j^{\tau(\mathcal{I},\Delta^{C,O})}$. Together with $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in r_j^{\mathcal{I}}$ (Def. 7.15), it follows that $\mathcal{I} \models (\exists r_j.E_j)(a)$. □

Remark 7.17. An application of Claim 3 in Lemma 7.16 is $(a,\ \mathcal{U}) \in \overline{\mathcal{U}}^{\tau(\mathcal{I},\Delta^{C,O})}$ *iff* $\mathcal{I} \models \overline{\mathcal{U}}(a)$. Even though $\overline{\mathcal{U}}$ is not an $\mathcal{EL}_\perp$ concept, the "*iff* " in that claim allows to conclude the following. Suppose $\mathcal{I} \models \overline{\mathcal{U}}(a)$, then for all $E \sqsubsetsim F \in \mathcal{U}$ it holds that $\mathcal{I} \not\models E(a)$ or $\mathcal{I} \models F(a)$. Claim 3 of Lem. 7.16 implies that both are equivalent to $(a,\ \mathcal{U}) \notin E^{\tau(\mathcal{I},\Delta^{C,O})}$ and $(a,\ \mathcal{U}) \in F^{\tau(\mathcal{I},\Delta^{C,O})}$, respectively. Hence, showing equivalence of $\mathcal{I} \models \overline{\mathcal{U}}(a)$ and $(a,\ \mathcal{U}) \in \overline{\mathcal{U}}^{\tau(\mathcal{I},\Delta^{C,O})}$.

The transformation in Def. 7.15 is designed carefully, to maintain satisfaction of the strict part of a KB and to construct a typicality interpretation that is standard. Nevertheless, the resulting typicality interpretation is only a standard model of a KB, if the original interpretation (over an arbitrary domain) complies with the defeasible information associated to individual representatives. While this seems like a convoluted prerequisite, the following result is only applied under circumstances that naturally satisfy these requirements, i.e. for models of an extended ABox.

Lemma 7.18. *Let $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ be a KB, $\mathbb{C}, \mathbb{O}$ a quantification closed context containing $\mathcal{K}$, and $\Delta^{C,O}$ some typicality domain with* $\mathtt{TMod}(\mathcal{K}, \Delta^{C,O}) \neq \emptyset$. *For an interpretation $\mathcal{I}$ satisfying* (7.3) *as well as $\mathcal{I} \models \overline{\mathcal{U}}(a)$ for all $(a,\ \mathcal{U}) \in \Delta^{C,O}$, it holds that*

$$\mathcal{I} \in \mathtt{Mod}(\mathcal{K}) \implies \tau(\mathcal{I}, \Delta^{C,O}) \in \mathtt{TMod}(\mathcal{K}, \Delta^{C,O}). \tag{7.5}$$

Proof. From $\mathcal{I} \models \mathcal{A}$, Definition 7.15 (covering role assertions) and Claim 3 of Lemma 7.16 (covering concept assertions) it follows that $\tau(\mathcal{I}, \Delta^{C,O}) \models \mathcal{A}$.

For all $E \sqsubseteq F \in \mathcal{T}$, it holds that

$$\begin{aligned}
&(G, \mathcal{U}) \in E^{\tau(\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})} \\
\implies &\mathcal{I} \models (\overline{\mathcal{U}} \sqcap G) \sqsubseteq E && \text{(Claim 2 of Lem. 7.16)} \\
\implies &\mathcal{I} \models (\overline{\mathcal{U}} \sqcap G) \sqsubseteq F && (\mathcal{I} \models E \sqsubseteq F) \\
\implies &(G, \mathcal{U}) \in F^{\tau(\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})}. && \text{(Claim 2 of Lem. 7.16)}
\end{aligned}$$

Likewise for individual representatives,

$$\begin{aligned}
&(a, \mathcal{U}) \in E^{\tau(\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})} \\
\implies &\mathcal{I} \models E(a) && \text{(Claim 3 of Lem. 7.16)} \\
\implies &\mathcal{I} \models F(a) && (\mathcal{I} \models E \sqsubseteq F) \\
\implies &(a, \mathcal{U}) \in F^{\tau(\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})}. && \text{(Claim 3 of Lem. 7.16)}
\end{aligned}$$

Therefore, $\tau(\mathcal{I}, \Delta^{\mathbb{C},\mathbb{O}})$ satisfies the TBox $\mathcal{T}$. Note that the prerequisite for Claim 2 and 3 of Lem. 7.16 (there: $\mathsf{Qc}(D) \subseteq \mathbb{C}$) is satisfied for E, F, because they appear in the TBox and $\mathbb{C}, \mathbb{O}$ is assumed to contain $\mathcal{K}$, i.e. $\mathsf{Qc}(\mathcal{K}) \subseteq \mathbb{C}$.

Claim 1 of Lem. 7.16 directly provides $(G, \mathcal{U}) \in \overline{\mathcal{U}}^{\tau(\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})}$. Using the argument in Remark 7.17 and the premise for $\mathcal{I}$, we also conclude $(a, \mathcal{U}) \in \overline{\mathcal{U}}^{\tau(\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})}$, hence showing that $\tau(\mathcal{I}, \Delta^{\mathbb{C},\mathbb{O}})$ is a typicality model of $\mathcal{K}$ (Def. 7.7).

Continuing to apply Definition 7.15 and Lemma 7.16 in the same way, allows to conclude that $\tau(\mathcal{I}, \Delta^{\mathbb{C},\mathbb{O}})$ is a standard typicality interpretation (Def. 7.6) and thus, $\tau(\mathcal{I}, \Delta^{\mathbb{C},\mathbb{O}}) \in \mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$. □

We are now ready to prove that simple materialisation[4] ($\mathtt{Mat}^{\mathcal{ALC}}()$) coincides with entailments obtained from the minimal typicality model for representative elements associated with the given set of DCIs. Lemma 7.16 is applied in the following proof, to derive a typicality model counterexample, when an arbitrary counterexample for a materialisation-based entailment exists. Defeasible subsumption and defeasible instance checking is treated separately.

Proposition 7.19. *For a KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, the finite, quantification closed context $\mathbb{C}, \mathbb{O}$ containing $\mathcal{K}$ and some typicality domain $\Delta^{\mathbb{C},\mathbb{O}}$ s.t.* $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \neq \emptyset$, *it holds for all $C, D \in \mathbb{C}$ that*

$$(C, \mathcal{U}) \in D^{\mathtt{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})} \text{ iff } \mathtt{Mat}^{\mathcal{ALC}}((\mathcal{A}, \mathcal{T}, \mathcal{U}), C \mathrel{\underset{\sim}{\sqsubset}} D) \tag{7.6}$$

Proof. We prove both directions of this claim separately, beginning with the *if*-direction. Because $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ is a standard model, it holds that $(C, \mathcal{U}) \in (\overline{\mathcal{U}} \sqcap C)^{\mathtt{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})}$. If $\overline{\mathcal{U}} \sqcap C \sqsubseteq D$ holds for all (classical) models of

4 Relying on $\mathcal{ALC}$ material implications greatly simplifies this proof, because no auxiliary concept names are involved. Theorem 4.22 allows to transfer this result to simple materialisation in $\mathcal{EL}_\bot$.

$(\mathcal{A},\mathcal{T})$, then it holds in particular for $\mathtt{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})$ (due to 1 in Definition 7.7), i.e. $(\mathrm{C},\mathcal{U}) \in \mathrm{D}^{\mathtt{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})}$.

For the *only-if*-direction, consider the intuition that, if a materialisation-based subsumption is not entailed by the classical knowledge base, then there is a standard typicality model as a counterexample. If any such model does not satisfy some subsumption, then the intersection of all standard typicality models (i.e. the minimal typicality model), clearly cannot satisfy the subsumption. In the first step, we need to show that from any arbitrary counterexample, we can derive a standard typicality model counterexample.

Suppose $\overline{\mathcal{U}} \sqcap \mathrm{C} \sqsubseteq \mathrm{D}$ is not entailed by $(\mathcal{A},\mathcal{T})$, hence, there is a model $\mathcal{I}$ of $(\mathcal{A},\mathcal{T})$, s.t. $\exists d \in (\overline{\mathcal{U}} \sqcap \mathrm{C})^{\mathcal{I}} \setminus \mathrm{D}^{\mathcal{I}}$. Take any $\mathcal{J} \in \mathtt{TMod}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})$ (recall: $\mathtt{TMod}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}}) \neq \emptyset$) and inspect the disjoint union $\mathcal{J} \uplus \mathcal{I}$, such that the individuals are interpreted as in $\mathcal{J}$, i.e. $a^{\mathcal{J}\uplus\mathcal{I}} = a^{\mathcal{J}}$. Even though $\Delta^{\mathbb{C},\mathbb{O}}$ is a subdomain of $\Delta^{\mathcal{J}\uplus\mathcal{I}}$, $\mathcal{J} \uplus \mathcal{I}$ can be regarded as an interpretation over an arbitrary domain. The benefit of this union is that some immediate properties of $\mathcal{J}$ are still satisfied for the domain elements of $\Delta^{\mathbb{C},\mathbb{O}}$ within $\mathcal{J} \uplus \mathcal{I}$. To be specific, $\mathcal{J} \uplus \mathcal{I}$ satisfies (7.3) in Lem. 7.16, as well as the prerequisites for the (arbitrary) input interpretation in Lem. 7.18. Furthermore, due to the construction of disjoint unions, it still holds that $\mathcal{J} \uplus \mathcal{I} \not\models (\overline{\mathcal{U}} \sqcap \mathrm{C}) \sqsubseteq \mathrm{D}$ and $\mathcal{J} \uplus \mathcal{I} \models (\mathcal{A},\mathcal{T})$ (cf. Sec. 2.3). It follows from Lemma 7.16 that the typicality interpretation transformation $\tau(\mathcal{J} \uplus \mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})$ does not satisfy $(\mathrm{C},\mathcal{U}) \in \mathrm{D}^{\tau(\mathcal{J}\uplus\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})}$. From Lemma 7.18 we know that $\tau(\mathcal{J} \uplus \mathcal{I},\Delta^{\mathbb{C},\mathbb{O}}) \in \mathtt{TMod}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})$. Therefore, Lemma 7.11, $(\mathrm{C},\mathcal{U}) \notin \mathrm{D}^{\tau(\mathcal{J}\uplus\mathcal{I},\Delta^{\mathbb{C},\mathbb{O}})}$, and Definition 7.13 imply $(\mathrm{C},\mathcal{U}) \notin \mathrm{D}^{\mathtt{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})}$. □

It remains to show the analogous result for defeasible instance checking. As before, this can be shown for general typicality domains. The difference for instance checking, is that defeasible information has to be regarded for all individual representatives, in case they are connected via role edges. In Proposition 7.19, the materialisation-based subsumption used the same set of DCIs that was associated with the concept representative of the typicality domain. More explicitly, in the interpretations that are considered to decide $\overline{\mathcal{U}} \sqcap \mathrm{C} \sqsubseteq \mathrm{D}$, only elements of C that comply with $\mathcal{U}$ are analysed. To comply with defeasible information associated with individual representatives, an appropriate ABox extension is used to restrict the considered classical models in the following result.

Proposition 7.20. *For a KB $\mathcal{K} = (\mathcal{A},\mathcal{T},\mathcal{D})$, the finite, quantification closed context $\mathbb{C},\mathbb{O}$ containing $\mathcal{K}$ and some typicality domain $\Delta^{\mathbb{C},\mathbb{O}}$ s.t. $\mathtt{TMod}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}}) \neq \emptyset$, it holds for all $\mathrm{C} \in \mathbb{C}, a \in \mathbb{O}$ that*

$$(a,\mathcal{U}) \in \mathrm{C}^{\mathtt{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})} \text{ iff } (\mathcal{A} \cup \{\overline{\mathcal{U}'}(b) \mid (b,\mathcal{U}') \in \Delta^{\mathbb{C},\mathbb{O}}\},\mathcal{T}) \models \mathrm{C}(a) \quad (7.7)$$

Proof. For a shorthand, let $\mathcal{K}' = (\mathcal{A} \cup \{\overline{\mathcal{U}'}(b) \mid (b,\mathcal{U}') \in \Delta^{\mathbb{C},\mathbb{O}}\},\mathcal{T})$. We prove both directions of this claim separately. For the *if*-direction, Remark 7.8 contains the arguments that are required to show $\mathtt{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}}) \in \mathtt{Mod}(\mathcal{K}')$, i.e. $\mathtt{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})$ satisfies the classical part of $\mathcal{K}$, and all individual representatives in $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ satisfy their associated set of DCIs. Therefore $\mathcal{K}' \models C(a)$ implies $(a, \mathcal{U}) \in C^{\mathtt{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})}$, because $a^{\mathtt{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})} = (a, \mathcal{U})$.

For the *only-if*-direction, we proceed by contraposition. Assuming $\mathcal{K}' \not\models C(a)$, there is a counterexample $\mathcal{I} \models \mathcal{K}'$ s.t. $a^{\mathcal{I}} \notin C^{\mathcal{I}}$. Using the same argument as in the proof of Proposition 7.19, we know that the disjoint union of $\mathcal{I}$ with a standard typicality model $\mathcal{J}$ of $\mathcal{K}$ over $\Delta^{\mathbb{C},\mathbb{O}}$, i.e. $\mathcal{I} \uplus \mathcal{J}$, is also a model of $\mathcal{K}$ (and $\mathcal{K}'$) and a counterexample to $\mathcal{K}' \models C(a)$. Therefore, $(a, \mathcal{U}) \notin C^{\tau(\mathcal{I}\uplus\mathcal{J},\Delta^{\mathbb{C},\mathbb{O}})}$ follows from Claim 3 of Lemma 7.16. As before, $\tau(\mathcal{I} \uplus \mathcal{J}, \Delta^{\mathbb{C},\mathbb{O}}) \in \mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ can be derived from Lemma 7.18 and thus, the intersection of all models in $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ cannot satisfy $(a, \mathcal{U}) \in C^{\mathtt{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})}$. □

The preceding propositions allow for an immediate comparison of materialisation-based entailments against the entailments produced by typicality model semantics of propositional coverage. To achieve rational and relevant strength, we proceed to construct the specific typicality domains.

7.2.2 *Propositional Rational Reasoning*

For materialisation-based rational reasoning, subsets of the input DBox that are consistent with concepts/individuals, are all selected from the rational chain. It is easy to see that the elements on the rational chain are totally ordered through $\subseteq$. Thus, for subsumption, the largest consistent subset on that chain (w.r.t. the query subject) is unique. For a defeasible instance check $C\{a\}$, the selected DBox subset for a depends also on the sets that have previously been selected for more preferred individuals $b \prec a$. Regardless, the set of DCIs selected for a will be unique for every order $\prec$. From Definition 7.10 and 7.4 it can be seen that if concept and individual representatives adhere to the DBox subsets, selected by $\mathtt{Cons}_{\mathsf{rat}}()$ (or finally $\mathtt{ext}(\mathcal{K}_{\mathsf{rat}_\prec}, a)$), then the standard models over such a typicality domain can produce the same entailments as the materialisation framework instantiated with $\mathtt{rat}$. Recall the notation $\mathtt{ext}(\mathcal{K}_{\mathbf{s}_\prec}, a) = \{\mathcal{U} \subseteq \mathcal{D} \mid \overline{\mathcal{U}}(a) \in \mathcal{A}_{\mathbf{s}_\prec}\}$[5], providing the consistent DBox subsets that have been selected for individuals to obtain the ABox extension (Def. 4.16).

Definition 7.21 (Rational Typicality Domain). For a KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, a finite, quantification closed context $\mathbb{C}$, $\mathbb{O}$, containing $\mathcal{K}$, and a total order $\prec$ over N_I, the *rational typicality domain* $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_\prec}$ is defined as follows

$$\begin{aligned}\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_\prec} = &\{(C,\ \emptyset) \mid C \in \mathbb{C}\}\\ &\cup\{(C,\ \mathtt{Cons}_{\mathsf{rat}}(\mathcal{K}, C)) \mid C \in \mathbb{C}\}\\ &\cup\{(a,\ \mathtt{ext}(\mathcal{K}_{\mathsf{rat}_\prec}, a)) \mid a \in \mathbb{O}\}\end{aligned}$$

5 To be precise, for instantiations of the materialisation framework for $\mathcal{EL}_\bot$, the appropriate concept assertion $\widehat{\mathcal{U}}(a)$ should be used.

	$\mathcal{D}_0$	$\mathcal{D}_1$	$\cdots$	$\emptyset$
C	$\bullet$			$\bullet$
D		$\bullet$		$\bullet$
a	$\bullet$			
$\vdots$				

Figure 7.3: Abstract sketch of a rational typicality domain in matrix shape.

First of all, using Definition 7.21 and $\mathtt{TMod}_{\mathsf{prop}}()$, we obtain a full instantiation of Definition 7.10 as

$$\mathcal{K} \models^{(\mathsf{rat}_\prec,\mathsf{prop})} \alpha \textit{ iff } \forall \mathcal{I} \in \mathtt{TMod}_{\mathsf{prop}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_\prec}).\mathcal{I} \models \alpha \tag{7.8}$$

for defeasible inference problems $(\mathcal{K}, \alpha)$ and a total order $\prec$ over $\mathsf{N_I}$.

Let us inspect rational typicality domains $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_\prec}$ more closely. Consider Figure 7.3 for a sketch of an abstract rational typicality domain, representing at least the concepts C, D and an individual a over the rational chain $\mathtt{chain}(\mathcal{K}) = \langle \mathcal{D}_0, \mathcal{D}_1, \ldots, \emptyset \rangle$. The elements of a rational domain can be arranged in the shape of a matrix. Row labels $(C, D, a, \ldots)$ signify domain elements ($\bullet$) in this row to represent the respective concept or individual. Column labels $(\mathcal{D}_0, \mathcal{D}_1, \ldots, \emptyset)$ describe the associated sets of DCIs for the elements in that column. Notice, how there are at most 2 representatives for each concept in $\mathbb{C}$ and one for each individual in $\mathbb{O}$. Also, every associated set of DCIs is not only a subset of $\mathcal{D}$, but a member of $\mathtt{chain}(\mathcal{K})$. Interestingly enough, the domain of the interpretation in Figure 7.2 (Page 86) is precisely the rational domain of the smallest context containing $\mathcal{K}$ from Example 7.9, albeit not arranged as a matrix.

The ABox extension algorithm can treat individuals from $\mathbb{O} \setminus sig_{\mathsf{I}}(\mathcal{A})$, as explained in Remark 4.17 (last point). Consequently, $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_\prec}$ is a well-defined typicality domain if the context, in particular $\mathbb{O}$, is finite. For every concept $C \in \mathbb{C}$, there is a unique concept representative to check for the entailment of a defeasible subsumption (Def. 7.4), specifically, $(C, \mathtt{Cons}_{\mathsf{rat}}(\mathcal{K}, C))$. While $\mathtt{Cons}_{\mathsf{rat}}(\mathcal{K}, C) = \emptyset$ is entirely possible, it would simply mean that, in this case, only one C-representative is contained in $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_\prec}$. By associating DCIs produced by $\mathtt{Cons}_{\mathsf{rat}}()$ to representative elements, it is ensured that for a $\mathtt{QC}$-consistent KB $\mathcal{K}$ and a finite, consistent[6], quantification closed context $\mathbb{C}$, $\mathbb{O}$, containing $\mathcal{K}$, that $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_\prec})$ can never be empty.

6 Technically, if a consistent context contains a KB $\mathcal{K}$, then $\mathcal{K}$ is automatically $\mathtt{QC}$-consistent. Consistency of the context is the stronger property, allowing to omit $\mathtt{QC}$-consistency.

Proposition 7.19 and 7.20 can be easily applied for entailments based on $\mathtt{TMod}_{\mathtt{prop}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_\prec})$. As in the general case, $\mathtt{TMod}_{\mathtt{prop}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_\prec})$ contains a $\subseteq$-smallest element, the minimal (rational) typicality model $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_\prec})$. Therefore, $\mathcal{K} \models^{(\mathtt{rat}_\prec,\mathtt{prop})} \alpha$ holds *iff* $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_\prec}) \models \alpha$.

Theorem 7.22. *For a KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, a finite, consistent, quantification closed context $\mathbb{C}$, $\mathbb{O}$, containing $\mathcal{K}$, a total order $\prec$ over $\mathrm{N_I}$ and $C, D \in \mathbb{C}$, $a \in \mathbb{O}$, it holds that*

$$\mathcal{K} \models^{(\mathtt{rat}_\prec,\mathtt{prop})} \alpha \text{ iff } \mathcal{K} \models^{(\mathtt{rat}_\prec,\mathtt{mat})} \alpha \tag{7.9}$$

for $\alpha = C \mathrel{\underset{\sim}{\sqsubset}} D$ or $\alpha = C\{a\}$.

Proof. Proving (7.9) boils down to proving the following equivalences

$$\begin{aligned} &(C,\ \mathtt{Cons}_{\mathtt{rat}}(\mathcal{K}, C)) \in D^{\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_\prec})} \\ \textit{iff}\ \ &(\mathcal{A}, \mathcal{T}) \models \overline{\mathtt{Cons}_{\mathtt{rat}}(\mathcal{K}, C)} \sqcap C \sqsubseteq D \end{aligned} \tag{$\star$}$$

and

$$\begin{aligned} &(a,\ \mathtt{ext}(\mathcal{K}_{\mathtt{rat}_\prec}, a)) \in C^{\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_\prec})} \\ \textit{iff}\ \ &(\mathcal{A} \cup \{\overline{\mathtt{ext}(\mathcal{K}_{\mathtt{rat}_\prec}, b)}(b) \mid b \in \mathbb{O}\}, \mathcal{T}) \models C(a) \end{aligned} \tag{$\star\star$}$$

($\star$) is an instance of Proposition 7.19 and ($\star\star$) is an instance of Proposition 7.20. ☐

Theorem 7.22 can be used to directly capture the information contained in the minimal typicality model $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_\prec})$, using materialisation. In Section 8.1, this characterisation of the minimal typicality model is used as a starting point for the algorithmic characterisation of entailments under nested coverage. It has previously appeared in [PT'17a; PT'18] as the definition of the minimal typicality model for propositional rational semantics. Therefore, the earlier approach to reasoning with typicality models was essentially of algorithmic nature. Here, we have introduced an equivalent notion of the minimal (rational) typicality model with a model-theoretic approach, as confirmed by the following corollary.

Corollary 7.23. *Let $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ be a KB, $\mathbb{C}$, $\mathbb{O}$ a consistent, quantification closed context, containing $\mathcal{K}$, $\prec$ a total order over $\mathrm{N_I}$. For the minimal typicality model $\mathcal{I} = \mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_\prec})$ it holds that*

$$\begin{aligned} (C,\ \mathcal{U}) \in A^{\mathcal{I}} &\text{ iff } \mathtt{Mat}^{\mathcal{ALC}}((\mathcal{A}, \mathcal{T}, \mathcal{U}), C \mathrel{\underset{\sim}{\sqsubset}} A) \\ (a,\ \mathcal{U}) \in A^{\mathcal{I}} &\text{ iff } \mathtt{Mat}^{\mathcal{ALC}}(\mathcal{K}_{\mathtt{rat}_\prec}, A(a)) \\ ((C,\ \mathcal{U}), (D,\ \emptyset)) \in r^{\mathcal{I}} &\text{ iff } \mathtt{Mat}^{\mathcal{ALC}}((\mathcal{A}, \mathcal{T}, \mathcal{U}), C \mathrel{\underset{\sim}{\sqsubset}} \exists r.D) \\ ((a,\ \mathcal{U}), (D,\ \emptyset)) \in r^{\mathcal{I}} &\text{ iff } \mathtt{Mat}^{\mathcal{ALC}}(\mathcal{K}_{\mathtt{rat}_\prec}, (\exists r.D)(a)) \\ ((a,\ \mathcal{U}), (b,\ \mathcal{U}')) \in r^{\mathcal{I}} &\text{ iff } r(a, b) \in \mathcal{A} \end{aligned}$$

7.2.3 *Propositional Relevant Reasoning*

It would be naive to construct the relevant typicality domain in the same way as the rational domain, exchanging only $\mathtt{Cons}_{\mathtt{rat}}()$ with $\mathtt{Cons}_{\mathtt{rel}}()$. To understand why, recall the intuition on preferring sets of typicality models, which will allow to derive role edges with a more typical successor in terms of associated DCIs. Practically, if an edge $r(d, (D, \emptyset))$ is satisfied by all typicality models, then those models that also satisfy $r(d, (D, \mathcal{E}))$ (for some $\emptyset \subsetneq \mathcal{E}$) will be preferred, if there is at least one. For rational reasoning, it suffices to consider the most typical D representative in this case, because the intermediate ones (if associated with elements from $\mathtt{chain}(\mathcal{K})$) would hold the same information as $(D, \emptyset)$. This effect is consequential of the way $\mathtt{chain}(\mathcal{K})$ is defined. For relevant reasoning such an effect does not take hold. In particular, representatives associated with DCIs $\mathcal{E}$ between $\mathtt{Cons}_{\mathtt{rat}}(\mathcal{K}, D) \subsetneq \mathcal{E} \subsetneq \mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, D)$, may each hold distinct information. If only the least and most typical representatives were considered, the core idea of fine granularity would be lost for deriving typical information about role successors. To maintain this granularity, the relevant domain must contain concept representatives associated with every possible DBox subset between $\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, D)$ and $\emptyset$.

Definition 7.24 (Relevant Typicality Domain)**.** For a KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, a finite, quantification closed context $\mathbb{C}$, $\mathbb{O}$, containing $\mathcal{K}$, and a total order $\prec$ over N_I, let $\mathcal{K}_{\mathtt{rel}_\prec}$ be the extended KB produced by Def. 4.16 for the individuals in $\mathbb{O}$. The *relevant typicality domain* $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec}$ is defined as follows

$$\begin{aligned}\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec} = &\{(C, \mathcal{U}) \mid C \in \mathbb{C}, \mathcal{U} \subseteq \mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, C)\}\\ &\cup\{(a, \mathtt{ext}(\mathcal{K}_{\mathtt{rel}_\prec}, a)) \mid a \in \mathbb{O}\}\end{aligned}$$

Analogous to rational strength, we obtain a full instantiation of Definition 7.10, using Definition 7.24 and $\mathtt{TMod}_{\mathtt{prop}}()$, as

$$\mathcal{K} \models^{(\mathtt{rel}_\prec,\mathtt{prop})} \alpha \textit{ iff } \forall \mathcal{I} \in \mathtt{TMod}_{\mathtt{prop}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec}). \mathcal{I} \models \alpha \tag{7.10}$$

for defeasible inference problems $(\mathcal{K}, \alpha)$ and a total order $\prec$ over N_I. Note that to verify $\mathcal{I} \models \alpha$, there is also a unique concept/individual representative to check. It is easy to see that, for $\alpha = C \sqsubsetsim D$, $\mathcal{I} \models \alpha$ (Def. 7.4) holds *iff* $(C, \mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, C)) \in D^{\mathcal{I}}$.

Consider Figure 7.4 for an abstract illustration of a relevant domain. The relevant domain can be arranged in the shape of the full subset lattice of the given DBox (in Fig. 7.4 we assume $|\mathcal{D}| = 3$ to allow for a small illustration). Elements are grouped in boxes that describe the associated set of DCIs, according to its position in the subset lattice of $\mathcal{D}$. First of all, any concept $C \in \mathbb{C}$ (e.g. $\bullet, \circ$ in Fig. 7.4) is clearly consistent with $\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, C)$ and therefore must be consistent with every subset of $\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, C)$. Having representatives associated with *any* consistent subset of the DBox allows to prefer models satisfying defeasible information at role successors with the same granularity as it is achieved with materialisation for the top level of

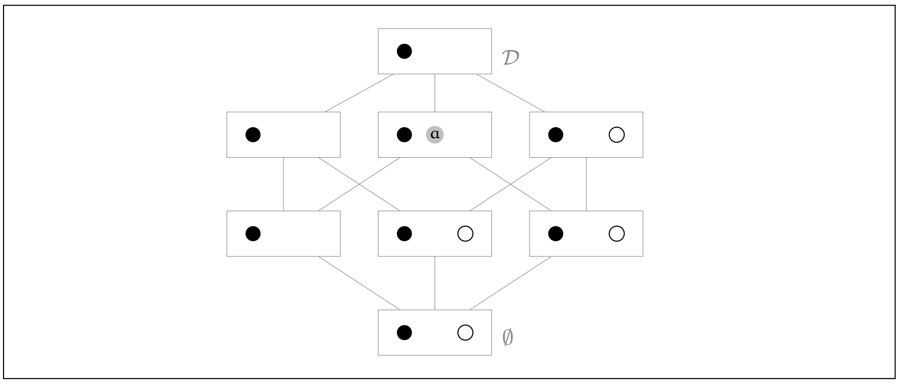

Figure 7.4: Abstract sketch of a relevant typicality domain in lattice shape.

a query concept. For individuals (e.g. a in Fig. 7.4), the relevant domain must still contain only one unique representative, which is of course, also associated to the "relevant amount" of defeasible knowledge. By only associating subsets of the DCIs produced by $\mathtt{Cons_{rel}}()$ to representative elements, it is ensured for a finite, consistent, quantification closed context $\mathbb{C}$, $\mathbb{O}$, containing $\mathcal{K}$, that $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec})$ can never be empty.

The implications of Proposition 7.19 and 7.20 for entailments based on $\mathtt{TMod_{prop}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec})$ are analogous to those for the rational domain. As in the general case, $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec}) \subseteq \mathcal{J}$ for all $\mathcal{J} \in \mathtt{TMod_{prop}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec})$. Therefore, $\mathcal{K} \models^{(\mathtt{rel}_\prec,\mathtt{prop})} \alpha$ holds *iff* $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec}) \models \alpha$ holds. We obtain the following explicit result as a consequence of the more general Proposition 7.19 and 7.20.

Theorem 7.25. *For a KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, a finite, consistent, quantification closed context $\mathbb{C}$, $\mathbb{O}$, containing $\mathcal{K}$, a total order $\prec$ over $\mathrm{N_I}$ and $\mathrm{C}, \mathrm{D} \in \mathbb{C}$, $a \in \mathbb{O}$, it holds that*

$$\mathcal{K} \models^{(\mathtt{rel}_\prec,\mathtt{prop})} \alpha \text{ iff } \mathcal{K} \models^{(\mathtt{rel}_\prec,\mathtt{mat})} \alpha \tag{7.11}$$

for $\alpha = \mathrm{C} \sqsubsim \mathrm{D}$ or $\alpha = \mathrm{C}\{a\}$.

Proof. Proving Eq. 7.11 requires to prove the following equivalences

$$\begin{aligned} &(\mathrm{C},\ \mathtt{Cons_{rel}}(\mathcal{K}, \mathrm{C})) \in \mathrm{D}^{\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec})} \\ \textit{iff}\ \ &(\mathcal{A}, \mathcal{T}) \models \overline{\mathtt{Cons_{rel}}(\mathcal{K}, \mathrm{C})} \sqcap \mathrm{C} \sqsubseteq \mathrm{D} \end{aligned} \tag{$\star$}$$

and

$$\begin{aligned} &(a,\ \mathtt{ext}(\mathcal{K}_{\mathtt{rel}_\prec}, a)) \in \mathrm{C}^{\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec})} \\ \textit{iff}\ \ &(\mathcal{A} \cup \{\overline{\mathtt{ext}(\mathcal{K}_{\mathtt{rel}_\prec}, b)}(b) \mid b \in \mathbb{O}\}, \mathcal{T}) \models \mathrm{C}(a) \end{aligned} \tag{$\star\star$}$$

($\star$) is proven with Proposition 7.19 and ($\star\star$) is proven with Proposition 7.20. □

The direct computation of the minimal typicality model over the rational domain (Cor. 7.23) can be trivially adopted to the relevant domain.

Correctness of this direct characterisation is an immediate consequence of Theorem 7.25. Note that even though Corollary 7.23 relies on classical reasoning (i.e. is polynomial for $\mathcal{EL}_\bot$), the size of the relevant domain is exponential in the size of the input DBox. Therefore an exponential number of classical entailment checks are required to compute $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec})$. While this is one approach to determine entailments under propositional relevant semantics, we refer to Chapter 8 for a tighter complexity result on (the equivalent) materialisation-based relevant reasoning.

7.3 PREFERRED SETS OF TYPICALITY MODELS

The downside of propositional coverage, i.e. the use of *all* models of the KB, is directly inherited from its equivalence to materialisation-based reasoning. Many expected entailments about quantified concepts cannot be inferred by considering all typicality models or using materialisation. The intuition on the cause of quantification neglect has been presented in Section 5.2 for the latter. In terms of typicality models, the conservative requirement for role-successor witnesses (Property 3 in Def. 7.6) to be associated with no DCIs, is the cause of this drawback. Intuitively, standard typicality models are not obligated to satisfy defeasible information at role successor elements. As a result, the only role edges that persist in *all* typicality models, are those with a non-typical successor (role edges between individuals aside). The missing entailments can be obtained by restricting the set of considered models to those satisfying appropriate DCIs at role successors. Such a restriction, however, is not trivially obtained, mainly for the following reason.

> Defeasible information can be satisfied at two different role-successors individually, but not necessarily at the same time. (7.12)

Therefore, there is no unique subset of all models of a KB, that is "as typical as possible" in terms of role successors satisfying as many DCIs as possible. In this section we introduce a preference relation on sets of typicality models, to define two different types of nested semantics using the *most preferred* sets of typicality models.

SCEPTICAL. Satisfaction of an inference problem is determined by *all* most preferred sets of models.

SELECTIVE. Satisfaction of an inference problem is determined by *a chosen* most preferred set of models.

The former has been introduced in [PT'17a] by an algorithmic construction of so-called *maximal typicality models*. In [PT'17a] we proposed to iteratively extend the minimal typicality model by role edges with more typical representatives for successors, until "maximal typicality" is achieved. The present approach achieves (almost[7]) the same entailments under sceptical nested semantics, but remains a model-theoretic characterisation.

7 Some refinements of the approach in [PT'17a; PT'17b] are introduced here. Motivation for the need of these refinements is given in Section 7.3.1.

As in Section 7.2, the notions required to introduce preference of sets of typicality models are introduced for generic typicality domains. After defining both nested semantics in general, we will explicitly discuss rational and relevant strength. The relation between both types of nested semantics is covered during the full overview of all of the discussed semantics in Chapter 9.

7.3.1 *Preference Options*

To understand the lack of defeasible information at role successors in entailments based on *all* typicality models, consider the implications of Definition 7.13 (Minimal Typicality Model) and Corollary 7.23 (direct construction of the minimal typicality model). The minimal typicality model only contains role successors to concept representatives that are associated with an empty set of DCIs. This makes the minimal typicality model an immediate counterexample for expected entailments such as $A \mathrel{\underset{\sim}{\sqsubset}} B \in \mathcal{D}$ implying $\exists r.A \mathrel{\underset{\sim}{\sqsubset}} \exists r.B$.

Example 7.26. Recall the KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ from Example 7.9 with

$$\begin{aligned}\mathcal{A} &= \{(\mathrm{Cat} \sqcap \exists \mathrm{friend}.\mathrm{Dog})(\mathrm{molli})\}\\ \mathcal{T} &= \{\mathrm{Cat} \sqsubseteq \exists \mathrm{likes}.\mathrm{Cat}\}\\ \mathcal{D} &= \{\mathrm{Dog} \mathrel{\underset{\sim}{\sqsubset}} \mathrm{Happy}, \mathrm{Cat} \mathrel{\underset{\sim}{\sqsubset}} \mathrm{Lazy}\}\end{aligned}$$

and the accompanying typicality model $\mathcal{J} = (\Delta^{\mathbf{C},\mathbf{O}}, \cdot^{\mathcal{J}})$ in Figure 7.2 (Page 86). As established before, $\mathcal{J}$ strictly extends the minimal typicality model for $\Delta^{\mathbf{C},\mathbf{O}}$ and $\mathcal{K}$, because the edge $\mathrm{friend}((\mathrm{molli}, \mathcal{D}), (\mathrm{Dog}, \mathcal{D}))$ is superfluous. At the same time, $\mathcal{J} \in \mathtt{TMod}(\mathcal{K}, \Delta^{\mathbf{C},\mathbf{O}})$ shows that it is plausible (formally: consistent), to expect the "more typical version" of the edge $\mathrm{friend}((\mathrm{molli}, \mathcal{D}), (\mathrm{Dog}, \emptyset))$, which is contained in all models in $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbf{C},\mathbf{O}})$. Intuitively, if this plausibility exists, a set of typicality models is preferred, if all its elements satisfy the more typical role edge. In that sense, the preference between two sets of typicality models is witnessed by some role edge. Note that it is important only to consider role edges whose non-typical version (i.e. successor associated with $\emptyset$) is already inferred, to avoid the support of arbitrary inferences.

Example 7.26 motivates the use of role edges with *typical* concept representatives for successors, as means to prefer one set of typicality models over another.

Roughly speaking, a set of models $\mathcal{M}_1$ is preferred over another set of models $\mathcal{M}_2$, if all interpretations in $\mathcal{M}_1$ satisfy a role edge that is not satisfied by some model in $\mathcal{M}_2$. As done intuitively in Example 7.26, we do not consider arbitrary such role edges, but only those that yield the desired effect and have an appropriate explanation to be inferred. The eligible role edges are classified as the *preference options* in a typicality domain.

Definition 7.27 (Preference Option). For a typicality domain $\Delta^{\mathbb{C},\mathbb{O}}$, some $r \in N_R$ and two elements $(\chi, \mathcal{U}), (F, \mathcal{U}') \in \Delta^{\mathbb{C},\mathbb{O}}$ $(\chi \in \mathbb{C} \cup \mathbb{O}, F \in \mathbb{C})$, the role edge $r((\chi, \mathcal{U}), (F, \mathcal{U}'))$ is a *preference option* in $\Delta^{\mathbb{C},\mathbb{O}}$ *iff*

1. $\mathcal{U}' \neq \emptyset$, (Typical Successor)
2. $\chi \in \mathbb{C} \implies \mathcal{U} \neq \emptyset$. (Typical Predecessor)

The set of all *preference options* over $\Delta^{\mathbb{C},\mathbb{O}}$ is $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$.

For a set of typicality interpretations $\mathcal{M}$ over $\Delta^{\mathbb{C},\mathbb{O}}$, a preference option $r(d, (F, \mathcal{U}')) \in \mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ is *admissible* for $\mathcal{M}$ *iff*

1. $\forall \mathcal{I} \in \mathcal{M}.(d, (F, \emptyset)) \in r^{\mathcal{I}}$, (Justified)
2. $\exists \mathcal{I} \in \mathcal{M}.(d, (F, \mathcal{U}')) \notin r^{\mathcal{I}}$, (Not-entailed)
3. $\exists \mathcal{I} \in \mathcal{M}.(d, (F, \mathcal{U}')) \in r^{\mathcal{I}}$. (Satisfiable)

When considering a specific KB $\mathcal{K}$ and its sets of typicality models, we assume the roles used in $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ to be finitely bound by $sig_R(\mathcal{K})$.

The properties that make role edges over a typicality domain into (admissible) preference options are selected with great care, as they determine certain effects in the resulting semantics.

TYPICAL SUCCESSOR. First of all, we do not consider role edges with individual representatives as the successor. Any such role edge can only ever be inferred from ABox role assertions due to the nature of standard typicality models. Because only one individual representative per individual in $\mathbb{O}$ is present, these edges are as typical as possible. What is still missing in the minimal typicality model, is defeasible information in (anonymous) concept representative successors, independent of the predecessor of a role edge. Thus, only those role edges are considered, that allow to capture a set of models satisfying a non-empty set of DCIs at some role successor.

TYPICAL PREDECESSOR. While the type of representative is irrelevant for the predecessor of a preference option, it is prohibited that this domain element is a non-typical concept representative. This is a refinement to the approach in [PT'18], where arbitrary domain elements are allowed as predecessors of a *more typical role edge* (cf. Def. 4.22 in [PT'18]). The following example illustrates the intuition why the successors of non-typical concept representatives should also be non-typical.

Example 7.28. Consider the extension $\mathcal{J}$ of the minimal typicality model $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ (for $\mathcal{K}$ and $\Delta^{\mathbb{C},\mathbb{O}}$ as in Example 7.9/7.26), depicted in Figure 7.5. Note the (typical) edge $((Cat, \emptyset), (Cat, \mathcal{D})) \in likes^{\mathcal{J}}$. If only models extending $\mathcal{J}$ are considered to determine entailments, then for typical Cats ($(Cat, \mathcal{D})$) it would be derived that their immediate $likes$ successors are not $Lazy$, but every even chain of $likes$ edges (strictly longer than 1) would point to a $Lazy$ Cat. Thus, it is implied that typical Cats like non-typical Cats, which in turn like typical Cats, somehow conflicting

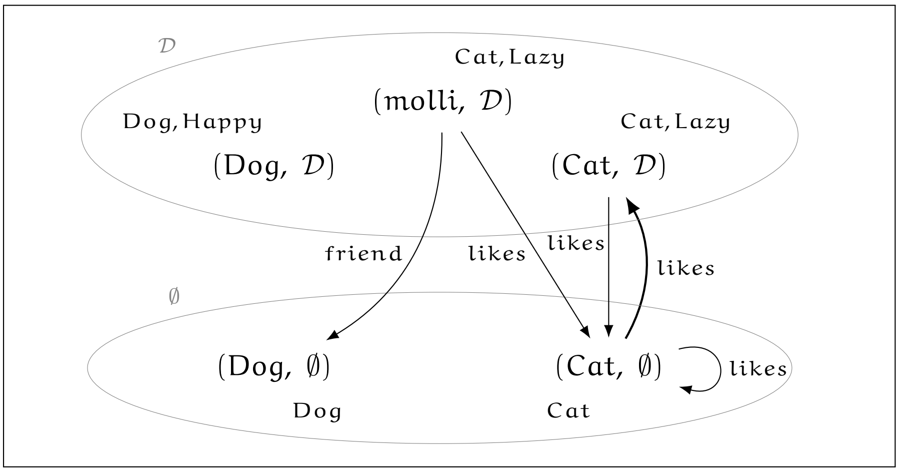

Figure 7.5: An extension of $\mathtt{M}(\mathcal{K}, \Delta^{\mathsf{C},\mathsf{O}})$ (cf. Example 7.9), with an undesired more typical role edge.

with their lack of typicality. Even worse, if the other two likes edges ($\mathsf{likes}((\mathsf{molli},\ \mathcal{D}), (\mathsf{Cat},\ \mathcal{D}))$, $\mathsf{likes}((\mathsf{Cat},\ \mathcal{D}), (\mathsf{Cat},\ \mathcal{D}))$) would be inferred for a set of models extending $\mathcal{J}$, then, even the classical subsumption $\mathsf{Cat} \sqsubseteq \exists \mathsf{likes}.\mathsf{Lazy}$ could be inferred from this set of models. Therefore, the intuition is that, if an element is non-typical, it may only be related to other non-typical elements.

ADMISSIBLE PREFERENCE OPTION. Admissibility of a preference option is determined with respect to a set of typicality models. It is used to identify role edges capturing "holes" in a set of models, that (if metaphorically "filled") could lead to a more refined set of models in terms of overall DCIs satisfied at role successors, in the following sense.

1. No arbitrary role edge should be deduced. (Justified)
2. The preference option is not already satisfied in all members of the present set of models. (Not-entailed)
3. The entailment of the preference option is possible in at least one of the present models. (Satisfiable)

The third aspect of the admissible property is used in particular to determine maximality of sets of models within the preference relation. It is also crucial for detecting potential *conflicts*, prohibiting two role edges with typical successors to be satisfied at the same time (i.e. in the same model), as illustrated shortly, in Example 7.34.

7.3.2 *Nested Typicality Preference*

The intuitively described interaction of (admissible) preference options and the preference relation over sets of typicality models is made explicit by formalising the refinement of a set of interpretations to those satisfying a given role edge. Such a refinement operation relates two sets of interpretations over an arbitrary but shared domain. Therefore, it is not defined explicitly for typicality models.

Definition 7.29 (Interpretation-Set Refinement)**.** For a set of interpretations $\mathcal{M}$ over a (common) domain Δ and a role edge, $r(d, e)$ ($r \in N_R$, $d, e \in \Delta$), $\mathcal{M}$ is *refined* by $r(d, e)$ to

$$\mathcal{M}|_{r(d,e)} = \{\mathcal{I} \in \mathcal{M} \mid (d, e) \in r^{\mathcal{I}}\}$$

Interpretation-set refinements are used on sets of typicality models, to single out a preference option that is the witness to one set of typicality models being preferred over another.

Definition 7.30 (Typicality Preference)**.** For two sets of typicality interpretations $\mathcal{M}_1$, $\mathcal{M}_2$ over the same typicality domain $\Delta^{C,O}$, the *typicality preference* relation $<_t$ it defined such that

$\mathcal{M}_1 <_t \mathcal{M}_2$ *iff* $\mathcal{M}_2 = \mathcal{M}_1|_p$ for some $p \in \mathtt{PO}(\Delta^{C,O})$, admissible for $\mathcal{M}_1$.

A chain $\mathcal{M}_1 <_t \mathcal{M}_2 <_t \cdots <_t \mathcal{M}_n$ is called a *preference chain*. Such a chain is called

- *maximal iff* there does not exist a set of typicality interpretations $\mathcal{M}$ s.t. $\mathcal{M}_n <_t \mathcal{M}$, and
- *full iff* it is maximal and $\mathcal{M}_1 = \mathtt{TMod}(\mathcal{K}, \Delta^{C,O})$ (for a KB $\mathcal{K}$).

If $\mathcal{M}_1 <_t \mathcal{M}_2$ we say $\mathcal{M}_2$ is more preferred than $\mathcal{M}_1$. This notion of preference is incremental in terms of the number of satisfied role edges. A preference chain $\mathcal{M}_1 <_t \mathcal{M}_2 <_t \cdots$ exhibits an increase (at least) in role edges that are satisfied in all models in $\mathcal{M}_i$ (with increasing $i \geqslant 1$). Note the following immediate observations for nested typicality preference.

Remark 7.31.

1. If $\mathcal{M}_1 = \emptyset$, there are no admissible preference options for $\mathcal{M}_1$. If $\mathcal{M}_1 \neq \emptyset$ and $\mathcal{M}_1 <_t \mathcal{M}_2$, then $\mathcal{M}_2$ can also not be empty (every admissible preference option is Satisfiable in $\mathcal{M}_1$; Def. 7.27).
2. $\mathcal{M}_1 <_t \mathcal{M}_2$ implies $\mathcal{M}_2 \subset \mathcal{M}_1$. $\mathcal{M}_2 \subseteq \mathcal{M}_1$ follows promptly from Def. 7.29. $\mathcal{M}_2 \neq \mathcal{M}_1$ holds because an admissible preference option is Not-entailed (Def. 7.27), i.e. admissibility requires existence of some interpretation in $\mathcal{M}_1$ that cannot remain in $\mathcal{M}_2$. The inclusion $\mathcal{M}_2 \subset \mathcal{M}_1$ provides the desired effect of *refining* supported inferences, because $\mathcal{M}_2$ retains all consequences that are entailed by $\mathcal{M}_1$.
3. When considering only a finite signature and a finite typicality domain, the set of *all* standard typicality interpretations over said domain and signature, is finite. In that case, the previous observation and Def. 7.27 imply that for any set of standard typicality interpretations $\mathcal{M}$, there are only finitely many $\mathcal{M}'$ s.t. $\mathcal{M} <_t \mathcal{M}'$ and there are no infinite preference chains $\mathcal{M} <_t \mathcal{M}_1 <_t \cdots$.

We discuss three effects emerging from Definition 7.30, to understand its benefits and difficulties before defining explicit semantics based on this preference. The three subsequent examples illustrate the benefit of typicality preference (Exm. 7.32), a side-effect to be aware of (Exm. 7.33) and a situation in need of explicit attention (Exm. 7.34), in that order.

Example 7.32. Continuing on Example 7.26, it can be readily seen that $p = \mathit{friend}((\mathit{molli},\ \mathcal{D}), (\mathit{Dog},\ \mathcal{D}))$ is an admissible preference option for $\mathtt{TMod}(\mathcal{K}, \Delta^{C,O})$. It clearly satisfies the criteria to be a preference option in $\Delta^{C,O}$ and it is Satisfiable (witnessed by $\mathcal{J} \in \mathtt{TMod}(\mathcal{K}, \Delta^{C,O})$ in Fig. 7.2), Not-entailed, and Justified (both seen from $\mathtt{M}(\mathcal{K}, \Delta^{C,O})$, which is simply $\mathcal{J}$ without $\mathit{friend}((\mathit{molli},\ \mathcal{D}), (\mathit{Dog},\ \mathcal{D}))$). It is also not hard to see that $\mathcal{J}$ is the smallest standard typicality model extending $\mathtt{M}(\mathcal{K}, \Delta^{C,O})$ and satisfying p. Therefore, $\mathtt{TMod}(\mathcal{K}, \Delta^{C,O}) <_t \mathtt{TMod}(\mathcal{K}, \Delta^{C,O})|_p$ and

$$\mathtt{TMod}(\mathcal{K}, \Delta^{C,O})|_p = \{\mathcal{I} \in \mathtt{TMod}(\mathcal{K}, \Delta^{C,O}) \mid \mathcal{J} \subseteq \mathcal{I}\}.$$

Consequently, from all models in $\mathtt{TMod}(\mathcal{K}, \Delta^{C,O})|_p$ it can be concluded that (defeasibly) molli has a Happy friend.

By construction, $\mathcal{M}_1 <_t \mathcal{M}_2$ has (at least) one witness preference option $r(d, e)$. It is not necessarily the case, that $(d, e) \in r^{\mathcal{I}}$ is the only "additional" property that all $\mathcal{I} \in \mathcal{M}_2$ share, as illustrated in the following example.

Example 7.33. Extend the KB from Example 7.9 into

$$\mathcal{K}' = (\mathcal{A}, \mathcal{T} \cup \{\exists \mathit{friend}.\mathit{Happy} \sqsubseteq \exists \mathit{friend}.\mathit{Cat}\}, \mathcal{D}).$$

Recall that $\mathcal{J}$ in Figure 7.2 only extends the minimal typicality model $\mathtt{M}(\mathcal{K}, \Delta^{C,O})$ by the edge $p = \mathit{friend}((\mathit{molli},\ \mathcal{D}), (\mathit{Dog},\ \mathcal{D}))$. Because no element in $\mathtt{M}(\mathcal{K}, \Delta^{C,O})$ satisfies $\exists \mathit{friend}.\mathit{Happy}$, it trivially satisfies $\exists \mathit{friend}.\mathit{Happy} \sqsubseteq \exists \mathit{friend}.\mathit{Cat}$. This effectively shows that $\mathtt{M}(\mathcal{K}', \Delta^{C,O}) = \mathtt{M}(\mathcal{K}, \Delta^{C,O})$ and the defeasible consequences under propositional coverage for $\mathcal{K}$ and $\mathcal{K}'$ coincide.[8] However, using the admissible preference option p for $\mathtt{TMod}(\mathcal{K}', \Delta^{C,O})$ now results in a set $\mathtt{TMod}(\mathcal{K}', \Delta^{C,O})|_p$ that no longer contains $\mathcal{J}$, because $\mathcal{J}$ does not satisfy the new GCI. A small extension of $\mathcal{J}$, that does satisfy $\mathcal{K}'$ is easy to imagine. Extending $\mathcal{J}$ by the edge $\mathit{friend}((\mathit{molli},\ \mathcal{D}), (\mathit{Cat},\ \emptyset))$ does in fact yield the smallest member $\mathcal{J}' \in \mathtt{TMod}(\mathcal{K}', \Delta^{C,O})|_p$. There are two observations worth noting.

1. Expecting molli to have a typical Dog for a friend (by typicality preference), results in the implicit conclusion that molli also befriends a Cat.

2. With this typicality preference, there is a new admissible preference option (i.e. $\mathit{friend}((\mathit{molli},\ \mathcal{D}), (\mathit{Cat},\ \emptyset))$) that could be used to incrementally refine $\mathtt{TMod}(\mathcal{K}', \Delta^{C,O})|_p$.

8 With this extension to $\mathcal{K}'$, the context no longer contains the KB and several results from the previous sections would fail. However, entailment w.r.t. $\mathtt{TMod}(\mathcal{K}', \Delta^{C,O})$ and subsets thereof remains a valid notion. Only for illustrations sake, we continue referring to the minimal typicality model and put up with being partially imprecise in this (and the following) example.

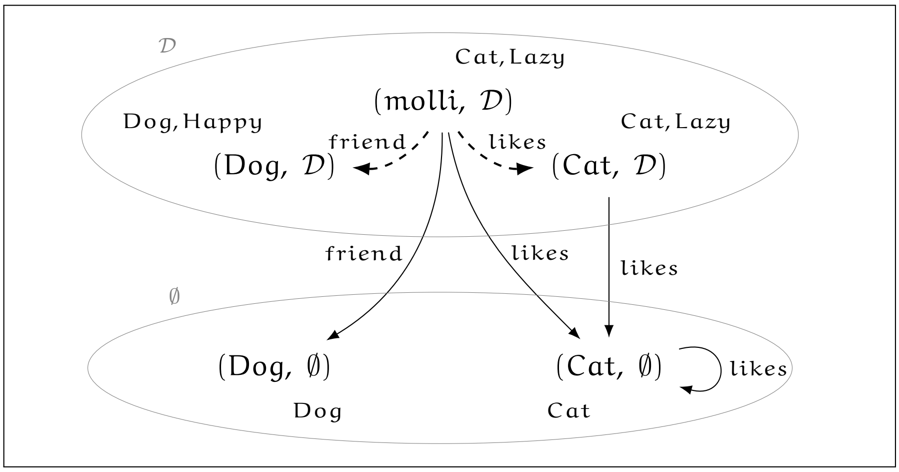

Figure 7.6: An extension of $\mathtt{M}(\mathcal{K}, \Delta^{\mathsf{C},\mathsf{O}})$ (cf. Example 7.9), violating the additional GCI $\exists\mathsf{friend}.\mathsf{Happy} \sqcap \exists\mathsf{likes}.\mathsf{Lazy} \sqsubseteq \bot$.

The effects described in Example 7.33 are implicit with the nested typicality preference, because we are only dealing with sets of *models* of a given KB. In [PT'17a; PT'17b; PT'18] these effects motivated an explicit, iterative construction (Sec. 4.3 in [PT'18]) of *typicality extensions* (add an admissible preference option to the current model) and *model completions* (extend the current interpretation minimally, to reach a model of the KB). This kind of algorithmic construction will be covered in Section 8.1, in the effort to capture a complexity upper bound for deciding entailments in semantics based on typicality preference.

Example 7.33 shows that the preference relation $<_{\mathsf{t}}$ needs only to be concerned with the refinement of models via admissible preference options. Implicit information (from the KB) and the property of admissibility for all preference options is somewhat "redetermined automatically" after every refinement. Nevertheless, finding a maximally preferred set of typicality models is not linear, because of the effect described in (7.12). Intuitively, two preference options can be in conflict with each other, this is illustrated with another small extension of the initial Example 7.9.

Example 7.34. Let

$$\mathcal{K}'' = (\mathcal{A}, \mathcal{T} \cup \{\exists\mathsf{friend}.\mathsf{Happy} \sqcap \exists\mathsf{likes}.\mathsf{Lazy} \sqsubseteq \bot\}, \mathcal{D}).$$

As in the previous examples, $\mathsf{p}_1 = \mathsf{friend}((\mathsf{molli}, \mathcal{D}), (\mathsf{Dog}, \mathcal{D}))$ is an admissible preference option for $\mathtt{M}(\mathcal{K}'', \Delta^{\mathsf{C},\mathsf{O}})$ (which again, for argument's sake, we consider to coincide with $\mathtt{M}(\mathcal{K}, \Delta^{\mathsf{C},\mathsf{O}})$). Another admissible preference option for $\mathtt{TMod}(\mathcal{K}'', \Delta^{\mathsf{C},\mathsf{O}})$ is $\mathsf{p}_2 = \mathsf{likes}((\mathsf{molli}, \mathcal{D}), (\mathsf{Cat}, \mathcal{D}))$. In fact, p_2 was admissible already for $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathsf{C},\mathsf{O}})$. Again, as in Example 7.32, refining $\mathtt{TMod}(\mathcal{K}'', \Delta^{\mathsf{C},\mathsf{O}})$ with p_1 allows to conclude $(\exists\mathsf{friend}.\mathsf{Happy})\{\mathsf{molli}\}$, but, precisely this consequence prohibits p_2 from being admissible for $\mathtt{TMod}(\mathcal{K}'', \Delta^{\mathsf{C},\mathsf{O}})|_{\mathsf{p}_1}$, as illustrated in Figure 7.6. $\mathcal{J}$ (Fig. 7.2) is still the smallest model in $\mathtt{TMod}(\mathcal{K}'', \Delta^{\mathsf{C},\mathsf{O}})|_{\mathsf{p}_1}$, and we can show that $\mathtt{TMod}(\mathcal{K}'', \Delta^{\mathsf{C},\mathsf{O}})|_{\mathsf{p}_1}$ contains no interpretation that satisfies p_2. All interpretations $\mathcal{I} \in$

$\texttt{TMod}(\mathcal{K}'', \Delta^{\mathbb{C},\mathbb{O}})|_{p_1}$ extend $\mathcal{J}$ (i.e. $\mathcal{J} \subseteq \mathcal{I}$) and satisfy $\mathcal{K}''$, in particular, $\mathcal{T} \cup \{\exists \mathit{friend}.\mathit{Happy} \sqcap \exists \mathit{likes}.\mathit{Lazy} \sqsubseteq \bot\}$. It is not hard to verify that every interpretation extending $\mathcal{J}$ by p_2 (e.g. Fig. 7.6) cannot satisfy $\exists \mathit{friend}.\mathit{Happy} \sqcap \exists \mathit{likes}.\mathit{Lazy} \sqsubseteq \bot$, due to the counterexample $(\mathit{molli}, \mathcal{D})$. Ergo, p_2 is not Satisfiable (cf. Def. 7.27) w.r.t. $\texttt{TMod}(\mathcal{K}'', \Delta^{\mathbb{C},\mathbb{O}})|_{p_1}$.

On the other hand, p_2 is Satisfiable w.r.t. $\texttt{TMod}(\mathcal{K}'', \Delta^{\mathbb{C},\mathbb{O}})$, showing that

$$\texttt{TMod}(\mathcal{K}'', \Delta^{\mathbb{C},\mathbb{O}}) <_t \texttt{TMod}(\mathcal{K}'', \Delta^{\mathbb{C},\mathbb{O}})|_{p_1} \tag{7.13}$$

as well as

$$\texttt{TMod}(\mathcal{K}'', \Delta^{\mathbb{C},\mathbb{O}}) <_t \texttt{TMod}(\mathcal{K}'', \Delta^{\mathbb{C},\mathbb{O}})|_{p_2} \tag{7.14}$$

but neither

$$\texttt{TMod}(\mathcal{K}'', \Delta^{\mathbb{C},\mathbb{O}})|_{p_1} <_t (\texttt{TMod}(\mathcal{K}'', \Delta^{\mathbb{C},\mathbb{O}})|_{p_1})|_{p_2}$$

nor

$$\texttt{TMod}(\mathcal{K}'', \Delta^{\mathbb{C},\mathbb{O}})|_{p_2} <_t (\texttt{TMod}(\mathcal{K}'', \Delta^{\mathbb{C},\mathbb{O}})|_{p_2})|_{p_1}$$

are true. Simply put, when $\mathcal{M}_1 <_t \mathcal{M}_2$, an admissible preference option for $\mathcal{M}_1$ need not be admissible for $\mathcal{M}_2$.

Example 7.34 shows that the transitive closure of $<_t$ is not confluent. Formally, $\mathcal{M} <_t \mathcal{M}_1$ and $\mathcal{M} <_t \mathcal{M}_2$ does not imply that there is some $\mathcal{M}'$ s.t. $\mathcal{M}_i <_t \ldots <_t \mathcal{M}'$ ($i \in \{1, 2\}$). This means that in general, there are multiple distinct sets of typicality interpretations that are maximal w.r.t. $<_t$ (starting from a particular set of interpretations). Maximality is characterised straightforward as, $\mathcal{M}$ is maximal *iff* $\neg\exists\mathcal{M}'.\mathcal{M} <_t \mathcal{M}'$. Nested semantics are defined over the maximally preferred subsets of the set of all standard typicality models.

Definition 7.35 (Maximally Preferred Sets of Typicality Models)**.** For the set of all typicality models $\texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ of $\mathcal{K}$ over $\Delta^{\mathbb{C},\mathbb{O}}$ and $\leqslant_t^*$, the reflexive transitive closure of $<_t$, the set of *maximally preferred sets of typicality models* (Max-TMs) is

$$\texttt{TMax}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) = \{\mathcal{M} \subseteq \texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \mid \texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \leqslant_t^* \mathcal{M} \\ \wedge \neg\exists\mathcal{M}'.\mathcal{M} <_t \mathcal{M}'\}.$$

We consider two options to capture entailments using Max-TMs. The approach taken in [PT'17a; PT'17b; PT'18] investigates the sceptical[9] variant, where entailments are only inferred, if supported by all *maximal typicality models*. As we are dealing with sets of interpretations here (rather than single models as in [PT'17a; PT'17b; PT'18]), this corresponds to considering the union of all Max-TMs for $\texttt{TMod}_{\mathrm{nest}}()$. One alternative, often

9 Throughout [PT'17a; PT'17b; PT'18] this is described as the *conservative* approach.

considered in inconsistency tolerant repair semantics (e.g. [Bou'16]), is to *bravely* check whether the query is entailed in *any* of the most preferred sets of models. In Reiter's default logic, this corresponds to *credulous* reasoning [Rei'80]. While for instance, in systems over conflicting sensory data it is plausible to be interested in *any* scenario where a potentially critical entailment can be derived, it makes little sense for our purpose, because contradictory knowledge would inevitably be inferred (e.g. Fig. 7.6). However, the motivation of defeasible reasoning is to determine *coherent* extensions of the original knowledge.

A third approach—originally introduced here—is to characterise the coverage of the semantics by considering a single member of $\mathtt{TMax}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$. This idea is aligned with the parametrisation for semantic strength, when reasoning over individuals. A total order on all preference options, can be used to uniquely identify (or *select*) some $\mathcal{M} \in \mathtt{TMax}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$. Intuitively, the incremental preference over sets of models is determinised by specifying for each set of models a (unique) most preferred, admissible preference option.

Definition 7.36 (Preference Chain Compliance)**.** For a finite typicality domain $\Delta^{\mathbb{C},\mathbb{O}}$, let $<_{\mathsf{po}}$ be a total order over $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$. A preference chain $\mathcal{M}_1 <_t \mathcal{M}_2 <_t \ldots$ (with $\mathcal{M}_{i+1} = \mathcal{M}_i|_{p_i}$ for $i \geqslant 1$) is *compliant* with $<_{\mathsf{po}}$ *iff* p_i is the $<_{\mathsf{po}}$-minimal element out of all admissible preference options for $\mathcal{M}_i$.

For any set of typicality interpretations over a finite domain $\Delta^{\mathbb{C},\mathbb{O}}$, the set $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ is clearly finite when considering only the roles in a finite signature $sig(\mathcal{K})$. Consequently, the unique $<_{\mathsf{po}}$-minimal admissible preference option exists for every set of typicality models that has at least one admissible preference option. If there is no admissible preference option for a set of interpretations $\mathcal{M}$, then there does not exist a more preferred (w.r.t. $<_t$) set of models to begin with. Note the following additional observations about preference chains and total orders over preference options.

Remark 7.37.

1. For a total order $<_{\mathsf{po}}$ over $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ and a finite set of typicality interpretations $\mathcal{M}_1$ over a finite domain $\Delta^{\mathbb{C},\mathbb{O}}$, the maximal preference chain $\mathcal{M}_1 <_t \ldots <_t \mathcal{M}_n$, compliant with $<_{\mathsf{po}}$, is unique. Finiteness of $\mathcal{M}_1$ ensures the existence of maximal preference chains, while finiteness of $\Delta^{\mathbb{C},\mathbb{O}}$ ensures that at every $\mathcal{M}_i$, a unique, $<_{\mathsf{po}}$-minimal admissible preference option exists (except for the maximally preferred set $\mathcal{M}_n$).

2. For a finite typicality domain $\Delta^{\mathbb{C},\mathbb{O}}$, every full preference chain $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) <_t \mathcal{M}_1 <_t \cdots <_t \mathcal{M}_n$ is compliant with some total order $<_{\mathsf{po}}$ over $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$. From such a preference chain, the preference options $p_0, p_1, \cdots, p_{n-1}$ are naturally enumerated, s.t. $\mathcal{M}_1 = \mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})|_{p_0}$ and $\mathcal{M}_{i+1} = \mathcal{M}_i|_{p_i}$ $(1 \leqslant i < n)$. This chain of preference options provides a total order $<_{\mathsf{po}}$ over $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ that the original preference chain must be compliant with.

3. Example 7.34 already contains somewhat of an illustration for preference chain compliance. Recall $\mathcal{K}''$, p_1 and p_2. For the preference $p_1 <_{\text{po}} p_2$, (7.13) occurs on a compliant preference chain and then no set of models further down on the preference chain could be refined by p_2. Thus, from the Max-TM selected by $<_{\text{po}}$, it will be possible to derive $(\exists \mathit{friend}.\mathit{Happy})\{\mathit{molli}\}$, but not $(\exists \mathit{likes}.\mathit{Lazy})\{\mathit{molli}\}$. The opposite could be inferred from any Max-TM induced by $p_2 <_{\text{po}} p_1$.

Using preference chain compliance and the sceptical approach, we introduce two sets of models to be considered for different kinds of nested entailment.

Definition 7.38 (Nested Coverage for Typicality Models)**.** For a KB $\mathcal{K}$, a finite typicality domain $\Delta^{\mathbb{C},\mathbb{O}}$ and $<_{\text{po}}$, a total order on $\texttt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$, let

1. $\texttt{TMod}_{\texttt{nest}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) = \bigcup_{\mathcal{M} \in \texttt{TMax}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})} \mathcal{M}$, and
2. $\texttt{TMod}_{\texttt{nest}_{<_{\text{po}}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) = \mathcal{M} \in \texttt{TMax}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$, where

$$\texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})) <_{\text{t}} \ldots <_{\text{t}} \mathcal{M}$$

is the preference chain, compliant with $<_{\text{po}}$.

Note that Definition 7.38 provides the sought after refinement of the set of models considered to obtain semantics of propositional coverage. More specifically,

$$\texttt{TMod}_{\texttt{nest}_{<_{\text{po}}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \subseteq \texttt{TMod}_{\texttt{nest}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \subseteq \texttt{TMod}_{\texttt{prop}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$$

shows that entailments under propositional coverage are preserved in both kinds of nested coverage, and entailments supported by a single Max-TM clearly extend those supported by all Max-TMs. Definition 7.38 is used to fully instantiate the definition of defeasible entailment in typicality model semantics (Def. 7.10) for sceptical and selective nested semantics for rational and relevant strength. In the following, we will formally show superiority of both nested coverages over materialisation-based reasoning, for both of the discussed strengths.

7.3.3 *Nested Rational Reasoning*

Naturally, nested rational semantics rely on the rational typicality domain $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\text{rat}_\prec}$ (Def. 7.21). This domain and the characterisation of considered models for nested coverage (Def. 7.38), provide two instantiations of Definition 7.10 as follows.

Definition 7.39 (Nested Rational Semantics)**.** Let $\mathbb{C}, \mathbb{O}$ ba a finite, consistent context, containing the KB $\mathcal{K}$ and $\prec$ be a total order on $\mathsf{N_I}$. A defeasible inference problem $(\mathcal{K}, \alpha)$ is true under

- *Sceptical Nested Rational Semantics*
  iff $\forall \mathcal{I} \in \texttt{TMod}_{\texttt{nest}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\text{rat}_\prec}).\mathcal{I} \models \alpha$ (write $\mathcal{K} \models^{(\text{rat}_\prec,\text{nest})} \alpha$), and

- *Selective Nested Rational Semantics*
  iff $\forall \mathcal{I} \in \mathtt{TMod}_{\mathsf{nest}_{<_{\mathsf{po}}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_{\prec}}).\mathcal{I} \models \alpha$ (write $\mathcal{K} \models^{(\mathsf{rat}_{\prec},\mathsf{nest}_{<_{\mathsf{po}}})} \alpha$), for a given total order $<_{\mathsf{po}}$ on $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_{\prec}})$.

With this characterisation of nested semantics, we can formally show that both nested semantics are superior to propositional or materialisation-based semantics, without sacrificing the entailments produced by those foundations. The following result shows that the main goal for these novel semantics is successfully reached.

Theorem 7.40. *Sceptical and selective nested rational semantics allow for strictly more entailments than materialisation-based semantics, i.e.*

1. $\mathcal{K} \models^{(\mathsf{rat}_{\prec},\mathsf{mat})} \alpha \implies \mathcal{K} \models^{(\mathsf{rat}_{\prec},\mathsf{nest})} \alpha$,
2. $\mathcal{K} \models^{(\mathsf{rat}_{\prec},\mathsf{mat})} \alpha \implies \mathcal{K} \models^{(\mathsf{rat}_{\prec},\mathsf{nest}_{<_{\mathsf{po}}})} \alpha$,
3. $\mathcal{K} \models^{(\mathsf{rat}_{\prec},\mathsf{mat})} \alpha \not\Longleftarrow \mathcal{K} \models^{(\mathsf{rat}_{\prec},\mathsf{nest})} \alpha$, *and*
4. $\mathcal{K} \models^{(\mathsf{rat}_{\prec},\mathsf{mat})} \alpha \not\Longleftarrow \mathcal{K} \models^{(\mathsf{rat}_{\prec},\mathsf{nest}_{<_{\mathsf{po}}})} \alpha$,

for a KB $\mathcal{K}$ and a defeasible subsumption or instance check α.

Proof. 1 and 2 are not too involved due to the construction of Max-TMs. Recall the equivalence of materialisation-based reasoning and propositional semantics (Thm. 7.22). Using Remark 7.31 (Observation 2), it is clear that $\mathtt{TMod}_{\mathsf{prop}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_{\prec}}) \leqslant^*_{\mathsf{t}} \mathcal{M}$ implies $\mathcal{M} \subseteq \mathtt{TMod}_{\mathsf{prop}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_{\prec}})$ and thus $\mathtt{TMod}_{\mathsf{nest}_{<_{\mathsf{po}}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_{\prec}}) \subseteq \mathtt{TMod}_{\mathsf{nest}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_{\prec}}) \subseteq \mathtt{TMod}_{\mathsf{prop}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_{\prec}})$. Therefore, if α is entailed by $\mathtt{TMod}_{\mathsf{prop}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_{\prec}})$, then clearly, all interpretations in $\mathtt{TMod}_{\mathsf{nest}_{<_{\mathsf{po}}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_{\prec}})$ and $\mathtt{TMod}_{\mathsf{nest}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_{\prec}})$ also satisfy α (for any $<_{\mathsf{po}}$). This argument also formally shows implication from sceptical to selective nested semantics. Superiority of the latter is illustrated with 3 in Remark 7.37.

For 3 and 4 we already established the simple counterexample in Example 7.32, using the KB $\mathcal{K}$ from Example 7.9. It is not difficult to verify that the domain $\Delta^{\mathbb{C},\mathbb{O}}$ in those examples corresponds to $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_{\prec}}$ for any order $\prec$ (because there is only one individual). As there are no emptiness constraints in $\mathcal{K}$ (i.e. no GCIs of the form $\cdots \sqsubseteq \bot$), all admissible preference options for $\mathtt{TMod}_{\mathsf{prop}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_{\prec}})$ can be satisfied simultaneously. Consequently, there is exactly one set of models $\mathcal{M} \in \mathtt{TMax}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_{\prec}})$. Thus, for any order $<_{\mathsf{po}}$, $\mathcal{M} = \mathtt{TMod}_{\mathsf{nest}_{<_{\mathsf{po}}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_{\prec}}) = \mathtt{TMod}_{\mathsf{nest}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_{\prec}})$. From satisfaction of $\mathit{friend}((\mathit{molli}, \mathcal{D}), (\mathit{Dog}, \mathcal{D}))$ in all models in $\mathcal{M}$, it follows that $\mathcal{K} \models^{(\mathsf{rat}_{\prec},\mathsf{nest})} (\exists \mathit{friend}.\mathit{Happy})\{\mathit{molli}\}$ and $\mathcal{K} \models^{(\mathsf{rat}_{\prec},\mathsf{nest}_{<_{\mathsf{po}}})} (\exists \mathit{friend}.\mathit{Happy})\{\mathit{molli}\}$, something that could not be derived from the materialisation framework. □

Unfortunately, the rational domain induced by $\mathcal{K}$ in Example 7.9 only exhibits two levels of typicality. For simplicity and brevity, this is enough to show the desired result and we shall not afford the space to make a more involved example. One could easily imagine a more complex DBox and some TBox constraints that would result in $|\mathtt{chain}(\mathcal{K})| > 2$. However, as only two concept representatives per element in $\mathbb{C}$ exist, role successors will be either fully typical or not typical at all. This is not the case for the relevant domain and allows for much more intricate examples.

7.3.4 *Nested Relevant Reasoning*

As for propositional semantics in Section 7.2.2 and 7.2.3, the investigation of nested relevant semantics is analogous to the preceding section.

Nested relevant semantics are naturally based on the relevant typicality domain $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec}$ (Def. 7.24). This domain, together with the definition of considered models for nested coverage (Def. 7.38), provides explicit nested semantics in term of Definition 7.10 as follows.

Definition 7.41 (Nested Relevant Semantics). A defeasible inference query $(\mathcal{K}, \alpha)$ is true under

- *Sceptical Nested Relevant Semantics*
  iff $\forall \mathcal{I} \in \mathtt{TMod}_{\mathtt{nest}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec}).\mathcal{I} \models \alpha$ (write $\mathcal{K} \models^{(\mathtt{rel}_\prec,\mathtt{nest})} \alpha$), and
- *Selective Nested Relevant Semantics*
  iff $\forall \mathcal{I} \in \mathtt{TMod}_{\mathtt{nest}_{<_{\mathtt{po}}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec}).\mathcal{I} \models \alpha$ (write $\mathcal{K} \models^{(\mathtt{rel}_\prec,\mathtt{nest}_{<_{\mathtt{po}}})} \alpha$), for a given total order $<_{\mathtt{po}}$ on $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec})$.

Before showing the analogous result to Theorem 7.40 for the rational case, we illustrate the power of the relevant domain, both in terms of nested coverage and a solution to inheritance blocking.

Example 7.42. For recognisability, we introduce a variant of the KB from Example 7.9, that exhibits more intricate features, in particular the ability to showcase inheritance blocking at role successors. Let $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, with

$$\begin{aligned}
\mathcal{A} = \{&(\mathit{Cat} \sqcap \exists \mathit{friend}.\mathit{Dog})(\mathit{molli})\} \\
\mathcal{T} = \{&\mathit{Cat} \sqsubseteq \exists \mathit{likes}.\mathit{Cat}, \\
&\exists \mathit{friend}.\mathit{Dog} \sqcap \mathit{Smart} \sqsubseteq \bot, && \text{(T1)} \\
&\exists \mathit{friend}.(\mathit{Dog} \sqcap \exists \mathit{hates}.\mathit{Cat}) \sqsubseteq \bot\} && \text{(T2)} \\
\mathcal{D} = \{&\mathit{Dog} \mathrel{\tilde{\sqsubseteq}} \mathit{Happy}, \mathit{Dog} \mathrel{\tilde{\sqsubseteq}} \exists \mathit{hates}.\mathit{Cat}, \\
&\mathit{Cat} \mathrel{\tilde{\sqsubseteq}} \mathit{Lazy}, \mathit{Cat} \mathrel{\tilde{\sqsubseteq}} \mathit{Smart}\}
\end{aligned}$$

First of all, let us analyse this KB. From the new emptiness-constraints in the TBox, no top-level (materialisation-based) conflict for the concept names Cat and Dog can be derived. That is, $\mathcal{K} \not\models \overline{\mathcal{D}} \sqcap \mathit{Cat} \sqsubseteq \bot$ (likewise for Dog), and thus $\mathtt{chain}(\mathcal{K}) = \langle \mathcal{D}, \emptyset \rangle$. However, we can imagine some domain elements that might be affected by those TBox constraints. For one, the individual representative for molli is in conflict with (T1) (in standard typicality models of $\mathcal{K}$). Also, any element with a friend successor in Dog, satisfying the entire DBox conflicts with (T2). These two types of conflicts

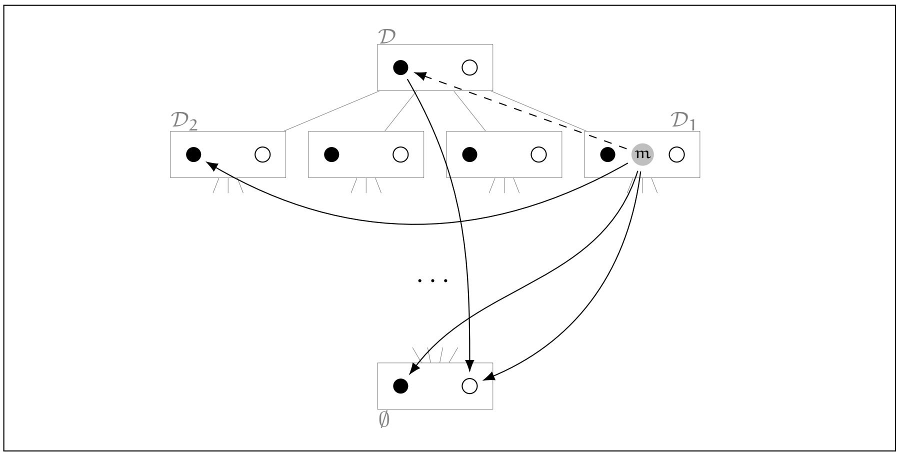

Figure 7.7: Labelled graph visualisation of a typicality interpretation over (an excerpt of) $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec}$.

are intuitively resolved, when considering domain elements (for the latter: the respective successor) associated to the DBox subsets

$$\mathcal{D}_1 = \{\mathrm{Dog} \sqsubsetsim \mathrm{Happy}, \mathrm{Dog} \sqsubsetsim \exists \mathit{hates}.\mathrm{Cat}, \mathrm{Cat} \sqsubsetsim \mathrm{Lazy}\} \text{ and}$$

$$\mathcal{D}_2 = \{\mathrm{Dog} \sqsubsetsim \mathrm{Happy}, \mathrm{Cat} \sqsubsetsim \mathrm{Lazy}, \mathrm{Cat} \sqsubsetsim \mathrm{Smart}\}, \text{ respectively.}$$

Figure 7.7 contains (part of) an interpretation over an excerpt of the relevant typicality domain in a somewhat abstracted form. To keep the visual representation in check, we use different symbols for the first component of representative elements. A full circle describes an element representing Dog, empty circles are Cat representatives and the grey circle with label $\mathfrak{m}$ represents molli. The next abstraction is, that all domain elements are associated with the DBox subset, labelling the surrounding rectangle, as previously described for Figure 7.4. For instance, the molli representative is associated with $\mathcal{D}_1$. To present a concise visualisation, we only care about elements associated to the sets in $\mathtt{chain}(\mathcal{K})$ and $\mathcal{D}_1$, $\mathcal{D}_2$, hence, excluding most of the subset lattice between $\mathcal{D}_1$, $\mathcal{D}_2$ and $\emptyset$. Also, we omit the labels in Fig. 7.7, because the respective entailments are very similar to previous examples. Note that edges pointing to Dog elements are always of type friend. molli's Cat successor is labelled with likes, and the most typical Dog hates Cats (i.e. $\mathit{hates}((\mathrm{Dog}, \mathcal{D}), (\mathrm{Cat}, \emptyset))$). Other Dog elements associated with DBox subsets containing $\mathrm{Dog} \sqsubsetsim \exists \mathit{hates}.\mathrm{Cat}$ should have such an edge too, but these and other irrelevant edges are also omitted in Fig. 7.7.

To see that no preference options are ever in conflict with one another, note that the TBox can only be violated for elements with a Dog friend. From $\mathcal{K}$, this relation will only ever be entailed for molli, never for any anonymous concept representative. As a matter of fact, the assertion in $\mathcal{A}$ ensures that molli is exceptional for the full DBox $\mathcal{D}$. As a result, $\mathtt{ext}(\mathcal{K}_{\mathtt{rel}_\prec}, \mathrm{molli}) = \mathcal{D}_1$ (for any $\prec$), whereas rational strength would have had to associate molli with the next biggest element in $\mathtt{chain}(\mathcal{K})$, i.e. $\emptyset$.

This allows to draw the conclusion $\mathsf{Lazy}\{\mathsf{molli}\}$, but this was derivable from materialisation-based relevant reasoning as well.

From the lack of conflicts between preference options, we know that there is only one unique Max-TM $\mathcal{M} \in \mathtt{TMax}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec})$ and all $\mathcal{I} \in \mathcal{M}$ contain at least the solid edges displayed in Fig. 7.7. The dashed edge from molli to the most typical Dog representative cannot be derived, as it would violate (T2). For rational strength, there are only two options for molli's friend successor, either typical (here: $(\mathsf{Dog}, \mathcal{D})$) or non-typical (here: $(\mathsf{Dog}, \emptyset)$). In $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec}$, there are all the representatives in between, in particular $(\mathsf{Dog}, \mathcal{D}_2)$. The edge $\mathsf{friend}((\mathsf{molli}, \mathcal{D}_1), (\mathsf{Dog}, \mathcal{D}_2))$ persists in all $\mathcal{I} \in \mathcal{M}$ as well, allowing to draw the conclusion $(\exists \mathsf{friend}.\mathsf{Happy})\{\mathsf{molli}\}$.

Example 7.42 illustrates that the defeasible information inferred for quantified concepts under nested relevant semantics, does not suffer from inheritance blocking. This successfully shows that two of the main drawbacks of materialisation-based Rational Closure (Sec. 5.1, 5.2) can be resolved using typicality models.

As for the rational case, we show that both nested semantics are stronger than propositional or materialisation-based semantics, in the amount of entailed inferences.

Theorem 7.43. *Sceptical and selective nested relevant semantics allow for strictly more entailments than materialisation-based semantics, i.e.*

1. $\mathcal{K} \models^{(\mathtt{rel}_\prec,\mathtt{mat})} \alpha \implies \mathcal{K} \models^{(\mathtt{rel}_\prec,\mathtt{nest})} \alpha$,
2. $\mathcal{K} \models^{(\mathtt{rel}_\prec,\mathtt{mat})} \alpha \implies \mathcal{K} \models^{(\mathtt{rel}_\prec,\mathtt{nest}_{<\mathtt{po}})} \alpha$,
3. $\mathcal{K} \models^{(\mathtt{rel}_\prec,\mathtt{mat})} \alpha \not\Longleftarrow \mathcal{K} \models^{(\mathtt{rel}_\prec,\mathtt{nest})} \alpha$, *and*
4. $\mathcal{K} \models^{(\mathtt{rel}_\prec,\mathtt{mat})} \alpha \not\Longleftarrow \mathcal{K} \models^{(\mathtt{rel}_\prec,\mathtt{nest}_{<\mathtt{po}})} \alpha$,

for a KB $\mathcal{K}$ and a defeasible subsumption or instance check α.

Proof. From Observation 2 in Remark 7.31 it is clear that $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec}) \leqslant^*_t \mathcal{M}$ implies $\mathcal{M} \subseteq \mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec})$ and thus $\mathtt{TMod}_{\mathtt{nest}_{<\mathtt{po}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec}) \subseteq \mathtt{TMod}_{\mathtt{nest}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec}) \subseteq \mathtt{TMod}_{\mathtt{prop}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec})$. Therefore, all entailments supported by $\mathtt{TMod}_{\mathtt{prop}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec})$ are also supported by $\mathtt{TMod}_{\mathtt{nest}_{<\mathtt{po}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec})$ and $\mathtt{TMod}_{\mathtt{nest}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec})$. The equivalence of propositional relevant reasoning and materialisation-based relevant reasoning (Theorem 7.25) therefore proves 1 and 2.

For 3 and 4, a counterexample is given in Example 7.42. □

For a full account of nested semantics and their relations towards each other, up to and including Theorem 7.43, we refer to the concluding remarks of this work in Chapter 9. We continue with the extensive investigation of the computational complexity for the newly developed typicality-model semantics.

8 COMPUTATIONAL COMPLEXITY

From a computational point of view, materialisation-based approaches to determine the defeasible consequences of a knowledge base, hold practical value. They all rely on a reduction of a defeasible inference problem to its classical counterpart. When considering $\mathcal{ALC}$, the complexity of classical reasoning (ExpTime) will be maintained for polynomial such reductions (e.g. Rational Closure), as well as exponential ones (e.g. Relevant Closure). If done appropriately, also for sub-boolean DLs such as $\mathcal{EL}_\bot$ (cf. Sec. 4.4), the complexity of defeasible query entailment under rational strength can be shown to remain polynomial altogether. Thus, materialisation-based reasoning allows for simple practical implementations ([CMMN'14; CMM+'15; CMMV'13; GGPR'17]), by employing highly efficient DL-reasoners [GHM+'14; KKS'14]. From Theorem 7.22 and 7.25, and especially Corollary 7.23 it is clear that determining entailments under propositional typicality model semantics is, roughly speaking, a reformulation of such a reduction algorithm.

It is much less intuitive, how high the computational effort is to determine sets of typicality models that are maximal under typicality preference. The best upper bound for the number of Max-TMs that we have been able to prove so far, is exponential in the size of the considered typicality domain. In addition to the different size of domain (rational vs. relevant), the complexity also varies between the different types of nested semantics:

SCEPTICAL: *All* of the Max-TMs must be investigated.

SELECTIVE: One specific Max-TM must be determined.

In this chapter we extensively study complexity upper bounds of all nested semantics for rational and relevant strength, and prove a complexity lower bound for sceptical nested rational entailment. We split the former into two parts, separating the presentation of algorithms (Sec. 8.1) from the analysis on their computational complexity (Sec. 8.2). The reduction from a known SAT problem, to obtain a tight lower bound for sceptical nested rational entailment (Sec. 8.3), is unfortunately not transferable to relevant strength. Lower bounds for selective nested semantics, on the other hand, are relatively easy to show for both rational and relevant strength.

8.1 ALGORITHMS FOR TYPICALITY MODELS

As in the preceding chapters, a large number of results in the algorithmic characterisation of typicality maximisation are general in terms of the underlying typicality domain. The following two such general results are used as the main motivation for the algorithms that we present in this section. First and foremost, we formalise the initial statement on the upper bound for the number of Max-TMs.

Proposition 8.1. *For a KB $\mathcal{K}$, a quantification closed context $\mathbb{C}, \mathbb{O}$ containing $\mathcal{K}$ and a typicality domain $\Delta^{\mathbb{C},\mathbb{O}}$, the set* $\mathtt{TMax}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ *contains at most exponentially many Max-TMs in the size of $\Delta^{\mathbb{C},\mathbb{O}}$.*

Proof. For $n = |\Delta^{\mathbb{C},\mathbb{O}}|$, there are at most $n^2 * |sig_R(\mathcal{K})|$ preference options (admissibility aside) in $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$. A simple upper bound for Max-TMs is clearly the number of distinct *sets* of preference options, i.e. exponentially many in $|\Delta^{\mathbb{C},\mathbb{O}}|$. □

Remark 8.2. Of course, if one of these sets of preference options (in the preceding proof), say $P \subseteq \mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$, being satisfied provides a Max-TM, then none of the (strict) subsets of P being satisfied provides a set of models that is maximal w.r.t. $<_t$. This practically reduces the upper bound to the number of subsets of $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ with $\frac{|\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})|}{2}$ elements, which remains exponential in $|\Delta^{\mathbb{C},\mathbb{O}}|$.

The main idea for an algorithmic characterisation rests in the intuition that a set $\mathcal{M} \in \mathtt{TMax}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ can be identified through a *preference chain* $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) <_t \mathcal{M}_1 <_t \ldots <_t \mathcal{M}$. However, Proposition 8.1 shows that it is not feasible to enumerate all Max-TMs when determining entailments under sceptical semantics. An alternate approach is to identify a single preference chain, leading to a Max-TM that does not satisfy the query, hence solving the complement of the decision problem at hand. For selective semantics, we have shown that the additional input identifies a unique preference chain, leading to a unique Max-TM. Thus, we present a maximisation algorithm, that essentially traverses a single preference chain (either specified by input or guessed by the algorithm), in Section 8.1.2.

To achieve this traversal, we need to provide the means to determine admissible preference options for sets of typicality models, as well as feasible means to handle the sets along a preference chain. To this end, the following consequence of the typicality preference relation is the second motivational result for this algorithmic characterisation.

Proposition 8.3. *If $\mathcal{M}_1 \subseteq \mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ is closed under intersection and $\mathcal{M}_1 <_t \mathcal{M}_2$, then $\mathcal{M}_2$ is closed under intersection.*

Proof. Let $r(d, e)$ be the admissible preference option of $\mathcal{M}_1$ s.t. $\mathcal{M}_2 = \mathcal{M}_1|_{r(d,e)}$. Clearly, $\mathcal{M}_2 \subseteq \mathcal{M}_1$, thus for $\mathcal{J}_1, \mathcal{J}_2 \in \mathcal{M}_2$, it follows that $\mathcal{J}_1 \cap \mathcal{J}_2 \in \mathcal{M}_1$, because $\mathcal{M}_1$ is closed under intersection. Both $\mathcal{J}_1$ and $\mathcal{J}_2$ satisfy $r(d, e)$ which implies that $\mathcal{J}_1 \cap \mathcal{J}_2$ satisfies $r(d, e)$ Consequently, $\mathcal{J}_1 \cap \mathcal{J}_2 \in \mathcal{M}_1|_{r(d,e)} = \mathcal{M}_2$. □

From Proposition 8.3 and the fact that $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ is closed under intersection (for a finite $\Delta^{\mathbb{C},\mathbb{O}}$, cf. Prop. 7.12), we know that every $\mathcal{M}_i$ along a preference chain $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) <_t \mathcal{M}_1 <_t \mathcal{M}_2 \ldots$, including Max-TMs are closed under intersection and trivially contain a $\subseteq$-smallest typicality model. This smallest element $\mathcal{J}_i$ of $\mathcal{M}_i$ is canonical for its respective $\mathcal{M}_i$, in the sense that $\mathcal{J}_i \subseteq \mathcal{I}$ for all $\mathcal{I} \in \mathcal{M}_i$. Hence, these canonical typicality models (starting from the minimal typicality model $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ for $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$) are sufficient to determine admissible preference options. Intuitively, for a preference $\mathcal{M}_1 <_t \mathcal{M}_2$, we algorithmically construct the canonical model for $\mathcal{M}_2$ from the canonical model for $\mathcal{M}_1$, in terms of a minimal model completion (Section 8.1.1). This construction has been one part of the original typicality model upgrade procedure in [PT'17a; PT'17b].

8.1.1 *Model Completion*

In Example 7.9, 7.32 and 7.33 we have intuitively already applied model completion. After expecting a preference option to be satisfied, it was clear that more information (e.g. role edges) must be contained in any $\subseteq$-extension of that interpretation, to satisfy the KB (Exm. 7.33). Thus, the idea of model completion is technically not tied to the preference relation $<_t$, but any typicality interpretation can be (attempted to be) completed into a model of the given KB. We introduce this completion in Algorithm 8.1 and continue to show its correctness, while its computational complexity is determined in Section 8.2.

The intuition behind this model completion algorithm is relatively simple. It iteratively determines if the current interpretation is a standard typicality model of the KB. If not, there must be some counterexample given by some domain element. The algorithm attempts to mend this counterexample, by extending the interpretation appropriately. For instance, some domain element d might belong to the left-hand side of some GCI (or associated DCI), but not to the right-hand side in the current interpretation. If this interpretation satisfies 1 of Definition 7.6 ($\forall C \in \mathbb{C}.(C, \mathcal{U}) \in C^{\mathcal{I}}$), then it is easy to extend it in such a way that this particular domain element and GCI (or DCI) do not present a counterexample anymore. The element d will be added to the appropriate concept name extensions and appropriate role edges are introduced, so that afterwards, d belongs to the right-hand side of the violated GCI (or DCI). Once fixed, this specific counterexample can never be a counterexample again, if the interpretation is only ever extended. The counterexamples with GCIs and DCIs are treated in the exact same way. Formally, we introduce the symbol $\bowtie \in \{\sqsubseteq, \sqsubseteq\!\!\!\sim\}$ to cover both cases by using e.g. $E \bowtie F$.

The modification described above is formalised as the *promotion* of a domain element within an interpretation, by a given concept. With this alteration, we need to be careful that (new) implicitly satisfied quantifications (e.g. $d \in (\exists r.E)^{\mathcal{J}}$) follow the principle of standard interpretations. In particular, they should satisfy Property 3 of Def. 7.6 in case $E \in \mathbb{C}$. This is achieved with a second kind of modification on typicality interpretations, called *standardisation*.

Definition 8.4 (Promotion and Standardisation). For a typicality interpretation $\mathcal{I} = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\mathcal{I}})$, an element $d \in \Delta^{\mathbb{C},\mathbb{O}}$ and an $\mathcal{EL}$ concept $E = \bigsqcap_{1 \leqslant i \leqslant n} A_i \sqcap \bigsqcap_{1 \leqslant j \leqslant m} \exists r_j.F_j$ with $\mathrm{Qc}(E) \subseteq \mathbb{C}$, the *promotion* of $\mathcal{I}$ w.r.t. d and E is $\mathcal{I}(d, E) = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\mathcal{I}(d,E)})$ s.t.

$$A^{\mathcal{I}(d,E)} = \begin{cases} A^{\mathcal{I}} \cup \{d\} & \text{, if } A \in \{A_1, \dots, A_n\} \\ A^{\mathcal{I}} & \text{, otherwise} \end{cases}$$

$$r^{\mathcal{I}(d,E)} = r^{\mathcal{I}} \cup \{(d, (F_j, \emptyset)) \mid 1 \leqslant j \leqslant m \wedge r = r_j\}$$

$$a^{\mathcal{I}(d,E)} = a^{\mathcal{I}}$$

for all $A \in N_C$, $r \in N_R$, $a \in N_I$.

The *standardisation* of $\mathcal{I}$ is $\mathsf{S}(\mathcal{I}) = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\mathsf{S}(\mathcal{I})})$ s.t.

$$A^{\mathsf{S}(\mathcal{I})} = A^{\mathcal{I}}$$

$$r^{\mathsf{S}(\mathcal{I})} = r^{\mathcal{I}} \cup \{(d, (F, \emptyset)) \mid F \in \mathbb{C} \wedge d \in (\exists r.F)^{\mathcal{I}}\}$$

$$a^{\mathsf{S}(\mathcal{I})} = a^{\mathcal{I}}$$

for all $A \in N_C$, $r \in N_R$, $a \in N_I$.

Combining both extensions, is called the *standard promotion* of $\mathcal{I}$ w.r.t. d, E, formally, $\mathsf{S}(\mathcal{I}(d, E))$. The term standardisation is for lack of a more specific term, because Def. 8.4 clearly only addresses the third property of the definition for standard typicality interpretations. Property 1 and 2 of Definition 7.6 will be assumed to be satisfied for the input interpretation $\mathcal{I}$ in the following. While it would be possible to define a true standardisation, explicitly taking care of the assumed properties, it is never required for the way standardisation is utilised here. The overall starting point for the typicality maximisation (Sec. 8.1.2) is the minimal typicality model $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$, which already satisfies 1 and 2 of Def. 7.6 and will only ever be extended throughout the algorithm.

Lemma 8.5. *Let $\mathcal{I} = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\mathcal{I}})$ be a typicality interpretation, $d \in \Delta^{\mathbb{C},\mathbb{O}}$ and $E = \bigsqcap_{1 \leqslant i \leqslant n} A_i \sqcap \bigsqcap_{1 \leqslant j \leqslant m} \exists r_j.F_j$ with $\mathrm{Qc}(E) \subseteq \mathbb{C}$. If $\mathcal{I}$ satisfies Property 1 and 2 of Definition 7.6, then the following hold*

1. $d \in E^{\mathcal{I}(d,E)}$
2. $\mathsf{S}(\mathcal{I})$ *is a standard typicality interpretation*
3. $\mathsf{S}(\mathcal{I}(d, E))$ *is a standard typicality interpretation with* $d \in E^{\mathsf{S}(\mathcal{I}(d,E))}$

Proof.

CLAIM 1 It follows immediately from Definition 8.4 that $d \in (\bigsqcap_{1 \leqslant i \leqslant n} A_i)^{\mathcal{I}(d,E)}$. By assumption, we know that $(F_j, \emptyset) \in F_j^{\mathcal{I}}$ (1 of Def. 7.6). Clearly $\mathcal{I} \subseteq \mathcal{I}(d, E)$, hence Lemma 2.12 implies $(F_j, \emptyset) \in F_j^{\mathcal{I}(d,E)}$ (for $1 \leqslant j \leqslant m$). Again, by Def. 8.4 $(d, (F_j, \emptyset)) \in r_j^{\mathcal{I}(d,E)}$, which, together with the preceding observation, implies $d \in (\bigsqcap_{1 \leqslant j \leqslant m} \exists r_j.F_j)^{\mathcal{I}(d,E)}$.

CLAIM 2 This is rather trivial, as the construction of $\mathsf{S}(\mathcal{I})$ forces for every $d \in (\exists r.F)^{\mathsf{S}(\mathcal{I})}$ ($r \in \mathsf{N_R}, F \in \mathbb{C}$) that $(d, (F, \emptyset)) \in r^{\mathsf{S}(\mathcal{I})}$, i.e. Property 3 of Def. 7.6. The other properties of Def. 7.6 are satisfied for $\mathcal{I}$ by assumption, and thus they are satisfied for all extensions of $\mathcal{I}$, in particular $\mathsf{S}(\mathcal{I})$, by Lemma 2.12.

CLAIM 3 Again, by Lemma 2.12, also $\mathcal{I}(d, E)$, as an extension of $\mathcal{I}$, satisfies the properties assumed for $\mathcal{I}$. Hence, Claim 2 applies also to $\mathcal{I}(d, E)$ and allows the combination of Claims 1 and 2. □

We rely on standard promotions in Algorithm 8.1 and on Lemma 8.5 for showing that this construction achieves to resolve counterexamples to the model property. However, not all such counterexamples can be resolved this way. If a domain element d satisfies the left-hand side of a GCI $E \sqsubseteq \bot$, then neither the current interpretation, nor any extension of it belong to $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$. These cases need to be treated specifically by Alg. 8.1. Intuitively, because all counterexamples are resolved by minimally extending the interpretation, violation of $E \sqsubseteq \bot$ implies that the original input interpretation of Alg. 8.1 cannot be completed into a model (proven in Prop. 8.8). Therefore, in this case, Algorithm 8.1 specifically returns `failure`. An example of this situation is already provided in Figure 7.6 (Page 107).

Algorithm 8.1: Minimal Model Completion

Input: KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, typicality interpretation $\mathcal{I} = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\mathcal{I}})$
Output: typicality interpretation $\mathcal{I}_n$ or `failure`

1 $\mathcal{I}_0 := \mathsf{S}(\mathcal{I})$;
2 $i := 0$;
3 **while** $\mathcal{I}_i \notin \mathtt{TMod}((\mathcal{T}, \mathcal{D}), \Delta^{\mathbb{C},\mathbb{O}})$ **do**
4   Let $d = (\chi, \mathcal{U}) \in \Delta^{\mathbb{C},\mathbb{O}}$ s.t. $E \bowtie F \in \mathcal{T} \cup \mathcal{U}$ and $d \in E^{\mathcal{I}_i} \setminus F^{\mathcal{I}_i}$;
5   **if** $F = \bot$ **then**
6   **return** `failure`;
7   **end**
8     $\mathcal{I}_{i+1} := \mathsf{S}(\mathcal{I}_i(d, F))$;
9     $i := i + 1$;
10 **end**
11 **return** $\mathcal{I}_i$;

Note that from Line 3 of Alg. 8.1 it becomes apparent that we are only interested in satisfying the TBox and the sets of DCIs that are associated with every domain element. We justify this in a similar way to the limited construction of standardisations. Clearly, it would be possible to extend an interpretation to "make it satisfy" assertions of an ABox. However, for our purposes, Algorithm 8.1 will only be used with input typicality interpretations extending the minimal typicality model. From Lemma 2.12 it follows that, if a typicality interpretation satisfies an ($\mathcal{EL}$) ABox (such as $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$), then every extension of that typicality interpretation satisfies this ABox as well.

By slight abuse of notation and in favour of simplicity in the following, we capture the set of all typicality models extending a given interpretation $\mathcal{I} = (\Delta^{C,O}, \cdot^{\mathcal{I}})$ as

$$\mathtt{TMod}(\mathcal{K},\mathcal{I}) = \{\mathcal{J} \in \mathtt{TMod}(\mathcal{K},\Delta^{C,O}) \mid \mathcal{I} \subseteq \mathcal{J}\}. \tag{8.1}$$

Remark 8.6. Note the following observations.

1. $\mathtt{TMod}(\mathcal{K},\mathcal{I})$ is not necessarily non-empty. In particular, if for some $\mathcal{I}_i$ in Alg. 8.1 (initialised with $\mathcal{I}$) a counterexample for $F \sqsubseteq \bot$ is encountered, then $\mathtt{TMod}(\mathcal{K},\mathcal{I}_i) = \emptyset$. Informally, this is an indication to $\mathcal{I}$ containing contradictory information to the KB.
2. If $\mathcal{I} \in \mathtt{TMod}(\mathcal{K},\Delta^{C,O})$, then not only $\mathcal{I} \in \mathtt{TMod}(\mathcal{K},\mathcal{I})$ holds, but $\mathcal{I}$ is clearly the $\subseteq$-smallest model in $\mathtt{TMod}(\mathcal{K},\mathcal{I})$.

Before proving correctness of Algorithm 8.1, we introduce a very powerful lemma, connecting the different interpretations $\mathcal{I}_i$ throughout the iterations of this algorithm.

Lemma 8.7. *Suppose Algorithm 8.1 terminates on the input $\mathcal{K} = (\mathcal{T},\mathcal{D})$ and $\mathcal{I}$. For any two typicality interpretations $\mathcal{I}_k$, $\mathcal{I}_j$ ($0 \leqslant k, j$) within the run of Alg. 8.1 on $\mathcal{K}$ and $\mathcal{I}$, it holds that*

$$\mathtt{TMod}(\mathcal{K},\mathcal{I}_k) = \mathtt{TMod}(\mathcal{K},\mathcal{I}_j) \tag{8.2}$$

Proof. To avoid confusion about numbers of iterations and indices of interpretations , note that we use the value that i (iteration counter in Alg. 8.1) would have if the algorithm reached Line 9 in the current iteration, as the iteration number. For example, in iteration 1, Line 3 checks $\mathcal{I}_0$ and for terminating in Line 6 at iteration $n+1$, the highest index on the constructed interpretations $\mathcal{I}_i$ is $i = n$.

For the special case that the input interpretation $\mathcal{I}$ (or $\mathcal{I}_0$) belongs to $\mathtt{TMod}(\mathcal{K},\Delta^{C,O})$ (i.e. the *while*-loop is never entered), the claim trivially holds for $k = j = 0$.

The second special case is, that in $\mathcal{I}_0$ there exists an element $(\chi, \mathcal{U}) \in E^{\mathcal{I}_0}$ and some $E \sqsubseteq \bot \in \mathcal{T}$ in the KB. In this case, Line 6 is reached in the first iteration of the *while*-loop and $\mathtt{TMod}(\mathcal{K},\mathcal{I}_0) = \emptyset$, because all typicality interpretations $\mathcal{J}$ with $\mathcal{I}_0 \subseteq \mathcal{J}$ will satisfy $(\chi, \mathcal{U}) \in E^{\mathcal{I}_0}$ (by Lemma 2.12), and therefore, not satisfy $E \sqsubseteq \bot$. Again, for $\mathcal{I}_0$ being the only interpretation to consider, the claim trivially holds for $k = j = 0$.

Assume the *while*-loop is entered and Line 6 is not reached in the first iteration. We show that

$$\mathtt{TMod}(\mathcal{K},\mathcal{I}_i) = \mathtt{TMod}(\mathcal{K},\mathcal{I}_{i+1}) \tag{8.3}$$

holds for $0 \leqslant i < n$, for the algorithm terminating either in Line 11 after n iterations, or in Line 6 during iteration $n+1$.[1] In both cases, the last

1 The case that Alg. 8.1 terminates in Line 6 in the first iteration is covered above, therefore, considering it to terminate not sooner than in iteration $n+1$ in Line 6 does not exclude any cases.

constructed interpretation is $\mathcal{I}_n$, regardless of whether $\mathcal{I}_n$, or `failure` will be returned.

The inclusion $\texttt{TMod}(\mathcal{K}, \mathcal{I}_i) \supseteq \texttt{TMod}(\mathcal{K}, \mathcal{I}_{i+1})$ is trivial, because by construction $\mathcal{I}_i \subseteq \mathcal{I}_{i+1}$ (see (8.1)).

For $\subseteq$, we show that $\mathcal{J} \in \texttt{TMod}(\mathcal{K}, \mathcal{I}_i)$ implies $\mathcal{I}_{i+1} \subseteq \mathcal{J}$. Because we know $\mathcal{I}_i \subseteq \mathcal{I}_{i+1}$ and $\mathcal{I}_i \subseteq \mathcal{J}$, we only need to show that the information not shared between $\mathcal{I}_i$ and $\mathcal{I}_{i+1}$ is contained in $\mathcal{J}$. By the assumption about termination, we know, for $E \bowtie F$ chosen in Line 4 during any iteration from 1 to n, that $F \neq \bot$. Hence, w.l.o.g.

$$F = \bigsqcap_{1 \leqslant h \leqslant m} A_h \sqcap \bigsqcap_{1 \leqslant l \leqslant o} \exists r_l.G_l$$

at every iteration before $n+1$. Then, by Definition 8.4, it holds for the promotion $\mathcal{I}_i(d, F)$ $(0 \leqslant i < n)$, that $d \in A_h^{\mathcal{I}_i(d,F)}$ for all $1 \leqslant h \leqslant m$ and $(d, (G_l, \emptyset)) \in r_l^{\mathcal{I}_i(d,F)}$ for all $1 \leqslant l \leqslant o$. These are the only containments that are (potentially) not shared between $\mathcal{I}_i$ and $\mathcal{I}_i(d, F)$ by construction of $\mathcal{I}_i(d, F)$. From

- $d \in E^{\mathcal{I}_i}$,
- $E^{\mathcal{I}_i} \subseteq E^{\mathcal{J}}$ (Lem. 2.12), and
- $\mathcal{J} \in \texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$,

it follows that $d \in F^{\mathcal{J}}$. This implies that $d \in A_h^{\mathcal{J}}$ $(1 \leqslant h \leqslant m)$ and because $\mathcal{J}$ is standard, $(d, (G_l, \emptyset)) \in r_l^{\mathcal{J}}$. Therefore, $\mathcal{I}_i(d, F) \subseteq \mathcal{J}$. The promotion $\mathcal{I}_i(d, F)$ being contained in the standard model $\mathcal{J}$, implies $(\exists r.H)^{\mathcal{I}_i(d,F)} \subseteq (\exists r.H)^{\mathcal{J}}$ for all $H \in \mathbb{C}$, $r \in N_R$ by Lemma 2.12. Thus, all role edges added to obtain $S(\mathcal{I}_i(d, F))$ must also be contained $\mathcal{J}$, showing that $\mathcal{I}_{i+1} = S(\mathcal{I}_i(d, F)) \subseteq \mathcal{J}$.

This proves (8.3) for iterations $0 \leqslant i < n$ (under our termination assumptions). If the assumption $\mathcal{J} \in \texttt{TMod}(\mathcal{K}, \mathcal{I}_i)$ is invalid, then $\texttt{TMod}(\mathcal{K}, \mathcal{I}_i) = \emptyset$ and $\texttt{TMod}(\mathcal{K}, \mathcal{I}_i) \subseteq \texttt{TMod}(\mathcal{K}, \mathcal{I}_{i+1})$ holds trivially.

For the original claim, assume w.l.o.g. $k < j$ ($k = j = 0$ is covered by the border cases). Transitive applications of (8.3) show $\texttt{TMod}(\mathcal{K}, \mathcal{I}_k) = \texttt{TMod}(\mathcal{K}, \mathcal{I}_j)$. □

Interestingly enough, Lemma 8.7 does not rely on the outcome of Algorithm 8.1, because (8.3) is shown for all iterations *before* termination. It is a simple consequence, that if the algorithm terminates in Line 6 at iteration $n+1$, with $\mathcal{I}_n$ being the last created interpretation, that $\texttt{TMod}(\mathcal{K}, \mathcal{I}_n) = \emptyset$ and thus $\texttt{TMod}(\mathcal{K}, \mathcal{I}_0) = \emptyset$. For the time being, we include termination of Alg. 8.1 as an assumption in Lemma 8.7, and refer the investigation for termination and complexity to Section 8.2. Incidentally, we will show that this model completion terminates on all finite inputs (Prop. 8.16), making this assumption in Lem. 8.7 moot.

We use Lemma 8.7 to show that the outcome of Algorithm 8.1 is always correct. That is, if it returns an interpretation, then this is the smallest model of the KB extending the input interpretation, and if it returns `failure`, then no extension of the input interpretation satisfies the KB.

Proposition 8.8. *Let* $\mathcal{K} = (\mathcal{A},\mathcal{T},\mathcal{D})$ *be a KB and* $\mathcal{I} = (\Delta^{C,O}, \cdot^{\mathcal{I}})$*, a typicality interpretation satisfying* $\mathcal{A}$ *and Properties 1, 2 of Def. 7.6.*

1. *If Algorithm 8.1 returns the typicality interpretation* $\mathcal{I}_n$*, from the input* $\mathcal{K},\mathcal{I}$*, then* $\mathcal{I}_n$ *is the* $\subseteq$*-smallest model in* $\mathtt{TMod}(\mathcal{K},\mathcal{I})$*.*
2. *If Algorithm 8.1 returns* `failure`*, then* $\mathtt{TMod}(\mathcal{K},\mathcal{I}) = \emptyset$

Proof. We assume for Claim 2, that the last created interpretation is also $\mathcal{I}_n$. For both cases, it follows from $\mathcal{I} \subseteq \mathcal{I}_n$ and Lemma 2.12, that $\mathcal{I}_n$ satisfies $\mathcal{A}$ as well as 1, 2 of Def. 7.6.

CLAIM 1. We show satisfaction of $\mathcal{K}$ and minimality w.r.t. $\subseteq$ separately.

STANDARD MODEL. From Lemma 8.5 and the requirements for $\mathcal{I}$, it follows that all $\mathcal{I}_i$ $(0 \leqslant i \leqslant n)$ are standard typicality interpretations. By assumption, the *while*-loop is only entered a finite number of times and Line 6 is never reached. As soon as the *while*-condition (Line 3) fails, i.e. for the returned interpretation $\mathcal{I}_n$, it is clear that $\mathcal{I}_n \in \mathtt{TMod}((\mathcal{T},\mathcal{D}), \Delta^{C,O})$. Because $\mathcal{I}$ already satisfies the ABox $\mathcal{A}$, it follows from $\mathcal{I} \subseteq \mathcal{I}_n$ and Lemma 2.12 that $\mathcal{I}_n \in \mathtt{TMod}(\mathcal{K}, \Delta^{C,O})$.

MINIMALITY. By Lemma 8.7 we know that $\mathtt{TMod}((\mathcal{T},\mathcal{D}),\mathcal{I}_0) = \mathtt{TMod}((\mathcal{T},\mathcal{D}),\mathcal{I}_n)$ and from the definition of standardisation, it is clear that $\mathtt{TMod}((\mathcal{T},\mathcal{D}),\mathcal{I}) = \mathtt{TMod}((\mathcal{T},\mathcal{D}),\mathcal{I}_0)$. Because all models in $\mathtt{TMod}((\mathcal{T},\mathcal{D}),\mathcal{I})$ extend $\mathcal{I}$ and $\mathcal{I}$ satisfies $\mathcal{A}$, it holds that $\mathtt{TMod}((\mathcal{T},\mathcal{D}),\mathcal{I}) = \mathtt{TMod}(\mathcal{K},\mathcal{I})$. Because $\mathcal{I} \subseteq \mathcal{I}_n$, the same argument shows that $\mathtt{TMod}((\mathcal{T},\mathcal{D}),\mathcal{I}_n) = \mathtt{TMod}(\mathcal{K},\mathcal{I}_n)$, and therefore $\mathtt{TMod}(\mathcal{K},\mathcal{I}) = \mathtt{TMod}(\mathcal{K},\mathcal{I}_n)$. As covered in Remark 8.6, for $\mathcal{I}_n \in \mathtt{TMod}(\mathcal{K},\Delta^{C,O})$, it follows that $\mathcal{I}_n \in \mathtt{TMod}(\mathcal{K},\mathcal{I}_n)$ and for all $\mathcal{J} \in \mathtt{TMod}(\mathcal{K},\mathcal{I}_n)$, it holds that $\mathcal{I}_n \subseteq \mathcal{J}$. This proves minimality of $\mathcal{I}_n$ in $\mathtt{TMod}(\mathcal{K},\mathcal{I})$.

CLAIM 2. If Algorithm 8.1 terminates in Line 6 at iteration $n+1$, then there exists some $E \sqsubseteq \bot \in \mathcal{T}$ s.t. $(\chi, \mathcal{U}) \in E^{\mathcal{I}_n}$. From Lemma 2.12 we know that $E^{\mathcal{I}_n} \subseteq E^{\mathcal{J}}$ for every extension $\mathcal{J}$ of $\mathcal{I}_n$ $(\mathcal{I}_n \subseteq \mathcal{J})$. Thus, no such $\mathcal{J}$ satisfies $E \sqsubseteq \bot$, showing $\mathtt{TMod}((\mathcal{T},\mathcal{D}),\mathcal{I}_n) = \emptyset$, which implies $\mathtt{TMod}((\mathcal{T},\mathcal{D}),\mathcal{I}) = \mathtt{TMod}(\mathcal{K},\mathcal{I}) = \emptyset$ by Lemma 8.7. □

Proposition 8.8 shows correctness of Algorithm 8.1 for all terminating executions. Intuitively Algorithm 8.1 achieves exactly what was intended, not only finding some extension of the input interpretation that satisfies the input KB, but the $\subseteq$-smallest such extension. In the following, we will denote the returned interpretation $\mathcal{I}_n$ from running Alg. 8.1 on $\mathcal{K}$, and $\mathcal{I}$ as the *minimal model completion* of $\mathcal{I}$ w.r.t. $\mathcal{K}$ (as in [PT'17a; PT'17b; PT'18]). Formally, we use $\mathtt{mmc}(\mathcal{K},\mathcal{I}) := \mathcal{I}_n$ and carefully include the phrase "if it exists", to account for Alg. 8.1 ending in `failure`.

We return from this excursion on generic model completion to present its implications in the context of typicality preference. As motivated by Proposition 8.3, the following main result of this section links the $\subseteq$-smallest, models of two sets of typicality models with $\mathcal{M}_1 <_t \mathcal{M}_2$, through the use of minimal model completion.

Theorem 8.9. *For a KB $\mathcal{K}$ and a finite typicality domain $\Delta^{C,O}$, let $\mathcal{M}_1, \mathcal{M}_2 \subseteq \mathtt{TMod}(\mathcal{K}, \Delta^{C,O})$ s.t. $\mathcal{M}_1$ is closed under intersection and $\mathcal{M}_1 \neq \emptyset$. Furthermore, let $\mathcal{I}_1 = \bigcap_{\mathcal{J} \in \mathcal{M}_1} \mathcal{J}$ and assume*

$$\forall \mathcal{J} \in \mathtt{TMod}(\mathcal{K}, \Delta^{C,O}).\mathcal{I}_1 \subseteq \mathcal{J} \implies \mathcal{J} \in \mathcal{M}_1. \tag{8.4}$$

For a preference option $r(d,e)$ over $\Delta^{C,O}$ that is Justified and Not-entailed in $\mathcal{M}_1$, let $\mathcal{I}_{r(d,e)} = (\Delta^{C,O}, \cdot^{\mathcal{I}_{r(d,e)}})$, with

$A^{\mathcal{I}_{r(d,e)}} = A^{\mathcal{I}_1}$ *(for $A \in N_C$),*
$a^{\mathcal{I}_{r(d,e)}} = a^{\mathcal{I}_1}$ *(for $a \in N_I$),*
$s^{\mathcal{I}_{r(d,e)}} = s^{\mathcal{I}_1}$ *(for $s \in N_R \setminus \{r\}$), and*
$r^{\mathcal{I}_{r(d,e)}} = r^{\mathcal{I}_1} \cup \{(d,e)\}$.

For the following three statements

1. *$\mathcal{M}_1 <_t \mathcal{M}_2$ with the admissible preference option $r(d,e)$ for $\mathcal{M}_1$ (i.e. $\mathcal{M}_2 = \mathcal{M}_1|_{r(d,e)}$),*
2. *$\mathtt{mmc}(\mathcal{K}, \mathcal{I}_{r(d,e)})$ exists, and*
3. *$\mathtt{mmc}(\mathcal{K}, \mathcal{I}_{r(d,e)})$ is the $\subseteq$-smallest element in $\mathcal{M}_2$,*

it holds that 1 is equivalent to 2; and 3 is implied by either 1 or 2.

Proof. Consider a few notes on the prerequisites of this theorem. By construction and $\mathcal{M}_1$ being closed under intersection, it trivially holds that $\mathcal{I}_1$ is the $\subseteq$-smallest member of $\mathcal{M}_1$. The requirement (8.4) does not affect $\mathcal{M}_1$ being intersection closed, it rather ensures that no arbitrary information can be inferred from $\mathcal{M}_1$ (and $\mathcal{M}_2$, assuming its construction in Statement 1). A full account of (8.4) is not substantial to this proof and is therefore referred to Remark 8.10.

$1 \implies 2$. First of all, $\mathcal{M}_2 \neq \emptyset$ follows from $r(d,e)$ being Satisfiable in $\mathcal{M}_1$. By assumption and model set refinement, it holds that $\mathcal{M}_2 \subseteq \mathcal{M}_1$. Therefore, $\mathcal{I}_1 \subseteq \mathcal{J}$ for all $\mathcal{J} \in \mathcal{M}_2$ implies $\mathcal{I}_{r(d,e)} \subseteq \mathcal{J}$ for all $\mathcal{J} \in \mathcal{M}_2$, by the construction of $\mathcal{M}_2$ and $\mathcal{I}_{r(d,e)}$. Because $\mathcal{M}_2 \subseteq \mathtt{TMod}(\mathcal{K}, \Delta^{C,O})$ it follows by the definition of $\mathtt{TMod}(\mathcal{K}, \mathcal{I}_{r(d,e)})$ (8.1), that $\mathcal{M}_2 \subseteq \mathtt{TMod}(\mathcal{K}, \mathcal{I}_{r(d,e)})$ and thus $\mathtt{TMod}(\mathcal{K}, \mathcal{I}_{r(d,e)}) \neq \emptyset$. Non-emptiness of $\mathtt{TMod}(\mathcal{K}, \mathcal{I}_{r(d,e)})$ ensures that Algorithm 8.1 cannot terminate with `failure` on the input $\mathcal{K}$, $\mathcal{I}_{r(d,e)}$. Its termination on all finite inputs (Prop. 8.16) ensures that $\mathtt{mmc}(\mathcal{K}, \mathcal{I}_{r(d,e)})$ exists.

$2 \implies 1$. From 2 and Requirement (8.4), we conclude $\mathtt{mmc}(\mathcal{K},\mathcal{I}_{r(d,e)}) \in \mathcal{M}_1$. It follows that, in addition to being Justified and Not-entailed in $\mathcal{M}_1$ by assumption, $r(d,e)$ is Satisfiable and therefore admissible w.r.t. $\mathcal{M}_1$. By Definition 7.30, it follows for $\mathcal{M}_2 = \mathcal{M}_1|_{r(d,e)}$ that $\mathcal{M}_1 <_t \mathcal{M}_2$.

$1 \implies 3$. Due to the equivalence of 1 and 2, we can w.l.o.g. assume both. As in the proof for $2 \implies 1$, it holds that $\mathtt{mmc}(\mathcal{K},\mathcal{I}_{r(d,e)}) \in \mathcal{M}_1$. Because $\mathcal{I}_{r(d,e)} \subseteq \mathtt{mmc}(\mathcal{K},\mathcal{I}_{r(d,e)})$, it clearly holds that $\mathtt{mmc}(\mathcal{K},\mathcal{I}_{r(d,e)}) \models r(d,e)$ and thus $\mathtt{mmc}(\mathcal{K},\mathcal{I}_{r(d,e)}) \in \mathcal{M}_2$. Furthermore, because $\mathtt{mmc}(\mathcal{K},\mathcal{I}_{r(d,e)}) \subseteq \mathcal{J}$ holds for all $\mathcal{J} \in \mathtt{TMod}(\mathcal{K},\mathcal{I}_{r(d,e)})$ (by Prop. 8.8) and $\mathcal{M}_2 \subseteq \mathtt{TMod}(\mathcal{K},\mathcal{I}_{r(d,e)})$ it follows that $\mathtt{mmc}(\mathcal{K},\mathcal{I}_{r(d,e)}) \subseteq \mathcal{J}$ for all $\mathcal{J} \in \mathcal{M}_2$. □

Even though Theorem 8.9 seems rather restricted in terms of its prerequisites, all of those requirements are naturally met for executions of Algorithm 8.1 during the upcoming typicality maximisation procedure. For now, consider the following remark about Requirement (8.4) in Theorem 8.9.

Remark 8.10. Assuming the set $\mathcal{M}_1$ in Thm. 8.9 to be "as large as possible", in terms of the typicality models extending its $\subseteq$-smallest model, has the support of two arguments. For one, this result does not hold for arbitrary sets $\mathcal{M}_1, \mathcal{M}_2 \subseteq \mathtt{TMod}(\mathcal{K},\Delta^{C,O})$. Consider $\mathcal{M}_1 = \{\mathcal{I}_1, \mathcal{I}_2\}$ and $\mathcal{M}_2 = \{\mathcal{I}_2\}$ s.t. $\mathcal{I}_1 \not\models r(d,e)$, $\mathcal{I}_2 \models r(d,e)$ and $\mathcal{I}_1 \subseteq \mathcal{I}_2$. Obviously, $\mathcal{M}_1$ is intersection closed with $\mathcal{I}_1$ as its $\subseteq$-smallest member. However, $\mathcal{I}_2$ could—aside from satisfaction of $r(d,e)$ and $\mathcal{I}_1 \subseteq \mathcal{I}_2$—contain arbitrary information, that is neither contained in $\mathcal{I}_1$, nor in *all* model extensions of $\mathcal{I}_1$ satisfying $r(d,e)$. For $\mathcal{I}_1 \subseteq \mathcal{I}_{r(d,e)}$ (as in Thm. 8.9), the minimal model completion $\mathtt{mmc}(\mathcal{K},\mathcal{I}_{r(d,e)})$ would then not contain such arbitrary information. Specifically, $\mathtt{mmc}(\mathcal{K},\mathcal{I}_{r(d,e)}) \subsetneq \mathcal{I}_2$ and thus, $\mathtt{mmc}(\mathcal{K},\mathcal{I}_{r(d,e)}) \notin \mathcal{M}_2$.

In addition, Theorem 8.9 will eventually be applied to such $\mathcal{M}_1$ and $\mathcal{M}_2$ that are members of a preference chain

$$\mathtt{TMod}(\mathcal{K},\Delta^{C,O}) <_t \cdots <_t \mathcal{M}_1 <_t \mathcal{M}_2 <_t \cdots .$$

Requirement (8.4) is naturally satisfied for $\mathtt{TMod}(\mathcal{K},\Delta^{C,O})$ and, by construction of each $\mathcal{M}$ in such a chain, it will also be naturally satisfied by those $\mathcal{M}$.

8.1.2 *Typicality Maximisation*

Minimal model completion, in particular Theorem 8.9, essentially provides two tools. Proposition 8.3 motivates, that consequences of a set of models $\mathcal{M}_i$ on a chain

$$\mathtt{TMod}(\mathcal{K},\Delta^{C,O}) <_t \mathcal{M}_1 <_t \mathcal{M}_2 <_t \cdots \tag{8.5}$$

can be determined with the $\subseteq$-smallest member of $\mathcal{M}_i$. Transitivity of Prop. 8.3 implies that such a smallest model exists in $\mathcal{M}_i$, if the chain begins at $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$. Theorem 8.9 shows that for two sets $\mathcal{M}_i$, $\mathcal{M}_{i+1}$ in this chain, the smallest member of $\mathcal{M}_{i+1}$ can be constructed algorithmically from the smallest model in $\mathcal{M}_i$. Once again, by transitivity, if this preference chain begins with $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$, then it induces a chain of smallest typicality models, beginning with $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$. Together with the direct construction of $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ in Corollary 7.23 by classical reasoning, this describes an algorithmic construction of the $\subseteq$-smallest model in any Max-TM. In the following, we will formalise this construction.

The algorithm we propose for typicality maximisation, effectively generates the chain of (smallest) typicality models induced by a preference chain. In general, starting from any given typicality model, Algorithm 8.2 proceeds to determine and select admissible preference options and incrementally extends the current typicality model by constructing the next smallest model that satisfies the selected preference option, using Algorithm 8.1. It continues this iteration until no more admissible preference options exist.

The crux in Alg. 8.2 is to determine admissibility of preference options. This is where the second, less obvious use of Theorem 8.9 is significant. Suppose $\mathcal{M}$ is any member of the chain in (8.5) and $\mathcal{I} \in \mathcal{M}$ is its $\subseteq$-smallest element. Obviously, the set of all preference options, $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$, is fixed for all interpretations over $\Delta^{\mathbb{C},\mathbb{O}}$. Also, for every preference option $p \in \mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$, we can determine its *Justified* and *Not-entailed* property w.r.t. $\mathcal{M}$, directly from $\mathcal{I}$. If the non-typical variant[2] of p is present in $\mathcal{I}$, then it is present in *all* models in $\mathcal{M}$ (Justified). If p itself is *not* present in $\mathcal{I}$, then $\mathcal{I}$ is directly a witness to show that p is Not-entailed in $\mathcal{M}$. The property *Satisfiable* is more involved to establish for p w.r.t. $\mathcal{M}$, because from $\mathcal{I}$, no conclusions about existence of information (role edges) in *some* model in $\mathcal{M}$ can be drawn. The implication $2 \implies 1$ in Theorem 8.9 provides a solution for this problem. Specifically, $p \in \mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ is admissible in $\mathcal{M}$, if it is Justified, Not-entailed and the minimal model completion of $\mathcal{I}$ extended by p exists (Satisfiable).

For brevity in Algorithm 8.2 and hereinafter, we refer to $\mathcal{I}[p]$ as the $\subseteq$-smallest extension of $\mathcal{I}$, satisfying the preference option $p = r(d, e)$, i.e. as in the construction of $\mathcal{I}_{r(d,e)}$ in Theorem 8.9.

Algorithm 8.2 relies additionally on the input of a partial function po. This function is used generically to select preference options for the sets of interpretations whose smallest members are the $\mathcal{I}_i$ within the algorithm. It allows to use Alg. 8.2 for deciding entailments under both sceptical and selective nested semantics, by appropriately instantiating po. We say the function po is *admissible iff*

1. $\mathsf{po}(\mathcal{J})$ is an admissible preference option for $\mathtt{TMod}(\mathcal{K}, \mathcal{J})$, and
2. $\mathsf{po}(\mathcal{J})$ is undefined *iff* $\mathtt{TMod}(\mathcal{K}, \mathcal{J})$ has no admissible preference options,

for all $\mathcal{J} \in \mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$. As usual, before considering explicit instantiations, we present general results for Algorithm 8.2.

2 For $p = r(d, (C, \mathcal{U}))$, its non-typical variant is $p' = r(d, (C, \emptyset))$.

Algorithm 8.2: Maximising Typicality

Input: KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, typicality interpretation $\mathcal{I} = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\mathcal{I}})$,
partial function $\mathsf{po} : \mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \nrightarrow \mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$

Output: typicality interpretation $\mathcal{I}_n$

1 $\mathcal{I}_0 := \mathcal{I}$;
2 $i := 0$;
3 **while** $\exists$ *admissible preference option for* $\mathtt{TMod}(\mathcal{K}, \mathcal{I}_i)$ **do**
4 $\quad p := \mathsf{po}(\mathcal{I}_i)$;
5 $\quad \mathcal{I}_{i+1} := \mathtt{mmc}(\mathcal{K}, \mathcal{I}_i[p])$;
6 $\quad i := i + 1$;
7 **end**
8 **return** $\mathcal{I}_i$;

The following intermediary lemma is a useful tool, connecting the sets of typicality models whose $\subseteq$-smallest members are interpretations $\mathcal{I}_i$, $\mathcal{I}_{i+1}$, during a run of Algorithm 8.2.

Lemma 8.11. *Let Algorithm 8.2 return* $\mathcal{I}_n$*, running on the inputs* $\mathcal{K}$, $\mathcal{I}$, *and an admissible* po*. It holds for* $0 \leqslant i < n$ *that*

$$\mathtt{TMod}(\mathcal{K}, \mathcal{I}_{i+1}) = \mathtt{TMod}(\mathcal{K}, \mathcal{I}_i)|_p$$

for p *selected at Line 4.*

Proof. We provide arguments for both inclusions of this claim separately. First of all, Proposition 8.8 implies that for p being admissible, $\mathtt{mmc}(\mathcal{K}, \mathcal{I}_i[p])$ exists. From $\mathcal{I}_i[p] \subseteq \mathcal{I}_{i+1}$, it follows that $\mathcal{I}_{i+1}$ and all models that extend $\mathcal{I}_{i+1}$ clearly belong to $\mathtt{TMod}(\mathcal{K}, \mathcal{I}_i)|_p$ (proving "$\subseteq$").

Because $\mathcal{I}_{i+1}$ is the $\subseteq$-smallest model that extends $\mathcal{I}_i$ and satisfies p (by Proposition 8.8), every model $\mathcal{J}$ in $\mathtt{TMod}(\mathcal{K}, \mathcal{I}_i)|_p$, must be an extension of $\mathcal{I}_{i+1}$ (proving "$\supseteq$"). □

We show that Algorithm 8.2 returns the $\subseteq$-smallest model of a Max-TM in $\mathtt{TMax}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ when given appropriate inputs.

Theorem 8.12. *For a KB* $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$*, a finite, consistent, and quantification closed context* $\mathbb{C}$, $\mathbb{O}$*, containing* $\mathcal{K}$*, and an admissible partial function* $\mathsf{po} : \mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \nrightarrow \mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$*, let* $\mathcal{I}_n$ *be the interpretation returned by Algorithm 8.2 on the input* $\mathcal{K}$, $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$, po*. It holds that*

1. $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \leqslant_t^* \mathtt{TMod}(\mathcal{K}, \mathcal{I}_n)$*, and*
2. $\neg\exists \mathcal{M} \subseteq \mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}).\mathtt{TMod}(\mathcal{K}, \mathcal{I}_n) <_t \mathcal{M}$.

Proof.

CLAIM 1. The chain of interpretations $\mathcal{I}_0, \mathcal{I}_1, \cdots, \mathcal{I}_n$ in Alg. 8.2 starting from $\mathcal{I}_0 = \mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$, provides the chain of sets of typicality models

$$\mathtt{TMod}(\mathcal{K}, \mathcal{I}_0), \mathtt{TMod}(\mathcal{K}, \mathcal{I}_1), \cdots, \mathtt{TMod}(\mathcal{K}, \mathcal{I}_n).$$

It is clear by the definition of $\mathsf{TMod}(\mathcal{K},\mathcal{I}_i)$ (see (8.1)), that for all $0 \leqslant i \leqslant n$, $\mathsf{TMod}(\mathcal{K},\mathcal{I}_i)$ is closed under intersection and satisfies (8.4) in Theorem 8.9. Note that the *while*-condition in Line 3 can be checked using Alg. 8.1 and Thm. 8.9. Also, by assumption on po, if the *while*-condition is satisfied, the selection of p in Line 4 is well-defined and the selected p is admissible for $\mathsf{TMod}(\mathcal{K},\mathcal{I}_i)$.

From Lem. 8.11, Def. 7.30 and the fact that p is admissible for $\mathsf{TMod}(\mathcal{K},\mathcal{I}_i)$, it follows that $\mathsf{TMod}(\mathcal{K},\mathcal{I}_i) <_t \mathsf{TMod}(\mathcal{K},\mathcal{I}_{i+1})$ holds for all $0 \leqslant i < n$. Consequently, $\mathsf{TMod}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}}) \leqslant_t^* \mathsf{TMod}(\mathcal{K},\mathcal{I}_n)$, because $\mathsf{TMod}(\mathcal{K},\mathcal{I}_0) = \mathsf{TMod}(\mathcal{K},\mathsf{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})) = \mathsf{TMod}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})$.

CLAIM 2. For the returned model $\mathcal{I}_n$, the *while*-condition in Line 3 clearly failed, hence by Definition 7.30, there cannot be a set of models that is more preferred than $\mathsf{TMod}(\mathcal{K},\mathcal{I}_n)$ w.r.t. $<_t$. □

Theorem 8.12 shows that Algorithm 8.2 provides the means to determine the $\subseteq$-smallest member $\mathcal{I}$ of a set $\mathcal{M} \in \mathsf{TMax}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})$. From Lemma 7.11 we know for a defeasible query α and this $\mathcal{M}$ and $\mathcal{I}$, that $\mathcal{M} \models \alpha$ *iff* $\mathcal{I} \models \alpha$. Recall the general characterisation of nested semantics from Definition 7.10 and 7.38. For sceptical nested semantics, entailment of a query α needs to be verified by all sets of models in $\mathsf{TMax}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})$. Therefore, non-entailment of a query α holds if a single Max-TM invalidates α. Entailments under selective nested semantics, only require to check a single element of $\mathsf{TMax}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})$, that is *uniquely identified* by a preference relation $<_{\mathsf{po}}$ over all preference options $\mathsf{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ (cf. Rem. 7.37). Both of these checks can be implemented by specific instantiations of po for the input of Algorithm 8.2.

SCEPTICAL NESTED SEMANTICS. We capture a complexity upper bound for deciding entailment under sceptical nested semantics in terms of the complement of this decision procedure. To decide non-entailment of a query, only a single counterexample in a single $\mathcal{M} \in \mathsf{TMax}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})$ needs to be found. This Max-TM $\mathcal{M}$ cannot be determined constructively without enumerating all (exponentially many) Max-TMs. However, a non-deterministic admissible partial function can allow a run of Alg. 8.2 to guess the construction of this $\mathcal{M}$. Formally, we define the partial function $\mathsf{po}_\mathsf{N} : \mathsf{TMod}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}}) \not\rightarrow \mathsf{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ as follows. Suppose $\mathsf{R} \subseteq \mathsf{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ is the set of all admissible preference options for $\mathsf{TMod}(\mathcal{K},\mathcal{I})$.

If $\mathsf{R} = \emptyset$ then $\mathsf{po}_\mathsf{N}(\mathcal{I})$ is undefined and otherwise, $\mathsf{po}_\mathsf{N}(\mathcal{I})$ guesses a $p \in \mathsf{R}$.

$\mathsf{po}_\mathsf{N}()$ is clearly an admissible partial function, hence using it as an input for Algorithm 8.2 does not invalidate Theorem 8.12. We show that the resulting algorithmic characterisation of sceptical nested semantics is sound and complete.

Proposition 8.13. *Let $\mathbb{C},\mathbb{O}$ be a finite, consistent, and quantification closed context, containing a KB $\mathcal{K}$. Let $\mathsf{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})$ be the minimal typicality model* *for $\mathcal{K}$ over $\Delta^{\mathbb{C},\mathbb{O}}$, and α a defeasible subsumption or instance query. The following are equivalent*

1. $\exists \mathcal{M} \in \texttt{TMax}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}).\mathcal{M} \not\models \alpha$,
2. *Algorithm 8.2 (on input* $\mathcal{K}$, $\texttt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$, $\texttt{po}_{\texttt{N}}$*)* can *return* $\mathcal{I}_n$ *s.t.* $\mathcal{I}_n \not\models \alpha$.

Proof. The implication 2 $\implies$ 1 follows quickly from Theorem 8.12. If Algorithm 8.2 returns $\mathcal{I}_n$, then $\texttt{TMod}(\mathcal{K}, \mathcal{I}_n) \in \texttt{TMax}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ and $\mathcal{I}_n \not\models \alpha$ implies $\texttt{TMod}(\mathcal{K}, \mathcal{I}_n) \not\models \alpha$.

For the implication 1 $\implies$ 2, recall Remark 7.37, stating that every full preference chain

$$\texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) <_t \mathcal{M}_1 <_t \cdots <_t \mathcal{M}_n$$

(by assumption $\mathcal{M}_n \not\models \alpha$), induces a chain of preference options

$$p_0, p_1, \cdots, p_{n-1}$$

such that each p_i is admissible for $\mathcal{M}_i$ ($i \geqslant 0$, in the following, let $\mathcal{M}_0 = \texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$). Let $\mathsf{R}_i \subseteq \texttt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ be the set of admissible preference options for $\mathcal{M}_i$ ($0 \leqslant i \leqslant n$). It holds for $0 \leqslant i < n$, that $p_i \in \mathsf{R}_i$ and $\mathsf{R}_n = \emptyset$. We show by induction on i, that Algorithm 8.2 can produce a chain of typicality interpretations

$$\mathcal{I}_0, \mathcal{I}_1, \cdots, \mathcal{I}_n$$

where $\mathcal{I}_0 = \texttt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ and $\mathcal{I}_n$ is returned in Line 8, such that $\mathcal{M}_i = \texttt{TMod}(\mathcal{K}, \mathcal{I}_i)$ ($1 \leqslant i \leqslant n$). For the induction start, $\texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ is clearly equivalent to $\texttt{TMod}(\mathcal{K}, \texttt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}))$, by the definition of $\texttt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ (Def. 7.13).

Suppose $\mathcal{M}_i = \texttt{TMod}(\mathcal{K}, \mathcal{I}_i)$. We show that $\mathcal{I}_{i+1}$ can be constructed (Line 5) such that $\mathcal{M}_{i+1} = \texttt{TMod}(\mathcal{K}, \mathcal{I}_{i+1})$. For $i < n$, $\mathcal{M}_i$ is not maximally preferred w.r.t. $<_t$, specifically, it has at least the admissible preference option p_i and so does $\texttt{TMod}(\mathcal{K}, \mathcal{I}_i)$ by the induction hypothesis. Because p_i is admissible for $\texttt{TMod}(\mathcal{K}, \mathcal{I}_i)$, it is possible for $\texttt{po}_{\texttt{N}}$ to select p_i in Line 4 and construct $\mathcal{I}_{i+1} := \texttt{mmc}(\mathcal{K}, \mathcal{I}_i[p_i])$ in Line 5. From Lemma 8.11 and the selection of p_i, we know that

$$\texttt{TMod}(\mathcal{K}, \mathcal{I}_{i+1}) = \texttt{TMod}(\mathcal{K}, \mathcal{I}_i)|_{p_i} = \mathcal{M}_i|_{p_i} = \mathcal{M}_{i+1}.$$

Thus, it follows that $\mathcal{M}_n = \texttt{TMod}(\mathcal{K}, \mathcal{I}_n)$. $\mathsf{R}_n = \emptyset$ implies that the *while*-condition (Line 3) fails for $\texttt{TMod}(\mathcal{K}, \mathcal{I}_n)$. Thus, Algorithm 8.2 *can* return $\mathcal{I}_n$ and if $\mathcal{M}_n \not\models \alpha$, then by Lemma 7.11 it follows that $\mathcal{I}_n$, as the $\subseteq$-smallest element of $\mathcal{M}_n$, also does not satisfy α. □

Keep in mind that this instantiates only the coverage of the resulting semantics. The obtained complexity upper bounds (Sec. 8.2) also vary for running Algorithm 8.1 and 8.2 on the rational or relevant domain.

SELECTIVE NESTED SEMANTICS. To determine entailments under selective nested semantics, an inquirer must supply the additional input of a total order over the set of preference options $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ for the underlying typicality domain $\Delta^{\mathbb{C},\mathbb{O}}$. In practice, it is likely not feasible to explicitly define this order, especially when the typicality domain is exponentially large in the input, as for $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec}$. One could imagine more practicable approaches, such as defining *only a strict partial order* over $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$, explicitly relating only role edges for which a preference is sensible or required. Such a strict partial order could then be automatically completed through "don't care" non-determinism into a total order. Another idea might be to define two or more total orders on role, concept and individual names, sets of DCIs, etc., and to derive a unique total order over $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ from those (smaller) relations. In favour of generality, we leave such considerations open for future work and the practitioners. We proceed—as with the total order on $\mathsf{N_I}$ for ABox extensions—assuming $<_{\mathsf{po}}$ to be part of the input for Algorithm 8.2.

From Remark 7.37 we know that a total order $<_{\mathsf{po}}$ over $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ uniquely determines a set $\mathcal{M} \in \mathtt{TMax}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$, i.e. $\mathcal{M} = \mathtt{TMod}_{\mathtt{nest}_{<_{\mathsf{po}}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$. This $\mathcal{M}$ can be deterministically constructed by Algorithm 8.2, given an appropriate partial function for selecting preference options in Line 4. The partial function $\mathsf{po}_{<_{\mathsf{po}}} : \mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \not\rightarrow \mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ is defined as follows. For a typicality interpretation $\mathcal{I}$, let $\mathsf{R} \subseteq \mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ be the set of admissible preference options for $\mathtt{TMod}(\mathcal{K}, \mathcal{I})$.

If $\mathsf{R} = \emptyset$, $\mathsf{po}_{<_{\mathsf{po}}}(\mathcal{I})$ is undefined, otherwise $\mathsf{po}_{<_{\mathsf{po}}}(\mathcal{I}) = \mathfrak{p}$, for $\mathfrak{p}$ being the $<_{\mathsf{po}}$-minimal element in R.

First of all, for a finite $\Delta^{\mathbb{C},\mathbb{O}}$, R is finite and, if it is non-empty, its $<_{\mathsf{po}}$-minimal element exists and is unique. Therefore, $\mathsf{po}_{<_{\mathsf{po}}}$ is clearly admissible as well as deterministic. Thus, a run of Algorithm 8.2 on the input $\mathsf{po}_{<_{\mathsf{po}}}$ will also be deterministic.

In contrast to sceptical semantics—where it was necessary to prove that any Max-TM can be reached (using $\mathsf{po_N}$)—for selective semantics, we need to show that the deterministic run of Algorithm 8.2 (using $\mathsf{po}_{<_{\mathsf{po}}}$), produces the $\subseteq$-smallest element in $\mathtt{TMod}_{\mathtt{nest}_{<_{\mathsf{po}}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$.

Proposition 8.14. *For a finite, consistent, and quantification closed context $\mathbb{C},\mathbb{O}$, containing a KB $\mathcal{K}$, a typicality domain $\Delta^{\mathbb{C},\mathbb{O}}$ and a total order $<_{\mathsf{po}}$ over $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$, let Algorithm 8.2 return $\mathcal{I}_n$ on the input $\mathcal{K}$, $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ and $\mathsf{po}_{<_{\mathsf{po}}}$. It holds that*

1. $\mathcal{I}_n \in \mathtt{TMod}_{\mathtt{nest}_{<_{\mathsf{po}}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$, *and*
2. $\forall \mathcal{J} \in \mathtt{TMod}_{\mathtt{nest}_{<_{\mathsf{po}}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) \,.\, \mathcal{I}_n \subseteq \mathcal{J}$.

Proof. In favour of comprehensibility, let $\mathcal{M}_k = \mathtt{TMod}_{\mathtt{nest}_{<_{\mathsf{po}}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ and $\mathcal{M}_0 = \mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$. Let $\mathcal{M}_0 <_t \mathcal{M}_1 <_t \cdots <_t \mathcal{M}_k$ be the full preference chain that is compliant with $<_{\mathsf{po}}$. Such a chain exists and is unique, by definition of $\mathtt{TMod}_{\mathtt{nest}_{<_{\mathsf{po}}}}()$. Given $<_{\mathsf{po}}$, there is a unique chain of preference options $p_0, p_1, \cdots, p_{k-1}$ induced by the above preference chain, s.t. $\mathcal{M}_{i+1} = \mathcal{M}_i|_{p_i}$ and p_i is the $<_{po}$-minimal admissible preference option for $\mathcal{M}_i$ (for $0 \leqslant i < k$). Note that we deliberately use distinct indices k and n, because we cannot assume $n = k$. We show by induction on i, for $i < n$ and $i < k$, that $\mathtt{TMod}(\mathcal{K}, \mathcal{I}_i) = \mathcal{M}_i$, for $\mathcal{I}_i$ as constructed in Algorithm 8.2 on the input $\mathcal{K}$, $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ and $\mathtt{po}_{<_{po}}$.

For $i = 0$, $\mathcal{I}_0 = \mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ and $\mathtt{TMod}(\mathcal{K}, \mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})) = \mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ holds by definition of the minimal typicality model (Def. 7.13). Suppose $\mathtt{TMod}(\mathcal{K}, \mathcal{I}_i) = \mathcal{M}_i$ holds when $i < n, k$. Let R be the set of admissible preference options for $\mathcal{M}_i$. Because R is also the set of admissible preference options for $\mathtt{TMod}(\mathcal{K}, \mathcal{I}_i)$, and $\mathsf{R} \neq \emptyset$ (because $i < k$), the function $\mathtt{po}_{<_{po}}$ selects p_i, the $<_{po}$-minimal element in R, in Line 4. From Lemma 8.11 it follows that $\mathtt{TMod}(\mathcal{K}, \mathcal{I}_{i+1}) = \mathtt{TMod}(\mathcal{K}, \mathcal{I}_i)|_{p_i} = \mathcal{M}_i|_{p_i} = \mathcal{M}_{i+1}$. It follows that $\mathtt{TMod}(\mathcal{K}, \mathcal{I}_i) = \mathcal{M}_i$ holds for $i \leqslant n, k$.

Assume for a contradiction that $k < n$. It follows that $\mathcal{M}_k$ has no admissible preference options, hence the *while*-condition (Line 3) fails for $\mathtt{TMod}(\mathcal{K}, \mathcal{I}_k)$, contradicting that Algorithm 8.2 returned $\mathcal{I}_n$. An analogous argument can be made to contradict $n < k$, therefore $n = k$ must hold, i.e. $\mathcal{M}_k = \mathtt{TMod}(\mathcal{K}, \mathcal{I}_n)$. Both claims of this proposition follow directly from $\mathtt{TMod}_{\mathtt{nest}_{<_{po}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}}) = \mathtt{TMod}(\mathcal{K}, \mathcal{I}_n)$. □

Proposition 8.14 proves that a deterministic run of Algorithm 8.2 produces the $\subseteq$-smallest element in $\mathtt{TMod}_{\mathtt{nest}_{<_{po}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$. As before, this model is canonical for $\mathtt{TMod}_{\mathtt{nest}_{<_{po}}}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ and therefore, can be used to determine entailments under selective nested semantics.

We continue to briefly and collectively instantiate decision procedures for sceptical and selective nested semantics, coupled with both, rational and relevant strength of reasoning. This concludes the algorithmic characterisation of all nested semantics in this thesis.

RATIONAL AND RELEVANT STRENGTH. As before, the difference between rational and relevant strength is manifested in the construction of the underlying typicality domain. Therefore, the specific structure of this domain plays no part in the algorithmic characterisation of Max-TMs. Constructing the minimal typicality model for the rational or relevant domain of a finite, consistent context can be achieved with a straightforward algorithm relying on classical reasoning, as encouraged by Corollary 7.23. The following explicitly shows that entailments under sceptical and selective, rational and relevant semantics can effectively be decided with Algorithm 8.2. It is an immediate consequence of Proposition 8.13, 8.14 and Theorem 8.12.

Corollary 8.15. *Let $\mathcal{K}$ be a KB, $\mathbb{C}$, $\mathbb{O}$ a finite, consistent, quantification closed context, containing $\mathcal{K}$, $\prec$ a total order on N_I, $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_\prec}$ and $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec}$ the rational and relevant typicality domains with $<^{\mathtt{rat}}_{po}$ and $<^{\mathtt{rel}}_{po}$ as total orders over $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_\prec})$ and $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec})$, respectively. For a defeasible subsumption or instance check α:*

1. $\mathcal{K} \not\models^{(\mathsf{s}_\prec,\mathtt{nest})} \alpha$ iff *Alg. 8.2 can return* $\mathcal{I}_n$ *on input* $\mathcal{K}$, $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathsf{O},\mathcal{K}}_{\mathbf{s}_\prec})$, $\mathtt{po}_{\mathtt{N}}$ *s.t.* $\mathcal{I}_n \not\models \alpha$

2. $\mathcal{K} \models^{(\mathsf{s}_\prec,\mathtt{nest}_{<^{\mathsf{s}}_{\mathsf{po}}})} \alpha$ iff *Alg. 8.2 returns* $\mathcal{I}_n$ *on input* $\mathcal{K}$, $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathsf{O},\mathcal{K}}_{\mathbf{s}_\prec})$, $\mathtt{po}_{<^{\mathsf{s}}_{\mathsf{po}}}$ *s.t.* $\mathcal{I}_n \models \alpha$

for $\mathbf{s} \in \{\mathtt{rat}, \mathtt{rel}\}$.

Once more, both items in Corollary 8.15 are impartial to the specific instantiation of $\mathbf{s}$. Recall the bigger picture. We specifically investigate rational strength for its solid foundation, decades of fortitude among a manifold of approaches to defeasible reasoning and its continued impact in this field. Relevant strength, we investigate for its ability to overcome inheritance blocking and for its strengthening of Rational Closure overall. Considering the generality of our construction and the characterisation of propositional and nested coverage through typicality models, both rational and relevant strength are mere examples for the power of typicality models. Other approaches to overcome inheritance blocking, such as Lexicographic Closure, could potentially be reconstructed in the typicality model paradigm. This supports our claim that typicality models can be employed for any Description Logic with the canonical model property, allowing to capture a variety of differently strong and nested defeasible semantics.

We continue to investigate termination and runtime of Algorithm 8.1 and 8.2, first, for general typicality domains. Finally, the specific complexity upper bound for rational and relevant strength, including the construction of their respective typicality domains, instantiates these generic results.

8.2 COMPLEXITY UPPER BOUNDS

Using the algorithms in Section 8.1, we provide complexity upper bounds for deciding entailments under all semantics identified through

$$\{\mathtt{rat}_\prec, \mathtt{rel}_\prec\} \times \{\mathtt{prop}, \mathtt{nest}, \mathtt{nest}_{<_{\mathsf{po}}}\}.$$

The runtime of Algorithm 8.1 and 8.2 can be analysed in terms of their generic inputs. More specifically, we initially provide complexity upper bounds that are parametric on the size of the typicality domain on which the input interpretations in Alg. 8.1 and 8.2 are based. To achieve rational and relevant strength, these algorithms are finally called with typicality models over the rational and relevant domain, respectively. Their size can be determined in terms of the other inputs, i.e. the given KB and context. Before investigating those typicality domains specifically, a few complexity results of the underlying materialisation foundations, such as computation of the rational chain, are required. Instantiation of the general runtime results for both algorithms, then provides specific complexity upper bounds for the decision problems of entailment under the discussed semantics.

MODEL COMPLETION. As introduced in Section 8.1.1, the main idea behind model completion is a systematic rectification for each counterexample, i.e. pair of domain element and GCI or DCI, in the input interpretation. While every iteration of the *while*-loop in Algorithm 8.1 fixes (at least[3]) one such counterexample, it can easily introduce new counterexamples. Nevertheless, the maximum number of possible counterexamples is polynomially bounded by the size of the input KB and the typicality domain of the input interpretation. The main idea in the following proof, is that each such counterexample may occur at most once during an entire run of Algorithm 8.1.

Proposition 8.16. *For a KB $\mathcal{K}$, a context $\mathbb{C}, \mathbb{O}$, containing $\mathcal{K}$, and a finite $\mathcal{I} = (\Delta^{\mathbb{C},\mathbb{O}}, \cdot^{\mathcal{I}})$ (satisfying 1 and 2 of Def. 7.6), Algorithm 8.1 always terminates and runs in polynomial time in $|\Delta^{\mathbb{C},\mathbb{O}}|$ and $|\mathcal{K}|$ on the input $\mathcal{K}, \mathcal{I}$.*

Proof. Checking the *while*-condition in Line 3 requires for every element $(\chi, \mathcal{U}) \in \Delta^{\mathbb{C},\mathbb{O}}$ to verify $(\chi, \mathcal{U}) \in E^{\mathcal{I}_i} \implies (\chi, \mathcal{U}) \in F^{\mathcal{I}_i}$ for every $E \bowtie F \in \mathcal{T} \cup \mathcal{U}$, i.e. at most $|\Delta^{\mathbb{C},\mathbb{O}}| * (|\mathcal{T}| + |\mathcal{D}|)$ checks. If the loop is entered, then a counterexample, as described in Line 4, has already been found and can be used at no extra computational cost. It is not hard to see that the order in which such elements in Line 4 are chosen is irrelevant for the outcome of the algorithm, hence rendering Alg. 8.1 deterministic, modulo "don't care" non-determinism. From Definition 8.4, it is clear that the standard promotion in Line 8 as well as the standardisation in Line 1 can be determined in polynomial time in the input.

Clearly, $\mathcal{I}_i \subseteq \mathcal{I}_{i+1}$ $(0 \leqslant i)$ holds. Therefore, if $d \in \Delta^{\mathbb{C},\mathbb{O}}$ and $E \bowtie F \in \mathcal{T} \cup \mathcal{U}$ are selected in iteration i (Line 4), then by Lemma 8.5 it holds that $d \in F^{\mathcal{I}_j}$ for all $j > i$. Hence the pair $d, E \bowtie F$ can never be selected at Line 4 in any iteration $j > i$. Therefore, for finite $\Delta^{\mathbb{C},\mathbb{O}}$ and $\mathcal{T} \cup \mathcal{D}$, an upper bound for the number of *while*-iterations in Algorithm 8.1 is $n = |\Delta^{\mathbb{C},\mathbb{O}}| * (|\mathcal{T}| + |\mathcal{D}|)$. Clearly, if the pool of pairs d, $E \bowtie F$ (i.e. at most n pairs) is depleted after some iteration, the *while*-condition must fail, and the algorithm terminates.

For every line taking at most polynomial time to compute and every line being executed only a polynomial number of times (at most n), it follows that Algorithm 8.1 always terminates on finite inputs $\mathcal{K}, \mathcal{I}$ in polynomial time in $|\Delta^{\mathbb{C},\mathbb{O}}| + |\mathcal{K}|$. □

TYPICALITY MAXIMISATION. A runtime analysis of Algorithm 8.2 is slightly more involved than that for model completion. The following lemma shows two important properties about the *while*-loop in Alg. 8.2. These properties are essential to determine the runtime of the typicality maximisation algorithm. In particular, the second property shows that its runtime is influenced by the complexity of determining the selection of

3 As seen in Example 7.33, altering an interpretation can have implicit consequences on the models that encompass it. Perhaps the domain element selected in the current iteration is also the counterexample to another GCI that is implicitly resolved by the current extensions.

$\text{po}(\mathcal{I}_i)$ in Line 4. The actual runtime (and non-/determinism) of Alg. 8.2 is therefore influenced by the input partial function, and must be analysed for every instantiation of po() explicitly.

Lemma 8.17. *Let $\mathbb{C}, \mathbb{O}$ be a context containing the KB $\mathcal{K}$, $\mathcal{I}$ be a finite typicality model $\mathcal{I} \in \texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ (satisfying 1 and 2 of Def. 7.6) and* po *be an admissible partial function. During a run of Algorithm 8.2 on the input $\mathcal{K}, \mathcal{I},$* po *the following holds.*

1. *The* while*-condition (Line 3) can be determined in polynomial time in $|\Delta^{\mathbb{C},\mathbb{O}}|$ and $|\mathcal{K}|$.*

2. *The* while*-loop (Line 3–7) is entered at most a polynomial number of times in $|\Delta^{\mathbb{C},\mathbb{O}}|$.*

Proof. For a finite typicality domain $\Delta^{\mathbb{C},\mathbb{O}}$, the set $\texttt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$ is finite and has the (loose) upper bound $m = |\Delta^{\mathbb{C},\mathbb{O}}|^2 * sig_R(\mathcal{K}, \mathbb{C})$.

CLAIM 1. Naively, it is enough to check for (up to) m preference options, whether they satisfy the properties for admissibility w.r.t. $\texttt{TMod}(\mathcal{K}, \mathcal{I}_i)$. From the definition of $\texttt{TMod}(\mathcal{K}, \mathcal{I}_i)$ (see (8.1)), it is clear that Justified and Not-entailed (Def. 7.27) can be determined directly with $\mathcal{I}_i$.[4] The technique was described in the beginning of Section 8.1.2 (Page 125). For the property Satisfiable, it follows from Theorem 8.9, that for any preference option $p \in \texttt{PO}(\Delta^{\mathbb{C},\mathbb{O}})$, $\texttt{mmc}(\mathcal{K}, \mathcal{I}_i[p])$ exists *iff* p is admissible for $\texttt{TMod}(\mathcal{K}, \mathcal{I}_i)$. Existence of $\texttt{mmc}(\mathcal{K}, \mathcal{I}_i[p])$ can be determined in polynomial time in $|\Delta^{\mathbb{C},\mathbb{O}}|$ and $|\mathcal{K}|$ (by Prop. 8.16). Therefore, at most m calls to Algorithm 8.1 are required to check the *while*-condition in Line 3, requiring overall polynomial time in $|\Delta^{\mathbb{C},\mathbb{O}}|$ and $|\mathcal{K}|$.

CLAIM 2. If po is admissible, then p is an admissible preference option for $\texttt{TMod}(\mathcal{K}, \mathcal{I}_i)$ (at any iteration i). From $\mathcal{I}_i[p] \subseteq \mathcal{I}_{i+1}$ it follows that the property *Not-entailed* (Def. 7.27) can never be satisfied for p w.r.t. any $\texttt{TMod}(\mathcal{K}, \mathcal{I}_j)$ for $j > i$. Hence, the *while*-loop can only be entered at most once for every (admissible) preference option, providing the upper bound on the number of iterations with m. □

Naturally, the actual number of calls to Algorithm 8.1 when checking the *while*-condition in Algorithm 8.2 is much lower than m, as only those p that are Justified and Not-entailed w.r.t. $\texttt{TMod}(\mathcal{K}, \mathcal{I}_i)$ need to be checked for the Satisfiable property. This remark goes towards optimising a run of Alg. 8.2 and is not essential to the proof of Lemma 8.17.

4 As a side note, this strategy works naturally in every iteration $i \geqslant 1$, because $\mathcal{I}_i$ will always be the outcome of Alg. 8.1, i.e. the $\subseteq$-smallest element in $\texttt{TMod}(\mathcal{K}, \mathcal{I}_i)$. It also works in the first iteration due to the assumption $\mathcal{I} \in \texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$ (which is naturally satisfied for $\texttt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O}})$).

MATERIALISATION-BASED FOUNDATIONS. To apply Proposition 8.16 and Lemma 8.17 for rational and relevant strength explicitly, the construction of the respective typicality domains needs to be investigated. From the definitions of $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_\prec}$ (Def. 7.21) and $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_\prec}$ (Def. 7.24), we can see that their construction heavily relies on the computation of $\mathtt{Cons}_{\mathbf{s}}(\mathcal{K},\chi)$ and the ABox extension $\mathcal{K}_{\mathbf{s}_\prec}$ ($\mathbf{s} \in \{\mathtt{rat},\mathtt{rel}\}$). While these computations diverge vastly between rational and relevant strength, both rest on the rational chain $\mathtt{chain}(\mathcal{K})$. Constructing $\mathtt{chain}(\mathcal{K})$ was originally claimed to always reside in the complexity of classical reasoning for the underlying DL in Casini and Straccia [CS'10]. This claim is only trivially true for DLs over the full boolean spectrum, such as $\mathcal{ALC}$ (allowing for full negation). Much later, in [PT'17c], we proved that the construction of the rational chain remains polynomial also for an $\mathcal{EL}_\bot$ KB $\mathcal{K}$. This result has also been independently achieved by the original authors, Casini et al. [CMS'18].

Proposition 8.18. *For an $\mathcal{EL}_\bot$ KB $\mathcal{K} = (\mathcal{A},\mathcal{T},\mathcal{D})$, the rational chain* $\mathtt{chain}(\mathcal{K})$ *can be computed in polynomial time and is of size at most $|\mathcal{D}|$.*

Proof. Let $\mathtt{chain}(\mathcal{K}) = \langle \mathcal{D}_0, \cdots, \mathcal{D}_n \rangle$. $\mathtt{chain}(\mathcal{K})$ must clearly be finite, because for $0 \leqslant i < n$, $\mathcal{D}_{i+1} \subset \mathcal{D}_i$ by Definition 4.6 (for well-separated KBs). For the same reason, $n \leqslant |\mathcal{D}|$. Given $\mathcal{D}_i$ ($0 \leqslant i < n$), it takes $|\mathcal{D}_i|$ number of classical entailment checks to determine $\mathcal{D}_{i+1}$. More explicitly, for every $E \sqsubsetsim F \in \mathcal{D}_i$ check entailment of $\widehat{\mathcal{D}_i} \sqcap E \sqsubseteq \bot$ by $(\mathcal{A}, \mathcal{T}^{\mathcal{D}})$. $\mathcal{T}^{\mathcal{D}}$ is clearly linearly big in $|\mathcal{T} \cup \mathcal{D}|$. Thus, the claim follows from classical entailment of subsumption in $\mathcal{EL}_\bot$ being decidable in polynomial time [Baa'03; Bra'04]. □

Finally, the computation of the minimal typicality model, can also be captured generically, on the input of a typicality domain $\Delta^{\mathbb{C},\mathbb{O}}$. Its algorithmic construction, and the complexity result thereof, are easily derived from Corollary 7.23.

Proposition 8.19. *For an $\mathcal{EL}_\bot$ KB $\mathcal{K}$, a finite, consistent context $\mathbb{C},\mathbb{O}$, containing $\mathcal{K}$, and a typicality domain $\Delta^{\mathbb{C},\mathbb{O}}$, s.t.* $\mathtt{TMod}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}}) \neq \emptyset$*, the minimal typicality model* $\mathtt{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})$ *can be computed in polynomial time in $|\Delta^{\mathbb{C},\mathbb{O}}|$.*

Proof. Corollary 7.23 describes how to construct, specifically, the minimal typicality model over the rational domain. We generalise the results of this corollary for arbitrary typicality domains, to separate the construction of the domain (relying on ABox extensions), from the construction of its minimal typicality model (not relying on ABox extensions). In general, for a typicality interpretation over $\Delta^{\mathbb{C},\mathbb{O}}$ to be the $\subseteq$-smallest element in $\mathtt{TMod}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})$, it must contain as little information as possible. Specifically, every domain element must follow Def. 7.6 to make the interpretation standard and satisfy $\mathcal{K}_{\mathtt{strict}}$ as well as the DCIs associated to it (Def. 7.7). This generalisation works for the rational and relevant domain in particular, because the associated DCIs already take $\texttt{Cons}_{\mathsf{s}}()$ and $\mathcal{K}_{\mathsf{s}_\prec}$ into account, guaranteeing $\texttt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{s}_\prec}) \neq \emptyset$.

Formally, let $\mathcal{K}_{\Delta^{\mathbb{C},\mathbb{O}}} = (\mathcal{A} \cup \{\widehat{\mathcal{U}}(a) \mid (a, \mathcal{U}) \in \Delta^{\mathbb{C},\mathbb{O}}\}, \mathcal{T}^{\mathcal{D}})$. It is easy to see that $\mathcal{K}_{\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_\prec}} = \mathcal{K}_{\mathsf{rat}_\prec}$, thus, $\mathcal{K}_{\Delta^{\mathbb{C},\mathbb{O}}}$ can be seen as a way to extract an ABox extension from a given typicality domain. The only consequences of Cor. 7.23 that need to be generalised for $\Delta^{\mathbb{C},\mathbb{O}}$ are the following (remember that $\mathcal{ALC}$ can be equivalently exchanged with $\mathcal{EL}_\bot$ in simple materialisation by Cor. 4.24):

$$(a, \mathcal{U}) \in A^{\mathsf{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})} \textit{ iff } \texttt{Mat}^{\mathcal{EL}_\bot}(\mathcal{K}_{\Delta^{\mathbb{C},\mathbb{O}}}, A(a))$$

$$((a, \mathcal{U}), (D, \emptyset)) \in r^{\mathsf{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})} \textit{ iff } \texttt{Mat}^{\mathcal{EL}_\bot}(\mathcal{K}_{\Delta^{\mathbb{C},\mathbb{O}}}, (\exists r.D)(a))$$

Note that $\mathcal{K}_{\Delta^{\mathbb{C},\mathbb{O}}}$ is always polynomial in $|\mathcal{K}| + |\mathbb{O}|$, because for every individual in $\mathbb{O}$, there is exactly one representative in $\Delta^{\mathbb{C},\mathbb{O}}$.

Now, it is not hard to see that for every domain element $(\chi, \mathcal{U})$, there are $|sig_C(\mathcal{K})|$ many classical entailment checks necessary to determine $(\chi, \mathcal{U}) \in A^{\mathsf{M}(\mathcal{K},\Delta^{\mathbb{C},\mathbb{O}})}$ (for $A \in sig_C(\mathcal{K})$). Likewise, all possible role edges (no more than $|\Delta^{\mathbb{C},\mathbb{O}}|^2 * |sig_R(\mathcal{K})|$) can each be determined with a single classical entailment check in polynomial time ([Bra'04]). □

Finally, the preceding results can be instantiated to capture complexity upper bounds for rational and relevant reasoning. Both are treated separately in Section 8.2.1 and 8.2.2, respectively investigating propositional, sceptical nested, and selective nested coverage.

8.2.1 *Rational Reasoning*

From Definition 7.21, it is not hard to see that $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_\prec}$ contains exactly 1 individual representative for each individual in the context and at most 2 concept representatives for each concept in $\mathbb{C}$. However, to construct the rational domain, it is necessary for many representatives, to determine a subset of the DBox that is consistent with the individual or concept to be represented. For individuals in particular, the full ABox extension (Def. 4.16) $\mathcal{K}_{\mathsf{rat}_\prec}$ must be computed beforehand.

Lemma 8.20. *For an $\mathcal{EL}_\bot$ KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, $\texttt{Cons}_{\mathsf{rat}}(\mathcal{K}, \chi)$ can be determined in polynomial time in the size of $\mathcal{K}$ for $\chi \in \mathfrak{C}(\mathcal{EL}_\bot) \cup \mathsf{N}_\mathsf{I}$. Likewise, for a set of individuals $\mathbb{O}$ (in general), $\mathcal{K}_{\mathsf{rat}_\prec}$ can be computed for a total order $\prec$ over N_I, in polynomial time in $|\mathcal{K}| + |\mathbb{O}|$.*

Proof. The first claim is straightforward. Let $\texttt{chain}(\mathcal{K}) = \langle \mathcal{D}_0, \cdots, \mathcal{D}_m \rangle$ (computed in polynomial time in $|\mathcal{K}|$ by Prop. 8.18). At most m classical entailment checks are required to determine a $\mathcal{D}_i$ that is consistent with χ, s.t. i is as small as possible ($0 \leqslant i \leqslant m$). Specifically, for $\chi = C \in \mathfrak{C}(\mathcal{EL}_\bot)$, check $(\mathcal{A}, \mathcal{T}^{\mathcal{D}}) \not\models \widehat{\mathcal{D}_i} \sqcap C \sqsubseteq \bot$, and for $\chi = a \in \mathsf{N}_\mathsf{I}$, check $(\mathcal{A} \cup \{\widehat{\mathcal{D}_i}(a)\}, \mathcal{T}^{\mathcal{D}}) \not\models \bot(a)$.

For the second claim, note that in the iterative ABox extension algorithm (Def. 4.16), $\mathcal{A}_{i+1}$ contains exactly one more concept assertion than $\mathcal{A}_i$.

For rational strength, $\mathtt{init}^{\mathcal{EL}_\bot}_{\mathsf{rat}_\prec}(\mathcal{K}) = (\mathcal{A}, \mathcal{T}^{\mathcal{D}}, \mathcal{D})$ comes at negligible extra computational cost and shows $\mathcal{A}_0$, $\mathcal{T}_0$ and $\mathcal{D}_0$ to be polynomial in $|\mathcal{K}|$. Furthermore, the iteration treats every individual in $\mathbb{O}$ exactly once, i.e. the final ABox is $\mathcal{A}_n$ with $n = |\mathbb{O}|$. Combining these arguments shows that throughout the iteration, $\mathtt{Cons}_{\mathsf{rat}}()$ is determined always with a polynomial input and used only a polynomial number of times, showing that $\mathcal{K}_{\mathsf{rat}_\prec}$ can be computed in polynomial time. □

Relying on Lemma 8.20, it is relatively easy to show that $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_\prec}$ can be constructed in polynomial time. Consequently, its size remains polynomial in the size of the KB and context.

Proposition 8.21. *For an $\mathcal{EL}_\bot$ KB $\mathcal{K}$, a finite context $\mathbb{C}, \mathbb{O}$ and a total order $\prec$ over $\mathrm{N_I}$, the rational typicality domain $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_\prec}$ can be computed in polynomial time in $|\mathcal{K}|$, $|\mathbb{C}|$ and $|\mathbb{O}|$.*

Proof. From Lemma 8.20, it immediately follows that determining each typical concept representative $(C,\ \mathtt{Cons}_{\mathsf{rat}}(\mathcal{K}, C))$ $(C \in \mathbb{C})$ can be done in polynomial time and must be done $|\mathbb{C}|$ many times. To determine the DBox subsets associated to individual elements, $\mathcal{K}_{\mathsf{rat}_\prec}$ must be computed once, and is then queried once for every $a \in \mathbb{O}$, i.e. extracting $\mathtt{ext}(\mathcal{K}_{\mathsf{rat}_\prec}, a)$. Lemma 8.20 shows that both can be achieved in polynomial time in $|\mathcal{K}| + |\mathbb{O}|$. □

It follows directly from Proposition 8.19 and Proposition 8.21, that the minimal typicality model over the rational typicality domain can be computed in polynomial time in the size of the $\mathcal{EL}_\bot$ knowledge base and the given context. We move on to show explicit complexity upper bounds for propositional as well as sceptical and selective nested rational semantics.

PROPOSITIONAL RATIONAL ENTAILMENT We include the results on propositional rational semantics here, mainly to provide a complete list of complexity upper bounds for the decision procedures supported by *all* investigated typicality model semantics at this point. There are two ways to derive polynomiality of propositional rational semantics from the previous results. For one, it follows from Theorem 7.22 that consequences under propositional rational semantics can be determined through $\mathcal{EL}_\bot$ materialisation. Lemma 8.20 shows that this reduction remains polynomial. On the other hand, we proved that $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_\prec})$ can be constructed in polynomial time, which is also sufficient to ensure the following result, due to the canonicity of $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_\prec})$ w.r.t. $\mathtt{TMod}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathsf{rat}_\prec})$ (cf. Def. 7.13 and Eq. (7.8)).

Theorem 8.22. *Entailment of defeasible subsumption and defeasible instance checks under $(\mathtt{rat}_\prec, \mathtt{prop})$ semantics can be determined in polynomial time for $\mathcal{EL}_\bot$.*

SCEPTICAL NESTED RATIONAL ENTAILMENT The difference between sceptical and selective nested semantics lies only in the input admissible partial function for Algorithm 8.2. The sceptical case is rather easy to prove, because computation of $\mathsf{po_N}$ is a single step, i.e. the guess of an admissible preference option. The main analysis of Alg. 8.2 is already presented in Lemma 8.17.

Theorem 8.23. *Entailment under* $(\mathtt{rat}_{\prec}, \mathtt{nest})$ *semantics can be determined in* CO-NPTIME *for* $\mathcal{EL}_{\bot}$.

Proof. Naturally, we prove non-entailment in NPTIME. Let $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ be the KB, $\mathbb{C}$, $\mathbb{O}$ a finite context, $\prec$ a total order over $\mathsf{N_I}$, and α a defeasible query for which to determine non-entailment by $\mathcal{K}$ under $(\mathtt{rat}_{\prec}, \mathtt{nest})$ semantics. W.l.o.g., we assume the context to contain $\mathcal{K}$ and be consistent and quantification closed, and α to be a query over the context. Appropriate extensions of an arbitrary finite context that do satisfy these requirements are at most polynomial in the size of $\mathcal{K}$ and α. From Proposition 8.19 and Proposition 8.21, it is clear that the minimal typicality model $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_{\prec}})$ can be constructed in polynomial time and is of polynomial size in the input.

Proposition 8.13 allows to determine non-entailment of α, by running Algorithm 8.2 on the input $\mathcal{K}$, $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_{\prec}})$ and $\mathsf{po_N}$. From Lemma 8.17 we know that all lines of Alg. 8.2 are executed at most a polynomial number of times in the size of the present inputs and the *while*-condition can also be checked in polynomial time. Proposition 8.16 (minimal model completion) ensures that Line 5 runs in polynomial time. Hence the interpretation $\mathcal{I}_n$ returned by Algorithm 8.2 can be computed in polynomial time in the size of $\mathcal{K}$, $\mathbb{C}$, $\mathbb{O}$ and α (also including the computation of the minimal typicality model). However, in Line 4, the algorithm relies on a non-deterministic choice, resulting in non-deterministic polynomial time to decide non-entailment of a query under $(\mathtt{rat}_{\prec}, \mathtt{nest})$ semantics. □

SELECTIVE NESTED RATIONAL ENTAILMENT For selective nested semantics, Algorithm 8.2 will be called with $\mathsf{po}_{<_{\mathsf{po}}}$ (cf. Corollary 8.15). As opposed to $\mathsf{po_N}$, there is some computational effort involved to determine the $<_{\mathsf{po}}$-minimal admissible preference option for the current iteration. Truth of the *while*-condition (Line 3 in Alg. 8.2) ensures only the existence of this preference option. The proof of the following result distinguishes itself from Thm. 8.23 mostly by showing how expensive it is to compute the $<_{\mathsf{po}}$-minimal admissible preference option. It turns out that this requires no more effort than checking the *while*-condition

Theorem 8.24. *Entailment under* $(\mathtt{rat}_{\prec}, \mathtt{nest}_{<_{\mathsf{po}}})$ *semantics can be determined in polynomial time for* $\mathcal{EL}_{\bot}$.

Proof. Let $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ be the KB, $\mathbb{C}$, $\mathbb{O}$ a finite context, $\prec$ a total order over $\mathsf{N_I}$, $<_{\mathsf{po}}$ a total order over $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_{\prec}})$ and α a defeasible query for which to determine entailment by $\mathcal{K}$ under $(\mathtt{rat}_{\prec}, \mathtt{nest}_{<_{\mathsf{po}}})$ semantics (assuming the same w.l.o.g. assumptions as in Thm. 8.23). As for Theorem 8.23, the minimal typicality model $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_\prec})$ is computed in polynomial time and is of polynomial size in the input. Likewise, a run of Algorithm 8.2 on the input $\mathcal{K}$, $\mathtt{M}(\mathcal{K}, \Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_\prec})$ and $\mathtt{po}_{<_{\mathsf{po}}}$ executes every line in Alg. 8.2 only polynomially many times.

As opposed to sceptical semantics, where Line 4 requires a simple non-deterministic guess, $\mathtt{po}_{<_{\mathsf{po}}}$ actually needs to determine the $<_{\mathsf{po}}$-minimal admissible preference option in every iteration. $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_\prec})$ is enumerated w.r.t. $<_{\mathsf{po}}$, starting at the $<_{\mathsf{po}}$-minimal preference option. When considering $<_{\mathsf{po}}$ as an original input, this enumeration requires only linearly many iterations, in which to check admissibility for $\mathtt{TMod}(\mathcal{K}, \mathcal{I}_i)$. This check is performed the same way as suggested by Claim 1 of Lemma 8.17 for the *while*-condition, i.e. it requires linearly many polynomial time executions (of Alg. 8.1). If the *while*-condition was true, then this enumeration must eventually find the $<_{\mathsf{po}}$-minimal admissible preference option for $\mathtt{TMod}(\mathcal{K}, \mathcal{I}_i)$.

Finally, we have shown that also Line 4, using $\mathtt{po}_{<_{\mathsf{po}}}$ can be computed in polynomial time in $|\mathcal{K}|$ and $|\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_\prec}|$, resulting in an overall polynomial complexity to determine entailment under $(\mathtt{rat}_\prec, \mathtt{nest}_{<_{\mathsf{po}}})$ semantics. □

While the enumeration of preference options in the preceding proof is sufficient to show the result, Alg. 8.2 can run a lot more efficient on the deterministic input $\mathtt{po}_{<_{\mathsf{po}}}$. Clearly, not all options in $\mathtt{PO}(\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rat}_\prec})$ must be enumerated, but only those that are Justified and Not-entailed for $\mathtt{TMod}(\mathcal{K}, \mathcal{I}_i)$. These can be determined from $\mathcal{I}_i$. Also, this enumeration can directly be used to determine the *while*-condition. Finally, once the *while*-condition is verified through this enumeration (as for Claim 1 of Lemma 8.17), it already produced the interpretation $\mathtt{mmc}(\mathcal{K}, \mathcal{I}_i[p]) = \mathcal{I}_{i+1}$, for p, the $<_{\mathsf{po}}$-minimal preference option for $\mathtt{TMod}(\mathcal{K}, \mathcal{I}_i)$. There is practically no need to determine the $<_{\mathsf{po}}$-minimal admissible preference option in Line 4 again.

8.2.2 *Relevant Reasoning*

We proceed through the instantiation of the generic results in Section 8.2 for relevant strength, analogous to the instantiation for rational strength. An analysis of the complexity to compute $\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, \chi)$ and $\mathcal{K}_{\mathtt{rel}_\prec}$, as the foundation for the relevant typicality domain, is followed by an investigation for the size and computability of the relevant domain itself. Finally, we show specific results for the three coverages $\mathtt{prop}$, $\mathtt{nest}$ and $\mathtt{nest}_{<_{\mathsf{po}}}$.

Looking at Definition 6.3, the obvious algorithm to determine $\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, \chi)$, would be to compute all justifications and remove their rank-minimal part, from the original DBox. However, seeing as in $\mathcal{EL}$ there can already be exponentially many justifications for a query w.r.t. a KB ([BPS'07]), we present a much better result, which adopts the problem of finding *some* justification and employs it linearly many times to determine which DCIs are to be disregarded.

Lemma 8.25. *For an $\mathcal{EL}_\bot$ KB $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$, $\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, \chi)$ and $\mathcal{K}_{\mathtt{rel}_\prec}$ can be determined in Δ^{P}_2.*

Proof. The following strategy is effective for computing $\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K},\chi)$. For every DCI $E \mathrel{\tilde{\sqsubseteq}} F \in \mathcal{D}_i$ for $\mathtt{chain}(\mathcal{K}) = \langle \mathcal{D}_0, \cdots, \mathcal{D}_n \rangle$ $(1 \leqslant i \leqslant n)$, check if there is some χ-justification $\mathcal{J} \subseteq \mathcal{D}$, s.t. $E \mathrel{\tilde{\sqsubseteq}} F \in \mathcal{J}$ and $G \mathrel{\tilde{\sqsubseteq}} H \notin \mathcal{J}$ for all $G \mathrel{\tilde{\sqsubseteq}} H \in \mathcal{D}_j$ with $j < i$. Theorem 7 and Corollary 8 (MINA-RELEVANCE) in [PS'17] show that such checks are NP-complete. Thus, our algorithm relies on polynomially many calls to an NP oracle to find (or not find) an appropriate justification for each $E \mathrel{\tilde{\sqsubseteq}} F$. To clarify, if a justification $\mathcal{J}$ (containing $E \mathrel{\tilde{\sqsubseteq}} F$) can be found such that it does not contain any DCIs of a strictly lower antecedent-concept-rank than $r_{\mathcal{K}}(E)$, then clearly $E \mathrel{\tilde{\sqsubseteq}} F$ belongs to some $\mathcal{J}_x^{min}$ (Def. 6.3) and cannot belong to $\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K},\chi)$. Hence, $\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K},\chi)$ can be computed in Δ_2^{P}.

The initialisation of the relevant ABox extension is polynomial by Lemma 8.20. The only other place where the relevant ABox extension algorithm differs from the rational variant, is the application of $\mathtt{Cons}_{\mathtt{rel}}()$, which is shown above to be computed in Δ_2^{P} and is called once for every $a \in \mathbb{O}$. □

Using Lemma 8.25, we can determine the size and required computational effort to create the relevant typicality domain. As usual, $\mathtt{ext}(\mathcal{K}_{\mathtt{rel}_\prec}, a)$ can be read off of $\mathcal{K}_{\mathtt{rel}_\prec}$ for every $a \in \mathbb{O}$, after computing it once. An increase in both size and complexity was to be expected, because the granularity for treating consistent subsets of the DBox (w.r.t. every representative) requires representation of (up to) the full subset lattice of the DBox.

Proposition 8.26. *For an $\mathcal{EL}_\bot$ KB $\mathcal{K}$, a finite context $\mathbb{C}, \mathbb{O}$ and a total order $\prec$ over $\mathrm{N_I}$, the relevant typicality domain $\Delta_{\mathtt{rel}_\prec}^{\mathbb{C},\mathbb{O},\mathcal{K}}$*

- *is of exponential size, and*
- *can be computed in exponential time*

in $|\mathcal{K}|$, $|\mathbb{C}|$ and $|\mathbb{O}|$.

Proof. By Definition 7.24, there is exactly one domain element for every $a \in \mathbb{O}$ and $2^{|\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K},C)|}$ domain elements for every $C \in \mathbb{C}$, clearly showing exponential size of $\Delta_{\mathtt{rel}_\prec}^{\mathbb{C},\mathbb{O},\mathcal{K}}$.

From Lemma 8.25 it follows that the sets of DCIs that each of the most typical representatives (concept and individual alike) are associated with, can be computed in less than exponential time. Therefore, constructing each of these exponentially many elements requires overall at most exponential time. □

It follows directly from Proposition 8.19 and Proposition 8.21, that the minimal typicality model over the relevant typicality domain can be computed in exponential time in the size of the $\mathcal{EL}_\bot$ knowledge base and the given context.

PROPOSITIONAL RELEVANT ENTAILMENT Even though the complexity result for computing the minimal typicality model over the relevant domain gives an immediate exponential upper bound for propositional relevant semantics, their equivalence to materialisation-based relevant reasoning (Thm. 7.25) allows to prove a better upper bound using Δ_2^{P} computations of $\mathtt{Cons}_{\mathtt{rel}}(\mathcal{K}, \chi)$.

Theorem 8.27. *Entailment under* $(\mathtt{rel}_{\prec}, \mathtt{prop})$ *semantics can be determined in* Δ_2^{P} *for* $\mathcal{EL}_{\bot}$.

Proof. From Theorem 7.25 we know that propositional and materialisation-based relevant semantics must share the same complexity upper bound for deciding entailments. From Lemma 8.25 it follows that the extended knowledge base $\mathcal{K}_{\mathtt{rel}_{\prec}}$ can be computed in Δ_2^{P} and is of polynomial size in the original KB. Thus, from Definition 4.3, 4.11, 4.16, Lemma 8.25 and the fact that classical reasoning in $\mathcal{EL}_{\bot}$ is polynomial, we conclude the Δ_2^{P} upper bound for the present claim. □

SCEPTICAL NESTED RELEVANT ENTAILMENT As in the instantiation for rational strength, the main difference between sceptical and selective nested semantics lies only in the input admissible partial function for Algorithm 8.2. Computation of $\mathtt{po}_{\mathtt{N}}$ is a single non-deterministic step, allowing to draw most conclusions directly from Lemma 8.17.

Theorem 8.28. *Entailment under* $(\mathtt{rel}_{\prec}, \mathtt{nest})$ *semantics can be determined in* CO-NEXPTIME *for* $\mathcal{EL}_{\bot}$.

Proof. Naturally, we prove non-entailment in NEXPTIME. Let $\mathcal{K} = (\mathcal{A}, \mathcal{T}, \mathcal{D})$ be the KB, $\mathbb{C}$, $\mathbb{O}$ a finite context, $\prec$ a total order over $\mathsf{N_I}$, and α a query for which to determine non-entailment by $\mathcal{K}$ under $(\mathtt{rel}_{\prec}, \mathtt{nest})$ semantics. W.l.o.g., we assume the context to be consistent, contain $\mathcal{K}$ and α and be quantification closed.

Proposition 8.13 allows to determine non-entailment of α, by running Algorithm 8.2 on the input $\mathcal{K}$, $\mathtt{M}(\mathcal{K}, \Delta_{\mathtt{rel}_{\prec}}^{\mathbb{C},\mathbb{O},\mathcal{K}})$ and $\mathtt{po}_{\mathtt{N}}$. From Lemma 8.17 we know that all lines of Alg. 8.2 are executed at most a polynomial number of times on the present inputs and the *while*-condition can also be checked in polynomial time in those inputs. As opposed to the rational domain, the relevant domain is of exponential size, resulting in overall exponential time to compute the *while*-condition and up to exponentially many executions of the *while*-loop in the original inputs of this decision problem ($\mathcal{K}$, α, $\mathbb{C}$, $\mathbb{O}$). Likewise, Proposition 8.16 shows that Line 5 (computing $\mathtt{mmc}()$) runs in exponential time. Hence the interpretation $\mathcal{I}_n$ returned by Algorithm 8.2 can be computed in exponential time in the size of $\mathcal{K}$, $\mathbb{C}$, $\mathbb{O}$ and α (also including the computation of the minimal typicality model).

In Line 4, Algorithm 8.2 relies on a non-deterministic choice, resulting in non-deterministic exponential time to decide non-entailment of a query under $(\mathtt{rel}_{\prec}, \mathtt{nest})$ semantics. □

SELECTIVE NESTED RELEVANT ENTAILMENT Recall, that the algorithms to compute entailment under selective and sceptical nested semantics diverge only in Line 4 of Algorithm 8.2. The function $\mathsf{po}_{<_{\mathsf{po}}}$ requires non-trivial computational effort, as opposed to po_{N}. Luckily, the argumentation in the proof of Thm. 8.24 for the rational case, applies analogously to the relevant case.

Theorem 8.29. *Entailment under* $(\mathtt{rel}_{\prec}, \mathtt{nest}_{<_{\mathsf{po}}})$ *semantics can be determined in* ExpTime *for* $\mathcal{EL}_{\bot}$.

Proof. Executions and checks concerning the *while*-loop in Alg. 8.2 are exponential in the input of the present decision problem (i.e. $\mathcal{K}$, $\mathbb{C}$, $\mathbb{O}$, α, $\prec$, $<_{\mathsf{po}}$, analogous to Thm. 8.28).

While $<_{\mathsf{po}}$ is polynomial in the size of the typicality domain and thus clearly exponential on the original input KB, it is still viewed as an original input for deciding selective nested relevant entailment. As such, the analogous argument as for Thm. 8.24 shows that an appropriate enumeration of preference options allows to determine the $<_{\mathsf{po}}$-minimal admissible preference option with linearly many calls to Algorithm 8.1, hence, it remains exponential in the original input. □

As briefly commented on in Section 8.1.2 (Page 129), one might consider deriving $<_{\mathsf{po}}$ from the other inputs, as a more practicable approach. Even if $<_{\mathsf{po}}$ is not considered an original input in future approaches, the argument in the proof of Thm. 8.29 needs only minor adjustment to remain intact. It holds that $|\mathsf{PO}(\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_{\prec}})| \leqslant |\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}}_{\mathtt{rel}_{\prec}}|^2 * |sig_{\mathsf{R}}(\mathcal{K})|$, thus, for an exponentially large typicality domain, $<_{\mathsf{po}}$ remains exponential in the other inputs. Consequently, Algorithm 8.1 will be called an exponential number of times to determine a $<_{\mathsf{po}}$-minimal preference option. However, this does not affect the number of loop-iterations in Alg. 8.2. Specifically, the resulting complexity is the square of an exponential rather than a double-exponential.

8.3 COMPLEXITY LOWER BOUNDS

Section 8.1 and 8.2 go hand in hand, because providing an algorithmic characterisation, together with its termination and correctness is an effective way to prove upper bounds for the computational complexity of a decision problem. Determining good complexity lower bounds requires a different approach. At this time, we are able to present tight lower bounds only for typicality model semantics of rational strength.

Just as for the upper bounds, there is no distinctive behaviour in terms of complexity for the typicality model formalism, between defeasible subsumption and instance checking. The results in this section either immediately apply to both types of queries or translate from one to the other with a simple argument. Additionally, the lower bounds for propositional and selective nested coverage trivially follow from their respective upper bounds (cf.

Sec. 8.2.1) in rational strength. To elaborate, deciding entailment of defeasible subsumption and instance checks under any rational semantics covers the entailment of classical consequences. Specifically, $(\mathcal{A}, \mathcal{T}) \models C \sqsubseteq D$ implies $(\mathcal{A}, \mathcal{T}, \mathcal{D}) \models^{(\mathsf{rat},\mathbf{c})} C \mathrel{\underset{\sim}{\sqsubset}} D$.[5] For $\mathcal{EL}_\bot$ KBs and queries, this provides an immediate polynomial lower bound for each of the discussed coverage types $\mathbf{c} \in \{\mathtt{prop}, \mathtt{nest}, \mathtt{nest}_{<_{\mathsf{po}}}\}$. For those coverages with the matching upper bound (Thm. 8.22 and 8.24) we formally express this result for completeness sake.

Theorem 8.30. *Deciding entailment of defeasible subsumption or instance checks under* propositional *and* selective nested *rational semantics is* P*-hard for* $\mathcal{EL}_\bot$*.*

While this lower bound applies for sceptical semantics as well, we have been able to close the remaining gap between P and CO-NP from below, using the reduction method. That is, taking a known hardness result of another decision problem, showing a suitable reduction to an instance of the decision problem at hand (i.e. defeasible subsumption and instance checking) and proving that one is solvable *iff* the other is solvable as well. In the remainder of this section we present this intricate and technical reduction from satisfiability of a (1-in-3)-positive 3SAT formula, to *non-entailment* of a defeasible subsumption query under sceptical nested rational semantics. The former is a SAT problem that is known to be NP-complete [GJ'79].

The upcoming exposition follows closely our initial investigation in [PT'18]. The main difference is that in [PT'18] we defined nested semantics on so-called *maximal typicality models*. From Section 8.1, it can be seen that these maximal typicality models correspond to the $\subseteq$-smallest elements of Max-TMs in the present formalisation. Hence, this proof for a complexity lower bound follows the same principles as [PT'18] but exhibits a variety of changes, from subtle adjustments in the reduction, to major differences in the proof of correspondence (Prop. 8.32). A reduction to defeasible instance checking is analogous and the required adjustments are described at the end of this section. The discussion in Chapter 9 contains several thoughts on lower bounds for semantics of relevant strength.

8.3.1 *Sceptical Nested Rational Reasoning*

We prove CO-NP-hardness for deciding sceptical nested rational subsumption entailment by a reduction from (1-in-3)-positive 3SAT. As this SAT problem is NP-complete ([GJ'79]), we provide a reduction to non-entailment of a defeasible subsumption query under $(\mathtt{rat}, \mathtt{nest})$-semantics.

(1-IN-3)-POSITIVE 3SAT. A (1-in-3)-positive 3SAT problem is given with a propositional formula φ in conjunctive normal form, with clauses of

5 For a side note, this implication holds for any entailment relation satisfying the KLM postulates (Ref) and (RW).

size 3. The general representation for such a problem with n clauses and k propositional variables $\mathcal{V} = \{x_1, \dots, x_k\}$ is

$$\varphi = \bigwedge_{i=1}^{n} (x_{i_1}, x_{i_2}, x_{i_3}),$$

s.t. $i_1, i_2, i_3 \in \{1, \dots, k\}$ ($i_1 \neq i_2 \neq i_3$) for all $1 \leqslant i \leqslant n$. $(x_{i_1}, x_{i_2}, x_{i_3})$ is called a *clause* in φ. Such an instance φ is satisfied by a truth assignment if in every clause in φ there is *exactly* one propositional variable that is assigned to true (note, all variables occur as a positive literal). More formally, a *truth assignment* is a function $\sigma : \mathcal{V} \to \{\top, \bot\}$. σ is extended to apply to a (1-in-3)-positive 3SAT formula φ s.t. $\widehat{\sigma}(\varphi) = \top$ *iff* for all $i \in \{1, \dots, n\}$ there is *exactly* one x_{i_j} ($j \in \{1, 2, 3\}$) s.t. $\sigma(x_{i_j}) = \top$.

THE REDUCTION FROM (1-IN-3)-POSITIVE 3SAT. Intuitively, the main correspondence of sceptical rational reasoning and satisfiability of (1-in-3)-positive 3SAT formulas lies between finding Max-TMs and finding satisfying assignments. Thus, the idea behind this reduction is the following. For a (1-in-3)-positive 3SAT formula φ, we shall construct a corresponding knowledge base (and context/query), such that a satisfying assignment for φ can be transformed into a full preference chain for the KB (and context) and vice versa. Correctness of this reduction, is then proven by the fact that this full preference chain provides a Max-TM counterexample to the generated query.

Given φ, an instance of (1-in-3)-positive 3SAT, the knowledge base $\mathcal{K}_\varphi$ is constructed over the signature

$$\mathsf{sig}_C(\mathcal{K}_\varphi) = \{A, B, X\} \cup \{C_i \mid 1 \leqslant i \leqslant n\} \tag{8.6}$$

$$\mathsf{sig}_R(\mathcal{K}_\varphi) = \{s, r_1, \dots, r_k\}, \tag{8.7}$$

where k is the number of distinct propositional variables and n the number of clauses occurring in φ. W.l.o.g. we assume a linear order on the clauses in φ, simply to reference them by indices $1 \leqslant i \leqslant n$. For explanation purposes, assume the domain elements that are referenced in the following to occur in the rational domain. This is formally confirmed after the full construction of $\mathcal{K}_\varphi$.

The correspondence between preference chains and truth assignments is provided by linking preference options for roles r_j to the propositional variable x_j. Specifically, $\sigma(x_j) = \top$ translates to a preference option $r_j((A, \mathcal{D}), (B, \mathcal{D}))$ being used in a preference chain. For a satisfying assignment σ that is translated to a preference chain in this way, it will turn out that the set of models $\mathcal{M}$ at the end of this chain is maximal w.r.t. $<_t$. Neither any of the r_i that have not yet been used in this chain, nor the extra role name s will provide admissible preference options for $\mathcal{M}$. The query will be constructed in such a way, that unsatisfiability of a specific preference option for s in $\mathcal{M}$ will ensure that $\mathcal{M}$ is a counterexample to the query's entailment.

The DBox in $\mathcal{K}_\varphi$ shall be rather small

$$\mathcal{D} = \{B \precsim X, \tag{8.8}$$
$$A \precsim \exists s.B, \tag{8.9}$$
$$A \precsim \exists r_1.B \sqcap \ldots \sqcap \exists r_k.B\}. \tag{8.10}$$

The remainder of the reduction can be achieved using the TBox and the query. The strict part of $\mathcal{K}_\varphi$ will not directly be in conflict with any defeasible information, i.e. A and B will not be exceptional w.r.t. $\mathcal{D}$ and $\mathcal{T}$ (formally shown after presenting $\mathcal{K}_\varphi$). Consequently, $\mathtt{chain}(\mathcal{K}) = \langle \mathcal{D}, \emptyset \rangle$, simplifies the structure of the rational domain greatly. The information contained in the resulting minimal typicality model $\mathtt{M}(\mathcal{K}_\varphi, \Delta_{\mathsf{rat}}^{\mathbb{C},\mathbb{O},\mathcal{K}_\varphi})$ is mostly influenced by this DBox. (8.9) is required "in the end", to ensure that the query is not entailed *iff* φ is satisfiable. (8.10) provides the property Justified (cf. Def. 7.27) for admissibility of all preference options that will be selected on the full preference chain. Most importantly, (8.8), ensures that the typical B representative will satisfy X in all typicality models of $\mathcal{K}_\varphi$ over the rational domain. We rely heavily on this containment in X to distinguish satisfied preference options, connecting the typical A representative with the typical B representative (i.e. $r_i((A, \mathcal{D}), (B, \mathcal{D}))$), from their atypical versions introduced by (8.10).

The TBox is mostly made up of two types of GCIs, each capturing information about the clauses of φ.

$$\mathcal{T}^i_{const} = \{\exists r_{i_1}.X \sqcap \exists r_{i_2}.X \sqsubseteq \bot, \tag{8.11}$$
$$\exists r_{i_2}.X \sqcap \exists r_{i_3}.X \sqsubseteq \bot, \tag{8.12}$$
$$\exists r_{i_1}.X \sqcap \exists r_{i_3}.X \sqsubseteq \bot\} \tag{8.13}$$

and

$$\mathcal{T}^i_{clause} = \{\exists r_{i_1}.X \sqsubseteq C_i, \tag{8.14}$$
$$\exists r_{i_2}.X \sqsubseteq C_i, \tag{8.15}$$
$$\exists r_{i_3}.X \sqsubseteq C_i\,\} \tag{8.16}$$

where the i-th clause in φ is $(x_{i_1}, x_{i_2}, x_{i_3})$ $(1 \leqslant i \leqslant n)$. $\mathcal{T}^i_{const}$ describes the disjointness constraints that will allow at most one of the corresponding preference options (e.g. $r_{i_1}((A, \mathcal{D}), (B, \mathcal{D}))$) to be used in any full preference chain. For instance the clause (x_1, x_3, x_5) prohibits any pair of these three variables to be set to $\top$ at the same time. Likewise, the constraints

$\exists r_1.X \sqcap \exists r_3.X \sqsubseteq \bot$,
$\exists r_3.X \sqcap \exists r_5.X \sqsubseteq \bot$, and
$\exists r_1.X \sqcap \exists r_5.X \sqsubseteq \bot$

prohibit two preference options, connecting the typical A representative with the typical B representative, using r_1, r_3, r_5 (Justified by (8.10)), to be satisfied by any typicality model in $\mathtt{TMod}(\mathcal{K}_\varphi, \Delta_{\mathsf{rat}}^{\mathbb{C},\mathbb{O},\mathcal{K}_\varphi})$. This conflict arises, because the typical B representative satisfies $B \precsim X$ (8.8).

The TBox $\mathcal{T}^i_{\mathtt{clause}}$ creates a kind of marker at the domain element $(A, \mathcal{D})$. If a typicality model $\mathcal{I}$ satisfies a preference option $r_{i_j}((A, \mathcal{D}),(B, \mathcal{D}))$ for the i-th clause using variables x_{i_j} $(j \in \{1,2,3\})$, then from $\mathcal{I} \models \mathcal{T}^i_{\mathtt{clause}}$ it follows that $(A, \mathcal{D}) \in C_i^{\mathcal{I}}$. This C_i marker represents that at least one preference option, corresponding to a variable of the i-th clause, is satisfied for a Max-TM. Together $\mathcal{T}^i_{\mathtt{const}}$ and $\mathcal{T}^i_{\mathtt{clause}}$ ensure that $C_1, \ldots, C_n$ are satisfied by $(A, \mathcal{D})$ when every clause in φ has *exactly one* variable set to $\top$.

The role edge $((A, \mathcal{D}),(B, \mathcal{D})) \in s^{\mathsf{M}(\mathcal{K}_\varphi, \Delta^{C,O,\mathcal{K}_\varphi}_{\mathsf{rat}})}$ (Justified by (8.9)) is used to invert the query entailment. To be explicit, the query is

$$A \sqsubsetsim \exists s.X \tag{8.17}$$

and the inversion is established by the following disjointness constraint.

$$\mathcal{T}_Q = \{\exists s.X \sqcap C_1 \sqcap \cdots \sqcap C_n \sqsubseteq \bot\} \tag{8.18}$$

Intuitively, $\mathcal{T}_Q$ ensures that the preference option $s((A, \mathcal{D}),(B, \mathcal{D}))$ can only be satisfied if *not* all clauses in φ are satisfied. Combining the above, results in the KB

$$\mathcal{K}_\varphi = (\bigcup_{i=1}^{n} (\mathcal{T}^i_{\mathtt{const}} \cup \mathcal{T}^i_{\mathtt{clause}}) \cup \mathcal{T}_Q, \mathcal{D}). \tag{8.19}$$

Recall that the existence of any Max-TM as a counterexample to the query is required to show non-entailment. It is therefore enough to show that a satisfying assignment for φ can be translated to a full preference chain providing such a counterexample, ignoring other full preference chains that might satisfy $s((A, \mathcal{D}),(B, \mathcal{D}))$.

COMPLEXITY AND CORRECTNESS OF THE REDUCTION. We proceed to show that this reduction is linear and that satisfiability of φ corresponds to non-entailment of the constructed query under nested rational semantics. Note that every GCI and DCI in $\mathcal{K}_\varphi$ is of constant size, with the exception of (8.10) and (8.18), which are linear in φ. Therefore, it suffices to analyse the number of axioms in $\mathcal{K}_\varphi$, rather than its size in terms of symbols[6]. The latter grows only linearly in the size of $\mathcal{K}_\varphi$, because only a constant number of axioms grows linearly in the input.

Proposition 8.31. *$\mathcal{K}_\varphi$ can be constructed in linear time in $|\varphi|$.*

Proof. The size of the input $\varphi = \bigwedge_{i=1}^{n}(x_{i_1}, x_{i_2}, x_{i_3})$, is $3 * n$. The number of constraints in $\mathcal{D}$, $\mathcal{T}_Q$ is constant and their size (as well as the size of the query) is at most linear in the size of the input.

An algorithm constructing $\mathcal{K}_\varphi$ needs to iterate over every clause in φ exactly once, generating $\mathcal{T}^i_{\mathtt{const}}$ and $\mathcal{T}^i_{\mathtt{clause}}$. The TBox parts $\mathcal{T}^i_{\mathtt{const}}$ and $\mathcal{T}^i_{\mathtt{clause}}$ have the constant size $|\mathcal{T}^i_{\mathtt{const}}| = |\mathcal{T}^i_{\mathtt{clause}}| = 3$ for every i. Therefore, the total size of $\mathcal{T}$ amounts to $\|\mathcal{T}\| = 6n + 1$ which can clearly be generated in linear time in the size of φ. □

6 The size of a concept is typically considered to be the number of symbols required to represent that concept, including ⊓, ∃, etc.

We continue to prove the main claim required for the desired hardness result. As typicality model semantics are defined under a given context, an appropriate, finite, consistent, quantification closed context, containing $\mathcal{K}_\varphi$, must be involved in this result.

Proposition 8.32. *A (1-in-3)-positive* 3SAT *formula* φ *is satisfiable* iff $\mathcal{K}_\varphi \not\models^{(rat,nest)} A \sqsubsetapprox \exists s.X$ *w.r.t.* $\mathbb{C} = \{A, B, X\}$, $\mathbb{O} = \emptyset$.

Proof. This proof consists of three parts. Before proving the claim in both directions separately, we provide the foundation for the typicality model semantics with the rational domain induced by $\mathcal{K}_\varphi$ and the context. Showing both directions is then relatively straightforward, by constructing appropriate preference chains or truth assignments, based on the respective premise. Let φ consist of n clauses, using k distinct variables.

For $\mathcal{K}_\varphi = (\mathcal{T}, \mathcal{D})$ we begin by showing $\mathtt{chain}(\mathcal{K}_\varphi) = \langle \mathcal{D}, \emptyset \rangle$ as a prerequisite for $\Delta^{\mathbb{C},\mathbb{O},\mathcal{K}_\varphi}_{\mathsf{rat}}$. Let $\mathcal{I} = (\Delta^\mathcal{I}, \cdot^\mathcal{I})$ be an interpretation with $\Delta^\mathcal{I} = \{d, e\}$, s.t. $d \in (A \sqcap B \sqcap X)^\mathcal{I}$, $e \in B^\mathcal{I}$ and $(d, e) \in r_i^\mathcal{I}$, $(d, e) \in s^\mathcal{I}$ for $1 \leqslant i \leqslant k$. Because no r_i- or s-successor satisfies X, it trivially holds that $\mathcal{I} \models \mathcal{T}$. Furthermore, it not hard to verify that $d \in (A \sqcap B \sqcap X \sqcap \overline{\mathcal{D}})^\mathcal{I}$, showing that the left-hand sides of all DCIs are not exceptional w.r.t. the entire DBox (and $\mathcal{T}$). Consequently, $\mathtt{chain}(\mathcal{K}_\varphi) = \langle \mathcal{D}, \emptyset \rangle$ by Def. 4.6.

The rational domain only contains concept representatives, because $\mathbb{O}$ is empty. As shown above, all concepts in $\mathbb{C}$ are consistent with $\mathcal{D}$. Thus, the rational domain for $\mathcal{K}_\varphi$ and $\mathbb{C}$ (omitting $\mathbb{O}$ hereinafter) is

$$\Delta^{\mathbb{C},\mathcal{K}_\varphi}_{\mathsf{rat}} = \{(E, \mathcal{U}) \mid E \in \mathbb{C}, \mathcal{U} \in \mathtt{chain}(\mathcal{K})\}.$$

For a more readable handle on the elements in $\Delta^{\mathbb{C},\mathcal{K}_\varphi}_{\mathsf{rat}}$, we denote the *typical* concept representatives (for $E \in \mathbb{C}$) with $t_E = (E, \mathcal{D})$ and the *non-typical* E representatives with $u_E = (E, \emptyset)$.

Clearly $(A, \mathcal{D})$ (i.e. t_A) is the most typical representative of A. Therefore, the right-hand side of this proposition, $\mathcal{K}_\varphi \not\models^{(rat,nest)} A \sqsubsetapprox \exists s.X$, is equivalent to

$$\exists \mathcal{M} \in \mathtt{TMax}(\mathcal{K}_\varphi, \Delta^{\mathbb{C},\mathcal{K}_\varphi}_{\mathsf{rat}}). \exists \mathcal{J} \in \mathcal{M}. t_A \notin (\exists s.X)^\mathcal{J}. \tag{8.20}$$

To prove this proposition, we show both directions separately:

(i) If φ is satisfiable then (8.20).

(ii) If (8.20) then there is a satisfying assignment for φ.

Towards (i), assume σ is a satisfying assignment for φ, i.e. $\widehat{\sigma}(\varphi) = \top$. Let $\mathcal{V}_\top = \{x \in \mathcal{V} \mid \sigma(x) = \top\}$ and for $|\mathcal{V}_\top| = m$, let $x_1, \cdots, x_m$ be any arbitrary (but from here on *fixed*) enumeration of all variables in $\mathcal{V}_\top$. From the set of variables that are assigned $\top$ under σ, we derive a preference chain, starting at $\mathtt{TMod}(\mathcal{K}_\varphi, \Delta^{\mathbb{C},\mathcal{K}_\varphi}_{\mathsf{rat}})$ and iteratively refining (Def. 7.29) the sets of models, with the preference options corresponding to those variables.

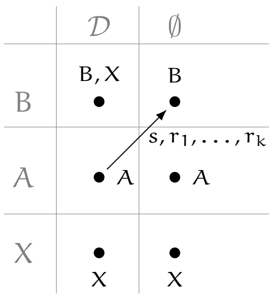

(a) Minimal typicality model

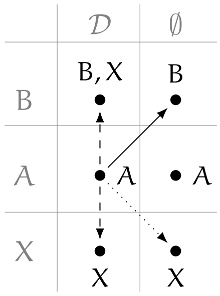

(b) Preference options

Figure 8.1: **(a)** is the labelled graph visualisation of $\mathtt{M}(\mathcal{K}_\varphi, \Delta_{\mathsf{rat}}^{\mathbb{C},\mathcal{K}_\varphi})$. **(b)** is an exemplification of inferred role edges from refining sets of models extending (a) (role labels are omitted). The dashed edges are explicitly satisfied preference options and the dotted edge illustrates inferred edges from refining with a preference option over (t_A, t_B).

Specifically, let $p_i^B = r_i(t_A, t_B)$ as well as $p_i^X = r_i(t_A, t_X)$ (for $1 \leqslant i \leqslant m$) and define

$$\mathcal{M}_1^B = \mathtt{TMod}(\mathcal{K}_\varphi, \Delta_{\mathsf{rat}}^{\mathbb{C},\mathcal{K}_\varphi})|_{p_1^B} \tag{8.21}$$

$$\mathcal{M}_i^B = \mathcal{M}_{i-1}^B|_{p_i^B} \qquad \text{(for } 1 < i \leqslant m\text{)} \tag{8.22}$$

$$\mathcal{M}_1^X = \mathcal{M}_m^B|_{p_1^X} \tag{8.23}$$

$$\mathcal{M}_i^X = \mathcal{M}_{i-1}^X|_{p_i^X} \qquad \text{(for } 1 < i \leqslant m\text{).} \tag{8.24}$$

p_i^B and p_i^X are the preference options corresponding to those variables in φ that are set to $\top$ under σ (dashed edges in Fig. 8.1 (b)). The sets of models $\mathcal{M}_i^B$ and $\mathcal{M}_i^X$ are the (intermediary) results when restricting the set of all models, say $\mathcal{M}_0 = \mathtt{TMod}(\mathcal{K}_\varphi, \Delta_{\mathsf{rat}}^{\mathbb{C},\mathcal{K}_\varphi})$, successively, by the chain of preference options $p_1^B, \cdots, p_m^B, p_1^X, \cdots, p_m^X$. We show admissibility (Def. 7.27) for each preference option to its appropriate set of models (e.g. p_2^B is admissible w.r.t. $\mathcal{M}_1^B$), thus, proving the following chain of typicality preferences:

$$\mathcal{M}_0 <_t \mathcal{M}_1^B <_t \cdots <_t \mathcal{M}_m^B <_t \mathcal{M}_1^X <_t \cdots <_t \mathcal{M}_m^X \tag{8.25}$$

ADMISSIBILITY OF p_i^B FOR $\mathcal{M}_{i-1}^B$. From (8.10), it is easy to see that $\mathcal{M}_0 \models A \mathrel{\tilde{\sqsubseteq}} \exists r_i.B$, or more specifically, $\forall \mathcal{J} \in \mathcal{M}_0.(t_A, u_B) \in r_i^{\mathcal{J}}$, holds for $1 \leqslant i \leqslant m$. Hence, all preference options p_i^B are Justified for $\mathcal{M}_0$ and all its subsets, in particular $\mathcal{M}_i^B$ for $1 \leqslant i \leqslant m$.

For the property Satisfiable, consider $\mathcal{J} = (\Delta_{\mathsf{rat}}^{\mathbb{C},\mathcal{K}_\varphi}, \cdot^{\mathcal{J}})$ with

- $A^{\mathcal{J}} = \{t_A, u_A\}$, $B^{\mathcal{J}} = \{t_B, u_B\}$, $X^{\mathcal{J}} = \{t_B, t_X, u_X\}$,
- $C_j^{\mathcal{J}} = \{t_A\}$ for $1 \leqslant j \leqslant n$,
- $s^{\mathcal{J}} = \{(t_A, u_B)\}$, and
- $r^{\mathcal{J}} = \begin{cases} \{(t_A, e) \mid e \in \{t_B, u_B, t_X, u_X\}\} & \text{if } r \in \{r_1, \cdots, r_m\} \\ \{(t_A, u_B)\} & \text{otherwise.} \end{cases}$

An illustration of role edges with $r \in \{r_1, \cdots, r_m\}$, are the edges in Fig. 8.1 (b). Clearly, $\mathcal{J} \models p_i^B$ for $1 \leqslant i \leqslant m$, by construction. We show $\mathcal{J} \in \mathcal{M}_m^B$, to derive that each p_i^B is Satisfiable for $\mathcal{M}_i^B$ for $1 \leqslant i \leqslant m$. $\mathcal{J} \models \mathcal{T}_Q$, because the only s-successor in $\mathcal{J}$ does not satisfy X. $\mathcal{J} \models \mathcal{T}^j_{clause}$ holds, because the only element with r-successors is t_A, and t_A satisfies all C_j $(1 \leqslant j \leqslant n)$. Assume w.l.o.g. that there is a GCI $\exists r_{j_1}.X \sqcap \exists r_{j_2}.X \sqsubseteq \bot$ in some $\mathcal{T}^j_{const}$ $(1 \leqslant j \leqslant n)$, s.t. $(\exists r_{j_1}.X \sqcap \exists r_{j_2}.X)^{\mathcal{J}} \neq \emptyset$. This implies $r_{j_1}, r_{j_2} \in \{r_1, \cdots, r_m\}$ by the construction of $\mathcal{J}$. However, from the construction of $\mathcal{T}^j_{const}$, it follows that the j-th clause in φ contains both x_{j_1} and x_{j_2}, contradicting $x_{j_1}, x_{j_2} \in \mathcal{V}_\top$ for the satisfying assignment σ, and thus contradicting $r_{j_1}, r_{j_2} \in \{r_1, \cdots, r_m\}$. As it is not hard to show that $\mathcal{J}$ is standard and satisfies the sets of DCIs associated to each domain element, it follows that $\mathcal{J} \in \mathtt{TMod}(\mathcal{K}_\varphi, \Delta_{\mathsf{rat}}^{\mathbb{C},\mathcal{K}_\varphi})$. Furthermore, it can be seen that $\mathcal{J} \in \mathcal{M}_i^B$ (as well as $\mathcal{J} \in \mathcal{M}_i^X$) for all $1 \leqslant i \leqslant m$, because $\mathcal{J}$ satisfies all preference options used for the refinements constructing $\mathcal{M}_m^X$ ((8.21)–(8.24)). Consequently, $\mathcal{J}$ is a witness for all p_i^B being Satisfiable w.r.t. $\mathcal{M}_{i-1}^B$ for all $1 \leqslant i \leqslant m$.

For Not-entailed, let $\mathcal{J}_l$ $(1 \leqslant l \leqslant m)$ coincide with $\mathcal{J}$ on everything but $r_l^{\mathcal{J}_l} = \{(t_A, u_B)\}$. Clearly, $\mathcal{J}_l \not\models p_l^B$ and it is not hard to verify that $\mathcal{J}_l \in \mathtt{TMod}(\mathcal{K}_\varphi, \Delta_{\mathsf{rat}}^{\mathbb{C},\mathcal{K}_\varphi})$. Therefore, for $1 \leqslant i < l$, it holds that $\mathcal{J}_l \in \mathcal{M}_i^B$, showing that $\mathcal{J}_l$ witnesses p_l^B being Not-entailed in $\mathcal{M}_{l-1}^B$. Likewise, $\mathcal{J}_1$ witnesses p_1^B being Not-entailed w.r.t. $\mathcal{M}_0$.

In conclusion, every p_i^B is admissible for $\mathcal{M}_{i-1}^B$ $(1 < i \leqslant m)$ and p_1^B is admissible for $\mathcal{M}_0$. Due to the construction of all $\mathcal{M}_i^B$, this directly proves all relations in the chain

$$\mathcal{M}_0 <_t \mathcal{M}_1^B <_t \cdots <_t \mathcal{M}_m^B$$

ADMISSIBILITY OF p_i^X FOR $\mathcal{M}_{i-1}^X$. Showing admissibility of every p_i^X for $\mathcal{M}_{i-1}^X$ $(1 < i \leqslant m)$ and p_1^X for $\mathcal{M}_m^B$ is slightly easier than the previous case, because, intuitively speaking, t_X and u_X coincide. The second half of the chain in (8.25) is mostly a technicality, because $\mathbb{C}$ must contain $\mathcal{K}_\varphi$, i.e. the rational domain must represent the concept $X \in \mathsf{Qc}(\mathcal{K}_\varphi)$. Every model in $\mathcal{M}_m^B$ satisfies all p_i^B $(1 \leqslant i \leqslant m)$ and because $t_B \in X^{\mathcal{I}}$ (for all $\mathcal{I} \in \mathcal{M}_m^B$) and models in $\mathcal{M}_m^B$ are standard, they all must satisfy $(t_A, u_X) \in r_i^{\mathcal{I}}$ (Property 3 of Def. 7.6). This already shows that all p_i^X are Justified for $\mathcal{M}_m^B$ and all $\mathcal{M}_i^X$.

The typicality model $\mathcal{J}$, constructed in the previous case, is also a witness for every p_i^X being Satisfiable w.r.t. $\mathcal{M}_{i-1}^X$, because $\mathcal{J} \in \mathcal{M}_m^X$ and $\mathcal{M}_m^X \subseteq \mathcal{M}_i^X$ $(1 < i \leqslant m)$. Likewise, p_1^X is Satisfiable for $\mathcal{M}_m^B$.

This time, for non-entailment, let $\mathcal{J}_l'$ coincide with $\mathcal{J}$ on everything but $r_l^{\mathcal{J}_l'} = \{(t_A, u_B), (t_A, t_B), (t_A, u_X)\}$. It is not hard to see that $\mathcal{J}_l' \in \mathcal{M}_m^B$ as well as $\mathcal{J}_l' \in \mathcal{M}_i^X$ for $1 \leqslant i < l$ and $\mathcal{J}_l' \not\models p_l^X$. As in the previous case, this shows that p_i^X is Not-entailed w.r.t.

$\mathcal{M}^X_{i-1}$ for every $1 < i \leqslant m$ (likewise for p^X_1 and $\mathcal{M}^B_m$), thus, proving admissibility of p^X_i and overall truth of (8.25).

It remains to show that there does not exist an admissible preference option for $\mathcal{M}^X_m$, in order to conclude $<_t$-maximality of $\mathcal{M}^X_m$, i.e. that the chain in (8.25) is *full*. It was shown before, that the constructed model $\mathcal{J}$ belongs to $\mathcal{M}^X_m$ and it is not hard to verify that $\mathcal{J}$ is the smallest model in $\mathtt{TMod}(\mathcal{K}_\varphi, \Delta^{C,\mathcal{K}_\varphi}_{\mathsf{rat}})$ that satisfies all p^B_i and p^X_i, i.e. $\mathcal{J} = \bigcap_{\mathcal{I} \in \mathcal{M}^X_m} \mathcal{I}$. The remaining Justified and Not-entailed preference options w.r.t. $\mathcal{M}^X_m$ can be read off of $\mathcal{J}$ as suggested in the beginning of Section 8.1.2 (Page 125). They are comprised of those $r_i(t_A, t_B)$ ($1 \leqslant i \leqslant k$) for which $\sigma(x_i) = \bot$, as well as $s(t_A, t_B)$. Because $t_A \in C^{\mathcal{J}}_j$ for all clauses in φ, i.e. $1 \leqslant j \leqslant n$, no model in $\mathcal{M}^X_m$ can satisfy $s(t_A, t_B)$, i.e. contradicting the Satisfiable property for admissibility of $s(t_A, t_B)$. Additionally, every x_i for which $\sigma(x_i) = \bot$ appears in the j-th clause in φ together with some x_l for which $\sigma(x_l) = \top$. For this x_l, the preference option $r_l(t_A, t_B)$ is satisfied in all models in $\mathcal{M}^X_m$ and no model can satisfy $r_l(t_A, t_B)$ and $r_i(t_A, t_B)$ at the same time, because of $\exists r_i.X \sqcap \exists r_l.X \sqsubseteq \bot \in \mathcal{T}^j_{const}$. Consequently, $\mathcal{M}^X_m$ has no admissible preference options and is $<_t$-maximal by definition. From the construction of $\mathcal{J}$ and $\mathcal{J} \in \mathcal{M}^X_m$, we can easily see that $\mathcal{J} \not\models A \mathrel{\underset{\sim}{\sqsubset}} \exists s.X$, concluding the proof of (i).

For showing (ii), assume $\mathcal{M} \in \mathtt{TMax}(\mathcal{K}_\varphi, \Delta^{C,\mathcal{K}_\varphi}_{\mathsf{rat}})$ such that $\mathcal{M} \not\models A \mathrel{\underset{\sim}{\sqsubset}} \exists s.X$. Let the assignment $\sigma_\mathcal{M}$ be as follows:

$$\sigma_\mathcal{M}(x_i) = \top \textit{ iff } \mathcal{M} \models A \mathrel{\underset{\sim}{\sqsubset}} \exists r_i.X \tag{8.26}$$

for every $1 \leqslant i \leqslant k$. Roughly speaking, we prove satisfaction of φ with $\sigma_\mathcal{M}$, by unravelling the reason for $s(t_A, t_B)$ not being satisfied by all models in $\mathcal{M}$. First of all, the only Justified preference option for the role s w.r.t. $\mathtt{TMod}(\mathcal{K}_\varphi, \Delta^{C,\mathcal{K}_\varphi}_{\mathsf{rat}})$ is $s(t_A, t_B)$, due to $A \mathrel{\underset{\sim}{\sqsubset}} \exists s.B \in \mathcal{D}$. From the TBox and DBox, it can be readily seen that no full preference chain starting from $\mathtt{TMod}(\mathcal{K}_\varphi, \Delta^{C,\mathcal{K}_\varphi}_{\mathsf{rat}})$, and reaching $\mathcal{M}$, can ever introduce another Justified preference option using s, thus, we can focus our attention on $s(t_A, t_B)$. $\mathcal{M} \not\models A \mathrel{\underset{\sim}{\sqsubset}} \exists s.X$ implies that $s(t_A, t_B)$ is Not-entailed w.r.t. $\mathcal{M}$. However, maximality of $\mathcal{M}$ w.r.t. $<_t$ implies $s(t_A, t_B)$ not to be admissible, which leaves $s(t_A, t_B)$ only to be not Satisfiable w.r.t. $\mathcal{M}$. We use this observation to derive that t_A is an element of C_j in all models in $\mathcal{M}$ for $1 \leqslant j \leqslant n$ as follows.

From $\mathtt{TMod}(\mathcal{K}_\varphi, \Delta^{C,\mathcal{K}_\varphi}_{\mathsf{rat}}) \leqslant^*_t \mathcal{M}$ and Prop. 8.3 we know a rather technical feature of $\mathcal{M}$. For one, $\mathcal{J} = \bigcap_{\mathcal{I} \in \mathcal{M}} \mathcal{I}$ is the $\subseteq$-smallest member of $\mathcal{M}$ and $\mathcal{M} = \mathtt{TMod}(\mathcal{K}_\varphi, \mathcal{J})$. While this seems insignificant, we can use it to conclude that ($\star$) $t_A \in C^{\mathcal{J}}_i$ (hence, the same holds for all $\mathcal{I} \in \mathcal{M}$) for $1 \leqslant i \leqslant n$. Suppose ($\star$) would not be true. There would be at least one model $\mathcal{I} \in \mathcal{M}$ and $i \in \{1, \cdots, n\}$ s.t. $t_A \notin C^{\mathcal{I}}_i$. From $\mathcal{M} = \mathtt{TMod}(\mathcal{K}_\varphi, \mathcal{J})$ ("$\mathcal{M}$ is as large as possible") we could then conclude that $\mathcal{M}$ also contains a model $\mathcal{I}' \supseteq \mathcal{I}$ with $t_A \notin C^{\mathcal{I}'}_i$ and $\mathcal{I}' \models s(t_A, t_B)$, without $\mathcal{I}'$ contradicting $\mathcal{T}_Q$. As this would contradict $s(t_A, t_B)$ not being Satisfiable, ($\star$) must be true.

From the containment of $t_A \in C_i^{\mathcal{I}}$ in all models $\mathcal{I} \in \mathcal{M}$ and their satisfaction of $\mathcal{T}^i_{clause}$ ($1 \leqslant i \leqslant n$), we conclude that (at least) one of the following must be true:

- $t_A \in (\exists r_{i_1}.X)^{\mathcal{I}}$,
- $t_A \in (\exists r_{i_2}.X)^{\mathcal{I}}$, or
- $t_A \in (\exists r_{i_3}.X)^{\mathcal{I}}$.

Consequently, by the construction of $\sigma_{\mathcal{M}}$, *at least one* variable per clause is assigned $\top$. $\mathcal{T}^i_{const}$ directly implies that for any set of models, in particular $\mathcal{M}$, no two preference options out of $r_{i_1}(t_A, t_B)$, $r_{i_2}(t_A, t_B)$ and $r_{i_3}(t_A, t_B)$ can be satisfied simultaneously for any $1 \leqslant i \leqslant n$. Hence, $\sigma_{\mathcal{M}}$ assigns *at most one* variable per clause to $\top$, concluding the proof of (ii). □

Even though in classical reasoning, a lower bound on subsumption checking immediately translates to a lower bound on instance checking, such a consequence is not trivial in defeasible semantics for DLs.

Remark 8.33. For a reduction to defeasible instance checking, consider a slight change in the context $\mathbb{C}$, $\mathbb{O}$, the KB $\mathcal{K}_\varphi$ and the query for which to decide non-entailment. The context $\mathbb{C} = \{B, X\}$ and $\mathbb{O} = \{a\}$ is still quantification closed and contains $\mathcal{K}_\varphi$. The KB is extended only with the ABox $\mathcal{A} = \{A(a)\}$, $\mathcal{T}$ and $\mathcal{D}$ remain as in (8.19). It is not hard to see that the rational domain contains $(a, \mathcal{U})$ instead of $(A, \mathcal{U})$ and $(A, \emptyset)$. Naturally, the new query must be a defeasible instance check, namely $\exists s.X\{a\}$. The proof of Proposition 8.32 remains intact, when the shorthand t_A refers to $(a, \mathcal{U})$. Hence, non-entailment of the defeasible instance check $\exists s.X\{a\}$ corresponds to satisfiability of φ.

The main result for the lower bound of defeasible reasoning under sceptical nested rational semantics is almost an immediate consequence of Proposition 8.32.

Theorem 8.34. *Deciding entailment of defeasible subsumption or instance checks under sceptical nested rational semantics is* co-NP*-hard for* $\mathcal{EL}_\bot$.

Proof. Satisfiability of a (1-in-3)-positive 3SAT formula is NP-hard, as shown by Garey and Johnson [GJ'79]. The reduction of this satisfiability problem to *non-entailment*, of defeasible subsumption (Prop. 8.32) and defeasible instance checking (Rem. 8.33), shows co-NP-hardness for the problem of deciding entailment of defeasible subsumption and instance checks under sceptical nested rational semantics. □

An overview and discussion of the complexity results obtained throughout this chapter is contained in the following, final chapter.

9 CONCLUSION, REFLECTION AND OUTLOOK

In this thesis we have ventured deep into the innards of rational reasoning in Description Logics, an area that provides a sizeable portion of the research towards non-monotonic DLs. We have reorganised and unified a variety of different entailment operations, that are all derived from the original materialisation-based KLM algorithm for deciding propositional entailment under Rational Closure. The generalisation of this algorithm into an abstract framework can be considered a reflection on the established literature, as well as a modernised contribution to this type of reasoning, accommodating several benefits. Interchangeability of single components in this framework shows almost trivially the commonalities and divergences among differently strong entailment operations that are defined in terms of materialisation. We have demonstrated the versatility of our framework by providing explicit instantiations that are capable of producing consequences under Rational, Relevant and Lexicographic Closure. For additional original contributions due to our framework, we lifted the reasoning service of defeasible instance checking to the Relevant Closure, and we have provided instantiations that rely entirely on $\mathcal{EL}_{\bot}$ concepts, a technique that is not trivially derived from $\mathcal{ALC}$ material implication. The latter proves the claim of Casini and Straccia [CS'10] that deciding entailments under Rational Closure in DLs remains in the complexity class of classical reasoning in the underlying DL. We have no doubt that this framework can be used to obtain further results and insights on the technique of materialisation-based defeasible reasoning. An instantiation of yet another type of Rational Closure, including a preprocessing by use of boolean inheritance networks ([CS'11]) should be as simple to reconstruct with our framework, as the lexicographic or relevant preprocessing step. Unifying these entailment operations allowed for a comprehensive discussion on their merits and drawbacks. Most notably, we discovered a neglect of quantified concepts that is inherent to all materialisation-based procedures. This fatal shortcoming for defeasible inference in DLs has remained unresolved for the better part of this decade.

We have identified the need to move beyond a *propositional* coverage of defeasible information, if meaningful defeasible consequences for Description Logics are to be derived. Taking advantage of representative interpretations and the canonical model property of the lightweight DL $\mathcal{EL}_{\bot}$, we have introduced a new kind of model formalism, explicitly attaching defeasible information (typicality) to subjects of a query and, more importantly, to *all* other domain elements that are potentially relevant in answering such a query. This allowed the formalisation of a new type of preference over sets of these *typicality models*, effectively eliminating evidence of low typicality. Preferring

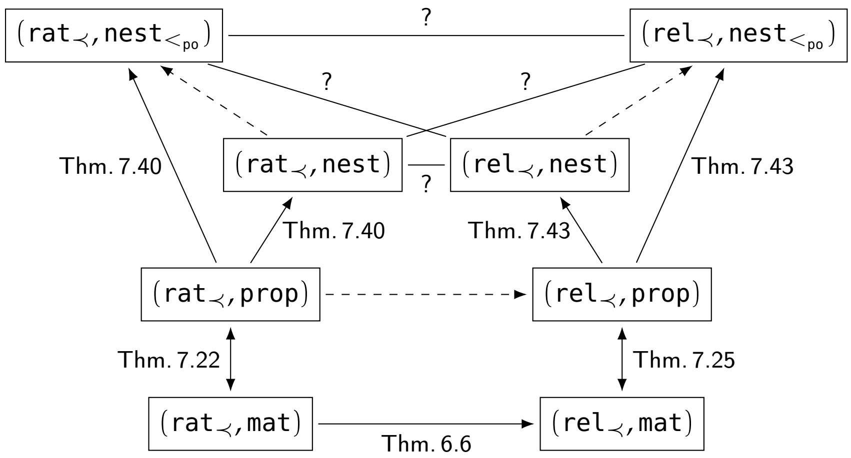

Figure 9.1: Implications among all of the investigated semantics of rational and relevant strength.

models of high typicality to be considered for deciding entailments, yields the desired effect of deriving defeasible information for concept elements in the *arbitrarily nested* relational neighbourhood of the query subject. Ultimately, we extended the toolbox for defeasible reasoning in DLs by composing previously investigated semantics of different strength with two types of *nested coverage*. Overall, we provided explicit results of superiority and complexity for all semantics captured by the tuples of strength and coverage

$$\{\mathtt{rat}_{\prec},\mathtt{rel}_{\prec}\}\times\{\mathtt{mat},\mathtt{prop},\mathtt{nest},\mathtt{nest}_{<_{\mathsf{po}}}\}$$

throughout this thesis.

We formally investigated relations among the covered semantics, by their virtue in terms of entailed queries, as illustrated in Figure 9.1. The direction of an arrow describes an implication for query entailment, e.g. if α is entailed by $\mathcal{K}$ under $(\mathtt{rat}_{\prec},\mathtt{prop})$ semantics, then it is also entailed under $(\mathtt{rat}_{\prec},\mathtt{nest})$ semantics. Bold arrows connect semantics whenever the proof for this result can be found in this thesis, dashed arrows show obvious implicit results and undirected edges indicate that the relation is unknown and could be addressed in the future. The edges in Fig. 9.1 are complete in the sense that a missing arrow indicates that this implication is not true, instead of unknown (modulo transitive implications). The only result in Fig. 9.1 that is not our own, is the implication of *defeasible subsumption entailment* from $(\mathtt{rat}_{\prec},\mathtt{mat})$-semantics to $(\mathtt{rel}_{\prec},\mathtt{mat})$-semantics (cf. [CMMN'14]). We first introduced defeasible instance checking in relevant semantics in [PT'18].

UNKNOWN ENTAILMENT IMPLICATIONS. Showing superiority of relevant over rational nested semantics, e.g. $(\mathtt{rel}_{\prec},\mathtt{nest})$ over $(\mathtt{rat}_{\prec},\mathtt{nest})$ simply requires to provide an appropriate example, similar to the proof for $(\mathtt{rel}_{\prec},\mathtt{mat})$ vs. $(\mathtt{rat}_{\prec},\mathtt{mat})$ (Thm. 6.6). Proving an implication from any rational nested semantics to its relevant variant, is much

	Rational	Relevant
Propositional	P-complete	in Δ_2^{P}
Selective nested	P-complete	in EXP
Sceptical nested	CO-NP-complete	in CO-NEXP

Table 9.1: Complexity for deciding defeasible subsumption and instance checking with typicality models.

more involved. For showing such a relation for selective nested semantics, one would have to be especially careful with the order $<_{\mathsf{po}}$, as it is defined on the underlying typicality domain. At least for concept representatives, it is clearly the case that all representatives of the rational domain also belong to the relevant domain, meaning, a total order for the rational domain (w.r.t. concept representatives) is a strict partial over the relevant domain. However, the same is not true if $\mathbb{O} \neq \emptyset$. The relevant and rational domain could potentially associate different DCIs with every individual representative, making it unclear how to derive a preference option order for the relevant domain, which "behaves" like the order over the rational domain. If no connection between the respective orders is imposed, this implication will surely not hold.

We suspect the main difficulty in the relation between sceptical nested rational and relevant semantics to be closely related to the reason for initialising the relevant ABox extension (Def. 4.16) with the rationally extended ABox (cf. Exm. 6.5). Preference options of "higher typicality than rationally possible" could be in conflict with preference options that are conflict free in the rational typicality domain. We leave a careful investigation to test this suspicion for future work.

COMPUTATIONAL COMPLEXITY. Any contribution to the reasoning capabilities in Description Logics is encouraged to discuss its computational properties in addition to its expressivity. In conclusion of our analysis for the discussed typicality-model semantics, observe Table 9.1 for a concise overview of the computational complexity for deciding defeasible subsumption and instance checks under propositional and nested coverage. Theorem 8.22, 8.24 and 8.30 provide matching lower and upper bounds, showing P-completeness for propositional and selective nested rational semantics. This result comes at no surprise for propositional semantics due to its equivalence to the materialisation reduction (Thm. 7.22). On the other hand, it is encouraging, that a deterministic preference of role-edges allows to draw conclusions based on maximal typicality, i.e. including defeasible information at quantified concepts, without increasing the *computational complexity*. When no such preference exists, an application might employ the more cautious approach with sceptical nested coverage. Theorem 8.23 and 8.34 show that this scepticism and the requirement to verify a query w.r.t. all possible typicality maximisations comes at considerable extra cost, being CO-NP-complete.

At this time, our results for relevant strength of all three investigated coverages remain upper bounds for deciding entailments (Sec. 8.2.2). PTime is of course a trivial lower bound for these entailment problems, as all typicality model semantics extend classical reasoning in $\mathcal{EL}_\bot$. While materialisation-based relevant semantics for $\mathcal{ALC}$ [CMMN'14] are more expensive already in the underlying classical reasoning (i.e. ExpTime-complete, [Sch'91a]), it remains exponential to determine all justifications of an entailment. This shows ExpTime-completeness of subsumption and instance checking for materialisation-based relevant entailment in $\mathcal{ALC}$. For $\mathcal{EL}_\bot$, classical reasoning is less complex than the computation of all (or some) justifications, hence, a tight lower bound for relevant strength does not transfer as in the case of $\mathcal{ALC}$. Even though the computation of justifications (as required here) is known to be NP-hard ([PS'17]), this lower bound does not trivially translate to any of our semantics, because the entailments we calculate justifications for, are of a very specific type (e.g. $\mathsf{C} \sqsubseteq \bot$). As it is customary, appropriate, known decision problems and individually tailored reductions would be required, to close these complexity gaps for relevant strength of defeasible reasoning with typicality models, in the future.

9.1 DIRECTIONS FOR FUTURE INVESTIGATIONS

While selective nested coverage is computationally less costly than its sceptical counterpart, one could argue that there is a different kind of extra cost, hidden in the required inputs to characterise specific semantics with $\mathtt{nest}_{<_{\mathsf{po}}}$. At this time, we presented this parametrisation as general as possible, simply assuming a total order $<_{\mathsf{po}}$ over the preference options of a typicality domain as input. It seems cumbersome to practically specify this order, especially as it might be unlikely that such a preference even exists from an application's point of view. Therefore, more practicable formalisations of this requirement (as briefly discussed in Section 8.1.2) could be an interesting direction for future research. Additionally, it is clearly the case that $<_{\mathsf{po}}$ influences the entailments under selective nested semantics by a great deal. We suspect that, if this order satisfies certain formal properties, it could be used to derive more intuitively understood characteristics for the resulting entailment relation, similar to KLM-style postulates.

On that topic, a full investigation of well-behaviour for nested semantics has purposefully not been addressed in this thesis. While we presented many results, typicality models are still in an early stage of their development. Refinements, such as the above, need further analysis, before a discussion of formal properties is useful. At this time, we can only point to the postulates that are potentially most difficult to derive for the typicality model formalism. Satisfaction of Rational Monotonicity (RM) and Cautious Monotonicity (CM) in representative models requires the entailment of corresponding properties for two distinct representative elements (e.g. $(\mathsf{C}, \mathcal{U})$ and $(\mathsf{C} \sqcap \mathsf{E}, \mathcal{U}'))$.

$$\text{(CM)}\ \frac{C \mathrel{\sqsubseteq\!\!\!\sim} D \quad C \mathrel{\sqsubseteq\!\!\!\sim} E}{C \sqcap E \mathrel{\sqsubseteq\!\!\!\sim} D} \qquad\qquad \text{(RM)}\ \frac{C \mathrel{\sqsubseteq\!\!\!\sim} D \quad C \not\mathrel{\sqsubseteq\!\!\!\sim} \neg E}{C \sqcap E \mathrel{\sqsubseteq\!\!\!\sim} D}$$

While the construction of the typicality domain guarantees this correspondence on a propositional level, the typicality preference relation does not necessarily behave appropriately.

It is worth noting that certain properties of the interplay between strict and defeasible knowledge, could result in the set of all Max-TMs to be a singleton. For example, if no two preference options contradict each other in the sense of Example 7.34, it is easy to see that any preference $<_{\mathsf{po}}$ produces the same unique and only Max-TM. A trivial instance with this property is simply any $\mathcal{EL}$ defeasible knowledge base, refraining from the use of $\bot$. Under such conditions, expressivity and complexity for semantics of sceptical and selective nested coverage would coincide. Once again, a more thorough analysis exceeds the scope of this work, but is certainly of great interest in the immediate future of typicality models.

Furthermore, with the introduction of any new non-monotonic formalism, its relation to surrounding approaches should be analysed. The comparison to the foundations of Rational and Relevant Closure is extensively studied in this work. As far as we know, no other KLM-style investigation has explicitly addressed practical means to derive defeasible information for quantified concepts in DLs. Somewhat related might be the approach to model defeasibility on the level of roles [BCMV'13; BV'16; BV'17a; BV'17b], but most of this research is studying non-practical representation results. It is not clear how the defeasible entailment of a role relation could lead to the entailment of defeasible concept properties in the filler of a quantification. Varzinczak [Var'18] investigates a concept/role constructor for typicality, that can be used anywhere in a concept, in particular also in nested quantification. This would provide the means to individually capture typicality of nested concepts, instead of a uniform maximisation approach such as ours. Once again, at this time only representation results are studied, positioning this approach in the semantic characterisability aspect of KLM-style research.

Finally, the typicality model formalism provides only the foundation of nested reasoning capabilities. In addition to the above and many more analytical questions, expressive extensions of typicality models and their capacities could be examined. Immediate options are the increase of expressivity within the lightweight framework of DLs, e.g. allowing for nominals or role constructions/hierarchies, or even considering more expressive reasoning services, such as conjunctive query answering. By remaining with DLs that satisfy the canonical model property, the representative approach towards typicality models requires potentially only slight adjustments. Surely, the most valuable and interesting direction however, is to generalise nested typicality coverage, to not rely on the canonical model property of lightweight DLs, and thus, to achieve truly meaningful defeasible consequence in expressive Description Logics.

[AGM'85] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions." In: *The Journal of Symbolic Logic* 50.2 (1985), pp. 510–530 (cit. on p. 30).

[AB'94] Krzysztof R. Apt and Roland N. Bol. "Logic Programming and Negation: A Survey." In: *The Journal of Logic Programming* 19 (1994), pp. 9–71 (cit. on p. 32).

[Baa'03] Franz Baader. "Terminological Cycles in a Description Logic with Existential Restrictions." In: *Proceedings of the 18th International Joint Conference on Artificial Intelligence, IJCAI*. 2003, pp. 325–330 (cit. on p. 134).

[BBL'05] Franz Baader, Sebastian Brandt, and Carsten Lutz. "Pushing the $\mathcal{EL}$ Envelope." In: *Proceedings of the 19th International Joint Conference on Artificial Intelligence, IJCAI*. 2005, pp. 364–369 (cit. on pp. 13, 22).

[BCM+'10] Franz Baader, Diego Calvanese, Deborah McGuinness, Peter Patel-Schneider, and Daniele Nardi. *The Description Logic Handbook: Theory, Implementation and Applications*. 2nd Edition. Cambridge University Press, 2010. ISBN: 978-0-521-15011-8 (cit. on pp. 13, 21).

[BH'92] Franz Baader and Bernhard Hollunder. "Embedding Defaults into Terminological Knowledge Representation Formalisms." In: *Principles of Knowledge Representation and Reasoning: Proceedings of the 3rd International Conference, KR*. 1992, pp. 306–317 (cit. on p. 32).

[BH'95a] Franz Baader and Bernhard Hollunder. "Embedding Defaults into Terminological Knowledge Representation Formalisms." In: *Journal of Automated Reasoning* 14.1 (1995), pp. 149–180 (cit. on p. 32).

[BH'95b] Franz Baader and Bernhard Hollunder. "Priorities on Defaults with Prerequisites, and Their Application in Treating Specificity in Terminological Default Logic." In: *Journal of Automated Reasoning* 15.1 (1995), pp. 41–68 (cit. on pp. 32, 34).

[BHLS'17] Franz Baader, Ian Horrocks, Carsten Lutz, and Ulrike Sattler. *An Introduction to Description Logic*. Cambridge University Press, 2017. ISBN: 978-0-521-69542-8 (cit. on pp. 13, 21).

[BLB'08] Franz Baader, Carsten Lutz, and Sebastian Brandt. "Pushing the $\mathcal{EL}$ Envelope Further." In: *Proceedings of the 4th OWLED Workshop on OWL: Experiences and Directions*. 2008 (cit. on p. 13).

[BP'10a] Franz Baader and Rafael Peñaloza. "Automata-Based Axiom Pinpointing." In: *Journal of Automated Reasoning* 45.2 (2010), pp. 91–129 (cit. on p. 69).

[BP'10b] Franz Baader and Rafael Peñaloza. "Axiom Pinpointing in General Tableaux." In: *Journal of Logic and Computation* 20.1 (2010), pp. 5–34 (cit. on p. 69).

[BPS'07] Franz Baader, Rafael Peñaloza, and Boontawee Suntisrivaraporn. "Pinpointing in the Description Logic $\mathcal{EL}$." In: *Proceedings of the 20th International Workshop on Description Logics, DL*. 2007 (cit. on p. 138).

[Bel'77] Nuel D. Belnap. "A Useful Four-Valued Logic." In: *Modern Uses of Multiple-Valued Logic*. Springer, 1977, pp. 5–37 (cit. on p. 30).

[BCD+'93] Salem Benferhat, Claudette Cayrol, Didier Dubois, Jérôme Lang, and Henri Prade. "Inconsistency Management and Prioritized Syntax-Based Entailment." In: *Proceedings of the 13th International Joint Conference on Artificial Intelligence, IJCAI*. 1993, pp. 640–647 (cit. on p. 73).

[BHS'05] Leopoldo Bertossi, Anthony Hunter, and Torsten Schaub. *Inconsistency Tolerance*. Springer, 2005. ISBN: 978-3-540-24260-4 (cit. on p. 30).

[BR'13] Meghyn Bienvenu and Riccardo Rosati. "Tractable Approximations of Consistent Query Answering for Robust Ontology-based Data Access." In: *Proceedings of the 23rd International Joint Conference on Artificial Intelligence, IJCAI*. 2013, pp. 775–781 (cit. on p. 30).

[Bon'19] Piero A. Bonatti. "Rational Closure for All Description Logics." In: *Artificial Intelligence, AIJ* 274 (2019), pp. 197–223 (cit. on pp. 24, 39, 49, 61, 62, 68, 79, 88).

[BFPS'15] Piero A. Bonatti, Marco Faella, Iliana M. Petrova, and Luigi Sauro. "A New Semantics for Overriding in Description Logics." In: *Artificial Intelligence, AIJ* 222 (2015), pp. 1–48 (cit. on pp. 32, 61, 62, 65).

[BFS'10] Piero A. Bonatti, Marco Faella, and Luigi Sauro. "$\mathcal{EL}$ with Default Attributes and Overriding." In: *The Semantic Web - 9th International Semantic Web Conference, ISWC, Revised Selected Papers, Part I*. 2010, pp. 64–79 (cit. on p. 32).

[BFS'11] Piero A. Bonatti, Marco Faella, and Luigi Sauro. "Defeasible Inclusions in Low-Complexity DLs." In: *Artificial Intelligence, AIJ* 42.1 (2011), pp. 719–764 (cit. on p. 31).

[BLW'06] Piero A. Bonatti, Carsten Lutz, and Frank Wolter. "Description Logics with Circumscription." In: *Principles of Knowledge Representation and Reasoning: Proceedings of the 10th International Conference, KR*. 2006, pp. 400–410 (cit. on p. 31).

[BLW'09] Piero A. Bonatti, Carsten Lutz, and Frank Wolter. "The Complexity of Circumscription in DLs." In: *Artificial Intelligence, AIJ* 35 (2009), pp. 717–773 (cit. on p. 31).

[BS'17] Piero A. Bonatti and Luigi Sauro. "On the Logical Properties of the Nonmonotonic Description Logic $\mathcal{DL}^{\mathsf{N}}$." In: *Artificial Intelligence, AIJ* 248 (2017), pp. 85–111 (cit. on pp. 32, 61).

[Bor'96] Alexander Borgida. "On the relative expressiveness of description logics and predicate logics." In: *Artificial Intelligence, AIJ* 82.1-2 (1996), pp. 353–367 (cit. on p. 14).

[Bou'16] Camille Bourgaux. "Inconsistency Handling in Ontology-Mediated Query Answering." PhD thesis. University of Paris-Saclay, France, 2016. URL: `https://tel.archives-ouvertes.fr/tel-01378723` (cit. on pp. 30, 109).

[Bou'94] Craig Boutilier. "Conditional Logics of Normality: A Modal Approach." In: *Artificial Intelligence, AIJ* 68.1 (1994), pp. 87–154 (cit. on p. 33).

[Bra'04] Sebastian Brandt. "Polynomial Time Reasoning in a Description Logic with Existential Restrictions, GCI Axioms, and - What Else?" In: *16th European Conference on Artificial Intelligence, ECAI*. 2004, pp. 298–302 (cit. on pp. 22, 134, 135).

[Bre'89] Gerhard Brewka. "Nonmonotonic Logics - A Brief Overview." In: *AI Communications* 2.2 (1989), pp. 88–97 (cit. on p. 29).

[BDK'97] Gerhard Brewka, Jürgen Dix, and Kurt Konolige. *Nonmonotonic Reasoning: An Overview*. Vol. 73. CSLI Lecture Notes. CSLI, 1997. ISBN: 978-1-881-52683-4 (cit. on p. 29).

[BCM+'13] Katarina Britz, Giovanni Casini, Thomas Meyer, Kody Moodley, and Ivan J. Varzinczak. *Ordered Interpretations and Entailment for Defeasible Description Logics*. Technical Report. Accessed on Dec. 13 2018, 12:45 (UTC+1). UKZN CSIR Meraka Centre for Artificial Intelligence Research, South Africa, 2013. URL: `http://www.cair.za.net/sites/default/files/outputs/DefeasibleDLs-TR1-2013_0.pdf` (cit. on pp. 39, 47, 50, 51, 61, 79).

[BCMV'13] Katarina Britz, Giovanni Casini, Thomas Meyer, and Ivan J. Varzinczak. "Preferential Role Restrictions." In: *Proceedings of the 26th International Workshop on Description Logics, DL*. 2013, pp. 93–106 (cit. on pp. 79, 155).

[BMV'11a] Katarina Britz, Thomas Meyer, and Ivan J. Varzinczak. "Preferential Reasoning for Modal Logics." In: *Electronic Notes in Theoretical Computer Science* 278 (2011), pp. 55–69 (cit. on p. 61).

[BMV'11b] Katarina Britz, Thomas Meyer, and Ivan J. Varzinczak. "Semantic Foundation for Preferential Description Logics." In: *AI 2011: Advances in Artificial Intelligence - 24th Australasian Joint Conference*. 2011, pp. 491–500 (cit. on pp. 33, 39, 47, 50, 51, 56).

[BV'16] Katarina Britz and Ivan J. Varzinczak. "Introducing Role Defeasibility in Description Logics." In: *Logics in Artificial Intelligence - 15th European Conference, JELIA*. 2016, pp. 174–189 (cit. on pp. 39, 79, 155).

[BV'17a] Katarina Britz and Ivan J. Varzinczak. "Context-based defeasible subsumption for d$\mathcal{SROIQ}$." In: *Proceedings of the 13th International Symposium on Commonsense Reasoning, COMMONSENSE*. 2017 (cit. on pp. 39, 49, 61, 67, 79, 155).

[BV'17b] Katarina Britz and Ivan J. Varzinczak. "Towards Defeasible $\mathcal{SROIQ}$." In: *Proceedings of the 30th International Workshop on Description Logics, DL*. 2017 (cit. on pp. 49, 61, 79, 155).

[BV'18a] Katarina Britz and Ivan J. Varzinczak. "From KLM-style conditionals to defeasible modalities, and back." In: *Journal of Applied Non-Classical Logics* 28.1 (2018), pp. 92–121 (cit. on p. 32).

[BV'18b] Katarina Britz and Ivan J. Varzinczak. "Rationality and Context in Defeasible Subsumption." In: *Foundations of Information and Knowledge Systems - 10th International Symposium, FoIKS*. 2018, pp. 114–132 (cit. on pp. 39, 67, 79).

[CMMN'14] Giovanni Casini, Thomas Meyer, Kodylan Moodley, and Riku Nortje. "Relevant Closure: A New Form of Defeasible Reasoning for Description Logics." In: *Logics in Artificial Intelligence - 14th European Conference, JELIA*. 2014, pp. 92–106 (cit. on pp. 7, 35, 39, 40, 47, 50, 51, 61, 63, 67, 68, 70, 72, 73, 79, 115, 152, 154).

[CMM+'15] Giovanni Casini, Thomas Meyer, Kodylan Moodley, Uli Sattler, and Ivan J. Varzinczak. "Introducing Defeasibility into OWL Ontologies." In: *The Semantic Web - 14th International Semantic Web Conference, ISWC*. 2015, pp. 409–426 (cit. on pp. 36, 39, 40, 50, 115).

[CMMV'13] Giovanni Casini, Thomas Meyer, Kodylan Moodley, and Ivan J. Varzinczak. "Towards Practical Defeasible Reasoning for Description Logics." In: *Proceedings of the 26th International*

Workshop on Description Logics, DL. 2013, pp. 587–599 (cit. on pp. 36, 39, 40, 50, 115).

[CMS'18] Giovanni Casini, Thomas Meyer, and Umberto Straccia. "A polynomial Time Subsumption Algorithm for Nominal Safe $\mathcal{ELO}_\bot$ under Rational Closure." In: *Information Sciences* (2018) (cit. on pp. 39, 49, 51, 61, 134).

[CMVM'13] Giovanni Casini, Thomas Meyer, Ivan J. Varzinczak, and Kodylan Moodley. "Nonmonotonic Reasoning in Description Logics: Rational Closure for the ABox." In: *Proceedings of the 26th International Workshop on Description Logics, DL*. 2013, pp. 600–615 (cit. on pp. 20, 36, 39, 40, 50–52, 55).

[CS'10] Giovanni Casini and Umberto Straccia. "Rational Closure for Defeasible Description Logics." In: *Logics in Artificial Intelligence - 12th European Conference, JELIA2010*. 2010, pp. 77–90 (cit. on pp. 7, 20, 36, 39, 40, 44, 50–53, 65, 134, 151).

[CS'11] Giovanni Casini and Umberto Straccia. "Defeasible Inheritance-Based Description Logics." In: *Proceedings of the 22nd International Joint Conference on Artificial Intelligence, IJCAI*. 2011, pp. 813–818 (cit. on pp. 7, 63, 67, 151).

[CS'12] Giovanni Casini and Umberto Straccia. "Lexicographic Closure for Defeasible Description Logics." In: *Proceedings of the 8th Australasian Ontology Workshop*. 2012, pp. 28–39 (cit. on pp. 7, 20, 36, 39, 40, 50–52, 61, 63, 67, 73–75, 79).

[CS'13] Giovanni Casini and Umberto Straccia. "Defeasible Inheritance-Based Description Logics." In: *Journal of Artificial Intelligence Research, JAIR* 48 (2013), pp. 415–473 (cit. on pp. 36, 39, 40, 50–52, 61, 63, 67, 79).

[CD'91] John W. Chinneck and Erik W. Dravnieks. "Locating Minimal Infeasible Constraint Sets in Linear Programs." In: *ORSA Journal on Computing* 3.2 (1991), pp. 157–168 (cit. on p. 69).

[Cla'78] Keith L. Clark. "Negation as Failure." In: *Logic and Data Bases*. Springer, 1978, pp. 293–322 (cit. on p. 32).

[CBH'92] William W. Cohen, Alexander Borgida, and Haym Hirsh. "Computing Least Common Subsumers in Description Logics." In: *Proceedings of the 10th National Conference on Artificial Intelligence*. 1992, pp. 754–760 (cit. on p. 23).

[Del'98] James P. Delgrande. "On first-order conditional logics." In: *Artificial Intelligence, AIJ* 105.1-2 (1998), pp. 105–137 (cit. on p. 33).

[DLN+'98] Francesco M. Donini, Maurizio Lenzerini, Daniele Nardi, Werner Nutt, and Andrea Schaerf. "An Epistemic Operator for Description Logics." In: *Artificial Intelligence, AIJ* 100.1-2 (1998), pp. 225–274 (cit. on p. 32).

[DLN+'92] Francesco M. Donini, Maurizio Lenzerini, Daniele Nardi, Andrea Schaerf, and Werner Nutt. "Adding Epistemic Operators to Concept Languages." In: *Principles of Knowledge Representation and Reasoning: Proceedings of the 3rd International Conference, KR*. 1992, pp. 342–353 (cit. on p. 32).

[DM'00] Francesco M. Donini and Fabio Massacci. "ExpTime tableaux for $\mathcal{ALC}$." In: *Artificial Intelligence, AIJ* 124.1 (2000), pp. 87–138 (cit. on p. 22).

[DNR'97] Francesco M. Donini, Daniele Nardi, and Riccardo Rosati. "Autoepistemic Description Logics." In: *Proceedings of the 15th International Joint Conference on Artificial Intelligence, IJCAI*. 1997, pp. 136–141 (cit. on p. 32).

[Eck'17] Andreas Ecke. "Quantitative Methods for Similarity in Description Logics." PhD thesis. Dresden University of Technology, Germany, 2017. URL: http://nbn-resolving.de/urn:nbn:de:bsz:14-qucosa-223626 (cit. on p. 23).

[EPT'15] Andreas Ecke, Maximilian Pensel, and Anni-Yasmin Turhan. "Elastiq: Answering Similarity-threshold Instance Queries in $\mathcal{EL}$." In: *Proceedings of the 26th International Workshop on Description Logics, DL*. 2015 (cit. on p. 23).

[GHN+'04] Harald Ganzinger, George Hagen, Robert Nieuwenhuis, Albert Oliveras, and Cesare Tinelli. "DPLL(T): Fast Decision Procedures." In: *International Conference on Computer Aided Verification*. Springer. 2004, pp. 175–188 (cit. on p. 69).

[Gär'92] Peter Gärdenfors. *Belief Revision*. Cambridge University Press, 1992. ISBN: 978-0-521-41260-5 (cit. on p. 30).

[GJ'79] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979. ISBN: 978-0-716-71044-8 (cit. on pp. 142, 150).

[GD'16] Laura Giordano and Daniele T. Dupré. "ASP for minimal entailment in a rational extension of $\mathcal{SROEL}$." In: *Theory and Practice of Logic Programming* 16.5-6 (2016), pp. 738–754 (cit. on pp. 61, 63).

[GD'18] Laura Giordano and Daniele T. Dupré. "Defeasible Reasoning in $\mathcal{SROEL}$: from Rational Entailment to Rational Closure." In: *Fundamenta Informaticae* 161.1-2 (2018), pp. 135–161 (cit. on p. 61).

[GG'18] Laura Giordano and Valentina Gliozzi. "Reasoning About Exceptions in Ontologies: from the Lexicographic Closure to the Skeptical Closure." In: *Proceedings of the 2nd Workshop on Logics for Reasoning about Preferences, Uncertainty, and Vagueness*. 2018 (cit. on pp. 39, 61, 63, 67, 79).

[GGO'18] Laura Giordano, Valentina Gliozzi, and Nicola Olivetti. "Towards a Rational Closure for Expressive Description Logics: the Case of $\mathcal{SHIQ}$." In: *Fundamenta Informaticae* 159.1-2 (2018), pp. 95–122 (cit. on pp. 61, 79).

[GGOP'07] Laura Giordano, Valentina Gliozzi, Nicola Olivetti, and Gian L. Pozzato. "Preferential Description Logics." In: *Logic for Programming, Artificial Intelligence, and Reasoning, 14th International Conference, LPAR*. 2007, pp. 257–272 (cit. on p. 39).

[GGOP'10a] Laura Giordano, Valentina Gliozzi, Nicola Olivetti, and Gian L. Pozzato. "Preferential vs Rational Description Logics: which one for Reasoning About Typicality?" In: *19th European Conference on Artificial Intelligence, ECAI*. 2010, pp. 1069–1070 (cit. on p. 61).

[GGOP'10b] Laura Giordano, Valentina Gliozzi, Nicola Olivetti, and Gian L. Pozzato. "Reasoning About Typicality in Preferential Description Logics: Preferential vs Rational Entailment." In: *Proceedings of the 25th Italian Conference on Computational Logic*. 2010 (cit. on pp. 39, 61).

[GGOP'12] Laura Giordano, Valentina Gliozzi, Nicola Olivetti, and Gian L. Pozzato. "A Minimal Model Semantics for Nonmonotonic Reasoning." In: *Logics in Artificial Intelligence - 13th European Conference, JELIA*. 2012, pp. 228–241 (cit. on p. 39).

[GGOP'15] Laura Giordano, Valentina Gliozzi, Nicola Olivetti, and Gian L. Pozzato. "Semantic characterization of rational closure: From propositional logic to description logics." In: *Artificial Intelligence, AIJ* 226 (2015), pp. 1–33 (cit. on pp. 36, 39, 40, 50, 52, 61).

[GGPR'17] Laura Giordano, Valentina Gliozzi, Gian L. Pozzato, and Riccardo Renzulli. "An Efficient Reasoner for Description Logics of Typicality and Rational Closure." In: *Proceedings of the 30th International Workshop on Description Logics, DL*. 2017, pp. 1–13 (cit. on p. 115).

[GOGP'09] Laura Giordano, Nicola Olivetti, Valentina Gliozzi, and Gian L. Pozzato. "$\mathcal{ALC}$ + T: a Preferential Extension of Description Logics." In: *Fundamenta Informaticae* 96.3 (2009), pp. 341–372 (cit. on p. 39).

[GHM+'14] Birte Glimm, Ian Horrocks, Boris Motik, Giorgos Stoilos, and Zhe Wang. "HermiT: An OWL 2 Reasoner." In: *Journal of Automated Reasoning* 53.3 (2014), pp. 245–269 (cit. on p. 115).

[HLP'08] Frank van Harmelen, Vladimir Lifschitz, and Bruce Porter. *Handbook of Knowledge Representation*. Vol. 1. Elsevier Science, 2008. ISBN: 978-0-444-52211-5 (cit. on p. 1).

[HO'48] Carl G. Hempel and Paul Oppenheim. "Studies in the Logic of Explanation." In: *Philosophy of Science* 15.2 (1948), pp. 135–175 (cit. on p. 69).

[Hew'72] Carl Hewitt. "Description and Theoretical Analysis (Using Schemata) of Planner: A Language for Proving Theorems and Manipulating Models in a Robot." PhD thesis. Massachusetts Institute of Technology, Cambridge Artificial Intelligence Lab, 1972. URL: `https://apps.dtic.mil/docs/citations/AD0744620` (cit. on p. 2).

[Hor'11] Matthew Horridge. "Justification Based Explanation in Ontologies." PhD thesis. University of Manchester, UK, 2011. URL: `http://www.manchester.ac.uk/escholar/uk-ac-man-scw:131699` (cit. on pp. 67–69).

[HPS'09] Matthew Horridge, Bijan Parsia, and Ulrike Sattler. "Explaining Inconsistencies in OWL Ontologies." In: *Scalable Uncertainty Management, 3rd International Conference, SUM*. 2009, pp. 124–137 (cit. on pp. 68, 69).

[HPH'03] Ian Horrocks, Peter Patel-Schneider, and Frank van Harmelen. "From $\mathcal{SHIQ}$ and RDF to OWL: the making of a Web Ontology Language." In: *Journal of Web Semantics* 1.1 (2003), pp. 7–26 (cit. on p. 3).

[Hor'94] John F. Horty. "Some Direct Theories of Nonmonotonic Inheritance." In: *Handbook of Logic in Artificial Intelligence and Logic Programming*. Ed. by Dov M. Gabbay, Christopher J. Hogger, and John A. Robinson. Vol. 3. 1994, pp. 111–187 (cit. on p. 67).

[KKS'14] Yevgeny Kazakov, Markus Krötzsch, and Frantisek Simancik. "The Incredible ELK - From Polynomial Procedures to Efficient Reasoning with $\mathcal{EL}$ Ontologies." In: *Journal of Automated Reasoning* 53.1 (2014), pp. 1–61 (cit. on p. 115).

[KLM'90] Sarit Kraus, Daniel J. Lehmann, and Menachem Magidor. "Nonmonotonic Reasoning, Preferential Models and Cumulative Logics." In: *Artificial Intelligence, AIJ* 44.1-2 (1990), pp. 167–207 (cit. on pp. 3, 32–34, 61).

[Leh'95] Daniel J. Lehmann. "Another Perspective on Default Reasoning." In: *Annals of Mathematics and Artificial Intelligence* 15.1 (1995), pp. 61–82 (cit. on pp. 63, 73, 79).

[LM'90] Daniel J. Lehmann and Menachem Magidor. "Preferential Logics: the Predicate Calculus Case." In: *Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge*. 1990, pp. 57–72 (cit. on pp. 33, 36, 61, 79).

[LM'92] Daniel J. Lehmann and Menachem Magidor. "What does a Conditional Knowledge Base Entail?" In: *Artificial Intelligence, AIJ* 55.1 (1992), pp. 1–60 (cit. on pp. 3, 33–36, 39, 47, 51, 61, 63, 67).

[LLR+'10] Domenico Lembo, Maurizio Lenzerini, Riccardo Rosati, Marco Ruzzi, and Domenico F. Savo. "Inconsistency-Tolerant Semantics for Description Logics." In: *Proceedings of Web Reasoning and Rule Systems - 4th International Conference (RR)*. 2010, pp. 103–117 (cit. on p. 30).

[LS'05] Mark H. Liffiton and Karem A. Sakallah. "On Finding All Minimally Unsatisfiable Subformulas." In: *International Conference on Theory and Applications of Satisfiability Testing*. Springer. 2005, pp. 173–186 (cit. on p. 69).

[Lou'87] Ronald P. Loui. "Theory and Computation of Uncertain Inference and Decision." PhD thesis. University of Rochester, USA, 1987. URL: `http://hdl.handle.net/1802/6159` (cit. on pp. 21, 34).

[LW'10] Carsten Lutz and Frank Wolter. "Deciding inseparability and conservative extensions in the description logic $\mathcal{EL}$." In: *Journal of Symbolic Computation* 45.2 (2010), pp. 194–228 (cit. on pp. 22, 24, 57).

[MMH'13] Frederick Maier, Yue Ma, and Pascal Hitzler. "Paraconsistent OWL and Related Logics." In: *Semantic Web* 4.4 (2013), pp. 395–427 (cit. on p. 30).

[McC'80] John McCarthy. "Circumscription - A Form of Non-Monotonic Reasoning." In: *Artificial Intelligence, AIJ* 13.1-2 (1980), pp. 27–39 (cit. on pp. 2, 31).

[Moo'85] Robert C. Moore. "Semantical Considerations on Nonmonotonic Logic." In: *Artificial Intelligence, AIJ* 25.1 (1985), pp. 75–94 (cit. on p. 32).

[Neb'90] Bernhard Nebel. *Reasoning and Revision in Hybrid Representation Systems*. Vol. 422. Lecture Notes in Computer Science. Springer, 1990. ISBN: 978-3-540-52443-4 (cit. on p. 23).

[Nut'01] Donald Nute. "Defeasible Logic." In: *Web Knowledge Management and Decision Support, 14th International Conference on Applications of Prolog, INAP*. 2001, pp. 151–169 (cit. on pp. 1, 21, 34).

[Peñ'09] Rafael Peñaloza. "Axiom-Pinpointing in Description Logics and Beyond." PhD thesis. Dresden University of Technology, Germany, 2009. URL: http://nbn-resolving.de/urn:nbn:de:bsz:14-qucosa-24743 (cit. on pp. 67, 69).

[PS'17] Rafael Peñaloza and Baris Sertkaya. "Understanding the complexity of axiom pinpointing in lightweight description logics." In: *Artificial Intelligence, AIJ* 250 (2017), pp. 80–104 (cit. on pp. 139, 154).

[PT'17a] Maximilian Pensel and Anni-Yasmin Turhan. "Including Quantification in Defeasible Reasoning for the Description Logic $\mathcal{EL}_\bot$." In: *Logic Programming and Nonmonotonic Reasoning, 14th International Conference, LPNMR*. Springer, 2017, pp. 78–84 (cit. on pp. 8, 23, 40, 61, 65, 79, 98, 101, 107, 108, 117, 122).

[PT'17b] Maximilian Pensel and Anni-Yasmin Turhan. "Making Quantification Relevant again—the Case of Defeasible $\mathcal{EL}_\bot$." In: *Proceedings of the 4th International Workshop on Defeasible and Ampliative Reasoning - DARe*. CEUR-WS.org, 2017 (cit. on pp. 8, 40, 61, 79, 101, 107, 108, 117, 122).

[PT'17c] Maximilian Pensel and Anni-Yasmin Turhan. *Making Quantification Relevant again—the Case of Defeasible $\mathcal{EL}_\bot$*. LTCS-Report 17-01. Accessed on Jun. 11 2019, 15:00 (UTC+2). Chair for Automata Theory, Institute for Theoretical Computer Science, Technische Universität Dresden, Germany, 2017. URL: https://lat.inf.tu-dresden.de/research/reports/2017/PeTu-LTCS-17-01.pdf (cit. on p. 134).

[PT'18] Maximilian Pensel and Anni-Yasmin Turhan. "Reasoning in the Defeasible Description Logic $\mathcal{EL}_\bot$—Computing Standard Inferences under Rational and Relevant Semantics." In: *International Journal of Approximate Reasoning (IJAR)* 103 (2018), pp. 28–70 (cit. on pp. 23, 61, 72, 98, 103, 107, 108, 122, 142, 152).

[QY'08] Guilin Qi and Fangkai Yang. "A Survey of Revision Approaches in Description Logics." In: *Proceedings of Web Reasoning and Rule Systems - 2nd International Conference (RR)*. Berlin, Heidelberg: Springer, 2008, pp. 74–88 (cit. on p. 30).

[RBG+'97] Alan L. Rector, Sean Bechhofer, Carole A. Goble, Ian Horrocks, Anthony Nowlan, and Danny Solomon. "The GRAIL Concept Modelling Language for Medical Terminology." In: *Artificial Intelligence in Medicine* 9.2 (1997), pp. 139–171 (cit. on p. 4).

[Rei'78] Raymond Reiter. "On Reasoning by Default." In: *Theoretical Issues in Natural Language Processing-2, TINLAP*. 1978 (cit. on p. 31).

[Rei'80] Raymond Reiter. "A Logic for Default Reasoning." In: *Artificial Intelligence, AIJ* 13.1-2 (1980), pp. 81–132 (cit. on pp. 2, 31, 34, 52, 109).

[Rei'87] Raymond Reiter. "A Theory of Diagnosis from First Principles." In: *Artificial Intelligence, AIJ* 32.1 (1987), pp. 57–95 (cit. on p. 70).

[San'72] Erik Sandewall. "An Approach to the Frame Problem, and its Implementation." In: *Machine Intelligence* 7.195-204 (1972), pp. 11–19 (cit. on p. 2).

[Sch'91a] Klaus Schild. "A Correspondence Theory for Terminological Logics: Preliminary Report." In: *Proceedings of the 12th International Joint Conference on Artificial Intelligence, IJCAI*. 1991, pp. 466–471 (cit. on pp. 22, 154).

[Sch'91b] Klaus Schild. "From Terminological Logics to Modal Logics." In: *Proceedings of the International Workshop on Terminological Logics*. 1991, pp. 101–104 (cit. on p. 21).

[SS'91] Manfred Schmidt-Schauß and Gert Smolka. "Attributive concept descriptions with complements." In: *Artificial Intelligence, AIJ* 48.1 (1991), pp. 1–26 (cit. on p. 13).

[SL'83] James G. Schmolze and Thomas A. Lipkis. "Classification in the KL-ONE Knowledge Representation System." In: *Proceedings of the 8th International Joint Conference on Artificial Intelligence, IJCAI*. 1983, pp. 330–332 (cit. on p. 23).

[Spa'00] Kent Spackman. "Managing clinical terminology hierarchies using algorithmic calculation of subsumption: Experience with SNOMED-RT." In: *Journal of the American Medical Informatics Association* (2000) (cit. on pp. 4, 6).

[TV'57] Alfred Tarski and Robert L. Vaught. "Arithmetical Extensions of Relational Systems." In: *Compositio Mathematica* 13 (1957), pp. 81–102 (cit. on p. 2).

[Tur'07] Anni-Yasmin Turhan. "On the Computation of Common Subsumers in Description Logics." PhD thesis. Dresden University of Technology, Germany, 2007. URL: `http://nbn-resolving.de/urn:nbn:de:bsz:14-ds-1212159862343-11479` (cit. on p. 23).

[Tur'50] Alan M. Turing. "Computing Machinery and Intelligence." In: *Mind* 59.236 (1950), pp. 433–460 (cit. on p. 29).

[Var'18] Ivan J. Varzinczak. "A Note on a Description Logic of Concept and Role Typicality for Defeasible Reasoning Over Ontologies." In: *Logica Universalis* 12.3-4 (2018), pp. 297–325 (cit. on pp. 39, 155).