

Towards a Knowledge Compilation Map for Heterogeneous Representation Languages*

Hélène Fargier
IRIT-CNRS, Univ. Paul Sabatier
Toulouse, France
fargier@irit.fr

Pierre Marquis and Alexandre Niveau
CRIL-CNRS, Univ. Artois
Lens, France
{marquis,niveau}@cril.fr

Abstract

The knowledge compilation map introduced by Darwiche and Marquis takes advantage of a number of concepts (mainly queries, transformations, expressiveness, and succinctness) to compare the relative adequacy of representation languages to some AI problems. However, the framework is limited to the comparison of languages that are interpreted in a homogeneous way (formulae are interpreted as Boolean functions). This prevents one from comparing, on a formal basis, languages that are close in essence, such as OBDD, MDD, and ADD. To fill the gap, we present a generalized framework into which comparing formally heterogeneous representation languages becomes feasible. In particular, we explain how the key notions of queries and transformations, expressiveness, and succinctness can be lifted to the generalized setting.

1 Introduction

There exist myriads of representation languages, with different capabilities; each one fits some applications, but is not suitable for others—there can be no best language. Choosing a good language for a given application is thus a fundamental problem in AI. Levesque and Brachman [1985] have shown that this problem boils down to a compromise between (computational) efficiency and expressiveness; for languages of equal expressiveness, efficiency must be balanced against succinctness [Gogic *et al.*, 1995]. The knowledge compilation map [Darwiche and Marquis, 2002] relies on these aspects to compare languages representing Boolean functions.

However, there are few works about the comparison of “heterogeneous” representation languages, based on different interpretation domains. It is well-known that ordered multi-valued decision diagrams (OMDDs) [Srinivasan *et al.*, 1990] can be transformed in polynomial time into ordered binary decision diagrams (OBDDs) [Bryant, 1986], using for example a log encoding (see Figure 1); these languages are considered somewhat equivalent. Yet, because of the specificity

of existing frameworks, this “equivalence” cannot be formally stated, and its use in proofs requires much prudence—we would like to use it as an elementary brick to infer results in an automated fashion. It can be noted that when the notion of a “representation language” is used in a broad sense, e.g., in surveys of the knowledge representation or artificial intelligence domains [Brachman and Levesque, 2004; van Harmelen *et al.*, 2008; Russell and Norvig, 2010], it is only described informally; and whenever there is a need for a formal definition, its scope is restricted to the case in point.

In this paper, we propose a generalized formal framework for representation languages, including a definition of equivalence suitable for heterogeneous languages, with the long-term purpose of casting them in a single knowledge compilation map. We prove that several properties usually expected in specific maps still hold in our general view, modulo certain modifications. We begin with a formal definition of a “representation language” and of important related notions, such as completeness and sublanguage, in Section 2. Then, we show how the usual concepts used for language comparison in the knowledge compilation map can be generalized to heterogeneous languages: Section 3 deals with notions related to representation efficiency (expressiveness, succinctness, polynomial translation), and Section 4 with notions related to computational power (notably queries and transformations). Section 5 concludes the paper, and illustrates a practical use of our framework on a case study (that of MDDs and BDDs).¹

2 Representation Language

2.1 Definitions

The knowledge compilation (KC) map is built upon the notion of *language*; it was not formally defined in the original map [Darwiche and Marquis, 2002], because it implicitly had the classical meaning of a *formal* language, i.e., a set of words over an alphabet. This “language” terminology only refers to syntax; semantics is given by an implicit *interpretation function* (that of propositional formulae). There is also a natural hierarchy on languages, namely *inclusion*: a *sublanguage*, or *fragment*, is simply a subset of a language.

In this classical setting, most features of languages are thus *implicit*, including their interpretation function and their

*This work is partially supported by the project BR4CP ANR-11-BS02-008 of the French National Agency for Research.

¹Proofs are omitted, but an extended version is available at <url: http://www.cril.univ-artois.fr/~niveau/publis/IJCAI13_FMN.pdf>.




Figure 1: An ordered interval automaton (OIA), an ordered multivalued decision diagram (OMDD), and an ordered binary decision diagram (OBDD). The OMDD is a discretization of the OIA, using partitions $\{[0, 3.45], (3.45, 5.1], [5.1, 8], (8, 10]\}$ for x and $\{[0, 2), [2, 4.6], (4.6, 7.7], (7.7, 10]\}$ for y . The OBDD is obtained from the OMDD by log encoding.

hierarchy, but also the notion of completeness and the domains of variables. This crucial limitation prevents one from directly applying this framework to a more general setting. Consider the bit sequence 10101010: interpreted as a binary-encoded natural integer, it corresponds to 170; as the “sign-and-magnitude” encoding of a relative integer, to -42 ; as a fixed-point real number, to 10.625; etc. These languages have the same syntax, but different interpretation functions. Our definition of a representation language builds upon that of Fargier and Marquis [2009], in that we use an explicit semantics; but we relax their assumption that it is restricted to the representation of Boolean functions, and consider languages that can represent potentially anything. For that purpose, let $\mathcal{U}$ be our universe of discourse, containing at least all “objects” that our languages aim at representing, notably real and natural numbers, Boolean functions,² etc. We also use a generic alphabet Σ , that we suppose countably infinite; we call *formula* a word over this alphabet, i.e., an element of Σ^* .³ Now, at the highest level of abstraction, a representation language is a *relation* linking formulae and objects in the universe.

Definition 2.1 (Representation language). A *representation language* is an ordered pair $L = \langle \Phi_L, \mathcal{I}_L \rangle$, where $\Phi_L \subseteq \Sigma^*$ is a set of formulae called the *syntax* of L , and $\mathcal{I}_L$, called the *semantics* of L , is a many-to-one relation $\mathcal{I}_L \subseteq \Sigma^* \times \mathcal{U}$ that is defined at least on all formulae in Φ_L .

The semantics of a language is a way of interpreting the symbols in Σ ; it can be seen as a function⁴ from Φ_L to $\mathcal{U}$, mapping any formula φ in the syntax of L (called an *L-representation*) to its interpretation $\llbracket \varphi \rrbracket_L$. For example, the language of propositional logic can be defined as $\text{PROP} = \langle \Phi_{\text{PROP}}, \mathcal{I}_{\text{PROP}} \rangle$, with Φ_{PROP} the set of well-formed formulae (with connectives restricted to $\neg$, $\vee$, and $\wedge$) and $\mathcal{I}_{\text{PROP}}$ the usual inductive interpretation function, associating with each formula φ its corresponding Boolean function $\llbracket \varphi \rrbracket_{\text{PROP}} \in \mathcal{B}$.

²We denote as $\mathcal{B}$ the set of Boolean functions of Boolean variables, and as $\mathcal{B}_{\mathbb{N}}$ and $\mathcal{B}_{\mathbb{R}}$ the sets of Boolean functions of variables having respectively an integer and a real domain.

³We use the Kleene star to denote unbounded Cartesian product: $S^* = \bigcup_{i \in \mathbb{N}} S^i$ (we consider words as tuples of symbols from Σ).

⁴In Def. 2.1, we use a many-to-one relation instead of a partial function to simplify later notation. The many-to-one requirement is not necessary (ambiguous semantics can be useful, a good example being natural languages) but we adopt it for the sake of simplicity.

This semantics is used in lots of languages, such as $\text{CNF} = \langle \Phi_{\text{CNF}}, \mathcal{I}_{\text{PROP}} \rangle$ (with Φ_{CNF} the set of formulae in conjunctive normal form) or $\text{HORN-C} = \langle \Phi_{\text{HORN-C}}, \mathcal{I}_{\text{PROP}} \rangle$, (with $\Phi_{\text{HORN-C}}$ the set of Horn-CNFs). Other examples include BDD, MDD, and IA, the languages of binary and multivalued decision diagrams and of interval automata [Niveau *et al.*, 2010] (see Figure 1), and the language of algebraic decision diagrams, ADD [Bahar *et al.*, 1997]. For space reasons, the paper is focused on these languages (in particular, the relationship between the MDD and BDD families is presented in Section 5), but the framework is much more general.⁵

Remark that there can be elements $\omega \in \mathcal{U}$ with which $\mathcal{I}_L$ does not associate *any* formula, neither inside Φ_L nor outside. Those elements are thus completely unrelated to the language; for example, $\mathcal{I}_{\text{PROP}}$ ignores any object that is not a Boolean function. The elements of $\mathcal{U}$ with which $\mathcal{I}_L$ associates at least one formula are called *L-interpretations*.

Definition 2.2 (Interpretation space). The *interpretation space* of a representation language L , denoted Ω_L , is the codomain of $\mathcal{I}_L$, i.e., $\Omega_L = \{ \omega \in \mathcal{U} \mid \exists \varphi \in \Sigma^*, \varphi \mathcal{I}_L \omega \}$.

Since CNF and HORN-C have the same semantics, they also have the same interpretation space, which is $\mathcal{B}$. Note that elements in the interpretation space of a language L are not guaranteed to have a *representation* in Φ_L ; e.g., the syntax of HORN-C is *not expressive enough* to cover the whole interpretive power of its semantics—the language is *incomplete*.

Definition 2.3 (Completeness). A representation language L is *complete* if and only if $\forall \omega \in \Omega_L, \exists \varphi \in \Phi_L, \llbracket \varphi \rrbracket_L = \omega$.

2.2 Sublanguages

A representation language is not simply a set of formulae, therefore the notion of *sublanguage* cannot simply be based on a restriction of syntax, as it is in the classical KC map.

Definition 2.4 (Sublanguage). Let L and L' be two representation languages; L' is a *sublanguage* of L , denoted $L' \subseteq L$, if and only if $\Phi_{L'} \subseteq \Phi_L$ and $\mathcal{I}_{L'} \subseteq \mathcal{I}_L$. Moreover, if $L' \subseteq L$ and $\mathcal{I}_{L'} = \mathcal{I}_L$, L' is said to be a *fragment* of L .

⁵It notably includes VNNF [Fargier and Marquis, 2007], and languages outside the usual scope of KC, like all flavors of constraint networks (discrete or continuous, weighted, fuzzy, etc.) [see, e.g., Rossi *et al.*, 2006], quadrees [Finkel and Bentley, 1974], R^* -trees [Beckmann *et al.*, 1990] and, with appropriate interpretation spaces, first-order logic, ontologies, programming languages, etc.

The condition that all formulæ in the sublanguage must respect the syntax of the parent language remains; but they must also respect its semantics (which was implicit in the classical framework). Note that our definition of a representation language distinguishes two kinds of sublanguage hierarchies, that of *fragments*, which orders languages of a fixed semantics, and that of the generalized *sublanguages*, in which even heterogeneous languages can be compared: e.g., it holds that $\text{OBDD} \subseteq \text{OMDD} \subseteq \text{OIA}$, since OBDDs are specific OMDDs (restricted to Boolean variables), and OMDDs are specific OIAs (restricted to discrete variables). This natural property cannot be formally stated within the classical KC map.

It is interesting to notice that contrary to what is usually expected, incompleteness is not inherited by sublanguages. Thus, OIA is not complete (edge labels are restricted to closed intervals), but OBDD is, although $\text{OBDD} \subseteq \text{OIA}$. This comes from the fact that completeness is relative to the interpretation space of the language, which is not necessarily the same in a sublanguage. The expected property actually holds on fragments, which have stronger requirements.

Proposition 2.5. *Let L be a representation language, and L' a fragment of L . If L' is complete, then L also is.*

2.3 Operations on Languages

In the classical KC map, the hierarchy of fragments is induced by a set of syntactic properties verified by languages: considering only read-once BDDs yields FBDD, adding ordering yields OBDD, etc. These can be seen as *syntactic restrictions*; they can reduce expressiveness, since the interpretation space does not change. On the other hand, the hierarchy of decision diagrams in the broad sense ($\text{OBDD} \subseteq \text{OMDD}$) cannot be generated by only restricting syntax, since some interpretations must be removed: it can be built using *semantic restrictions*.

Definition 2.6 (Language restrictions). The (*syntactic*) restriction of a language L to a set of formulæ $\Phi \subseteq \Sigma^*$ is the fragment $L|_\Phi = \langle \Phi_L \cap \Phi, \mathcal{I}_L \rangle$.

The (*semantic*) restriction of a language L to an interpretation space $\Omega \subseteq \mathcal{U}$ is the sublanguage $L|^\Omega = \langle \{ \varphi \in \Phi_L \mid \llbracket \varphi \rrbracket_L \in \Omega \}, \mathcal{I}_L \cap (\Sigma^* \times \Omega) \rangle$.

Using this definition, HORN-C is the syntactic restriction of CNF to Horn-CNFs, for example, while BDD is the semantic restriction of MDD to functions of Boolean variables. Another way of building a language is to *combine* existing languages.

Definition 2.7 (Union and intersection). The *union* of two languages L and L' is $L \cup L' = \langle \Phi_L \cup \Phi_{L'}, \mathcal{I} \cup \mathcal{I}' \rangle$, and their *intersection* is $L \cap L' = \langle \Phi_L \cap \Phi_{L'}, \mathcal{I} \cap \mathcal{I}' \rangle$.

This allows for example to formally state that $\text{BDD} = \text{MDD} \cap \text{ADD}$, which expresses the fact that BDDs are both exactly the MDDs on Boolean variables, and the ADDs with two leaves.

We are now ready to define the usual concepts of the KC map: expressiveness, succinctness, and polynomial translatability, and support of queries and transformations.

3 Representation Efficiency

In the KC map, expressiveness has a fundamental role: a difference in succinctness between two languages is only significant if they have the same expressiveness. Languages with

different interpretation spaces are *a fortiori* not compared in the map. Yet, there are works about translations between heterogeneous languages [Walsh, 2000; Gottlob, 1995], and it would be interesting to cast known results in the map. This calls for a more general definition of expressiveness, succinctness, and polynomial translatability, that would not require a strict equality of interpretations, but only their equivalence modulo some *semantic correspondence*, i.e., some relation linking the two heterogeneous interpretation spaces—which induces a translation between formulæ.

Definition 3.1 (Semantic correspondence, translation). A *semantic correspondence* between two interpretation spaces Ω_1 and Ω_2 is a subset of $\Omega_1 \times \Omega_2$. We will call *translation from L_1 to L_2* a semantic correspondence $\mathcal{T} \subseteq \Omega_{L_1} \times \Omega_{L_2}$.

Examples of well-known translations include the various encodings of discrete constraint networks into propositional formulæ [Walsh, 2000], such as the direct encoding $\mathcal{T}_{\text{dir}}$ and the log encoding $\mathcal{T}_{\text{log}}$, which associate functions in $\mathfrak{B}_{\mathbb{N}}$ with functions in $\mathfrak{B}$; or the discretization translation $\mathcal{T}_{\text{discr}}$, which relates a continuous constraint network with the discrete constraint networks it can be discretized into. We also use the generic identity translation Id , which denotes any relation $\{ \langle \omega, \omega \rangle : \omega \in \Omega \}$ with $\Omega \subseteq \mathcal{U}$, in a slight abuse of notation.

3.1 Expressiveness and Succinctness

A first criterion for the comparison of representation languages is their relative expressiveness [Gogic *et al.*, 1995; Fargier and Marquis, 2008]. When they have the same semantics, expressiveness compares their ability to represent objects. However, we want to compare languages of different semantics, therefore we will define expressiveness as depending on a semantic correspondence, and similarly extend relative succinctness [Gogic *et al.*, 1995; Darwiche and Marquis, 2002], which compares the ability of languages to represent objects *compactly*.

Definition 3.2 (Expressiveness, succinctness). Let L_1 and L_2 be two languages, and $\mathcal{T}$ a translation from L_1 to L_2 .

L_2 is *at least as expressive* as L_1 modulo $\mathcal{T}$, denoted $L_1 \geq_e^{\mathcal{T}} L_2$, if and only if for each L_1 -representation φ_1 , there exists an L_2 -representation φ_2 such that $\llbracket \varphi_1 \rrbracket_{L_1} \mathcal{T} \llbracket \varphi_2 \rrbracket_{L_2}$.

L_2 is *at least as succinct* as L_1 modulo $\mathcal{T}$, denoted $L_1 \geq_s^{\mathcal{T}} L_2$, if and only if there exists a polynomial $P(\cdot)$ such that for each L_1 -representation φ_1 , there exists an L_2 -representation φ_2 such that $|\varphi_2| \leq P(|\varphi_1|)$ and $\llbracket \varphi_1 \rrbracket_{L_1} \mathcal{T} \llbracket \varphi_2 \rrbracket_{L_2}$.

Comparable expressiveness is necessary for a translation to be well-adapted to the comparison of two languages. It is clearly not a sufficient condition: for any two languages L_1 and L_2 , the trivial translation $\mathcal{T} = \Omega_{L_1} \times \Omega_{L_2}$ verifies $L_1 \geq_e^{\mathcal{T}} L_2$, yet it does not allow one to infer any result about L_1 and L_2 . Sufficient conditions will be studied in Section 4.

As in the classical KC framework, succinctness is a refinement of expressiveness: $L \geq_s^{\mathcal{T}} L' \implies L \geq_e^{\mathcal{T}} L'$. How-

⁶We also use (i) $L_1 \leq_e^{\mathcal{T}} L_2$ to mean $L_2 \geq_e^{\mathcal{T}^{-1}} L_1$, (ii) $L_1 \sim_e^{\mathcal{T}} L_2$ to mean that both $L_1 \geq_e^{\mathcal{T}} L_2$ and $L_1 \leq_e^{\mathcal{T}} L_2$, and (iii) $L_1 >_e^{\mathcal{T}} L_2$ to mean that $L_1 \geq_e^{\mathcal{T}} L_2$ but $L_1 \not\leq_e^{\mathcal{T}} L_2$. This notation applies, *mutatis mutandis*, to succinctness and polynomial translatability.

ever, while this implies that the classical succinctness relation, which corresponds to $\geq_s^{\text{Id}}$, is not very informative when applied to heterogeneous languages, our generalization does not suffer from this drawback: it is possible to formally state succinctness results like $\text{MDD} \not\geq_s^{\mathcal{T}_{\text{dir}}} \text{CNF}$. Note that we can apply the KC notation to any representation language: for example, denoting RADIX_n the language of natural integers represented in radix n , it is well-known that $\text{RADIX}_1 \geq_s^{\text{Id}} \text{RADIX}_2$, and that for any $n > 1$, $\text{RADIX}_n \sim_s^{\text{Id}} \text{RADIX}_2$.

Note also that specific translations can be used to compare the succinctness of languages having the same interpretation space but different expressiveness. For example, languages HORN-C and $\text{KROM-C} = \langle \Phi_{\text{KROM-C}}, \mathcal{I}_{\text{PROP}} \rangle$, where $\Phi_{\text{KROM-C}}$ is the set of 2-CNF formulae, are incomparable with respect to $\geq_e^{\text{Id}}$, therefore they are also incomparable with respect to classical succinctness $\geq_s^{\text{Id}}$, however, it can be interesting to know which one is the most succinct when restricted to those Boolean functions that both can represent. It is possible to express this within the generalized framework, defining an ad-hoc translation $\mathcal{T}$, for which $\text{HORN-C} \geq_s^{\mathcal{T}} \text{KROM-C}$ holds if and only if every HORN-C -representation is either not representable in KROM-C , or representable in polynomial size.⁷

3.2 Polynomial Translatability

Succinctness requires the existence of a polynomial-size translation, but not the existence of an algorithm building it, let alone its tractability. The last refinement of expressiveness is polynomial translatability [Fargier and Marquis, 2009].

Definition 3.3 (Polynomial translatability). Let L_1 and L_2 be two representation languages, and $\mathcal{T}$ a translation from L_1 to L_2 ; L_1 is *polynomially translatable* into L_2 modulo $\mathcal{T}$, denoted $L_1 \geq_p^{\mathcal{T}} L_2$, if and only if there exists a polynomial-time algorithm mapping any L_1 -representation ϕ_1 to an L_2 -representation ϕ_2 such that $\llbracket \phi_1 \rrbracket_{L_1} \mathcal{T} \llbracket \phi_2 \rrbracket_{L_2}$. Moreover, if the output is guaranteed to be *at most* polynomially *smaller* than the input, that is, if there is a polynomial $P(\cdot)$ such that for any ϕ_1 , the output ϕ_2 verifies $|\phi_2| \leq P(|\phi_1|)$, then the translation is said to be *stable*, and we denote $L_1 \geq_p^{\mathcal{T}} L_2$.

It holds, for example, that $\text{OMDD} \geq_p^{\mathcal{T}_{\text{dir}}} \text{OBDD}$ and $\text{OMDD} \geq_p^{\mathcal{T}_{\log}} \text{OBDD}$ [Srinivasan *et al.*, 1990] (e.g., in Figure 1, the OBDD results from the log encoding of the OMDD). The stability condition is generally not hard to achieve; in particular, it discards special cases in which the translated formula is exponentially smaller than the original one.

In the classical KC framework, polynomial translation corresponds to $\geq_p^{\text{Id}}$, and has important consequences on the satisfaction of queries and transformations. For example, denoting MODS the restriction of PROP to smooth and deterministic DNFs [Darwiche and Marquis, 2002], the fact that $\text{MODS} \geq_p^{\text{Id}} \text{OBDD}$ (i.e., MODS -representations can be transformed into equivalent OBDDs in polynomial time) implies that MODS supports all queries supported by OBDD; and since $\text{NNF} \sim_p^{\text{Id}} \text{PROP}$ (propositional formulae with connectives limited to $\neg$, $\vee$, and $\wedge$ can be put in negation normal form in polynomial time), NNF and PROP support the exact same set

⁷One can define $\mathcal{T}$ as the set of pairs of Boolean functions $\langle f, g \rangle \in \mathcal{B}^2$ for which if g has a KROM-C -representation, then $f = g$.

of queries and transformations. In the latter case, the two languages are hence considered as strictly equivalent. However, these properties do not hold when a translation is used, because queries and transformations applying to one language do not necessarily apply to the other. For example, OIAs can be discretized into OMDDs (e.g., in Figure 1, the OMDD is a discretization of the OIA), i.e., $\text{OIA} \geq_p^{\mathcal{T}_{\text{discr}}} \text{OMDD}$, but while OMDD supports model enumeration, OIAs generally have an infinite number of models. Conditions under which this kind of inference is possible are studied in Section 4.

3.3 Properties of Comparison Relations

Let $\geq$ denote any of the three relations we have defined, $\geq_e$, $\geq_s$, or $\geq_p$. Using the Id translation, we recover the original definition of each relation; $\geq^{\text{Id}}$ is suitable only to homogeneous languages, but has the advantage of being a preorder. It is not the case for any $\mathcal{T}$, simply because interpretation spaces can be different (in which case $\geq^{\mathcal{T}}$ notably cannot be reflexive). However, when $\mathcal{T}$ is an endorelation, $\geq^{\mathcal{T}}$ inherits some of its properties.

Proposition 3.4. Let $\Omega \subseteq \mathcal{U}$, and $\mathcal{T} \subseteq \Omega^2$. If $\mathcal{T}$ is reflexive (resp. transitive), then $\geq^{\mathcal{T}}$ is also reflexive (resp. transitive).

Such translations can be used to compare languages over the same interpretation space, as shown with HORN-C and KROM-C . When $\mathcal{T}$ has no remarkable property, the next proposition nevertheless expresses a kind of “pseudo-transitivity”.

Proposition 3.5. If, for some representation languages L_1 , L_2 , and L_3 , and some semantic correspondences $\mathcal{T}$ and $\mathcal{T}'$, it holds that $L_1 \geq^{\mathcal{T}} L_2$ and $L_2 \geq^{\mathcal{T}'} L_3$, then $L_1 \geq^{\mathcal{T} \circ \mathcal{T}'} L_3$ (where $\circ$ denotes the composition of relations).

This proposition has a useful corollary, allowing succinctness results in a given fragment hierarchy to be extended to another, granted that a “reversible” polynomial translation exists between them (see Section 5 for an example).

Corollary 3.6. Let L_1 and L_2 (resp. L'_1 and L'_2) be two languages of interpretation space Ω (resp. Ω'). If there exists a bijective correspondence $\mathcal{T}$ between Ω and Ω' such that $L_1 \leq^{\mathcal{T}} L'_1$ and $L_2 \geq^{\mathcal{T}} L'_2$, then $L_1 \geq^{\text{Id}} L_2 \implies L'_1 \geq^{\text{Id}} L'_2$.

4 Computational Power

Representation languages are used to solve problems about the objects they represent. A solution to a problem is generally found thanks to an algorithm, that is, a sequence of operations on the objects; the practical implementation of the algorithm depends on the chosen representation languages.

4.1 Operations

Definition 4.1 (Semantic operation). A *semantic operation* on a universe $\Omega \subseteq \mathcal{U}$ is a piecewise total⁸ relation $\rho \subseteq \Omega^{\kappa} \times \mathcal{P} \times \mathcal{A}$, where $\mathcal{P} \subseteq \mathcal{U}$ and $\mathcal{A} \subseteq \mathcal{U}$ are sets containing *parameters* and *answers* respectively, and either $\kappa \in \mathbb{N}$ (the semantic operation is then *bounded*), or $\kappa = *$ (the semantic operation is then *unbounded*).

⁸A relation $\mathcal{R} \subseteq D_1 \times \dots \times D_n$ is piecewise total if and only if $\forall i \in \{1, \dots, n\}, \forall d_i \in D_i, \exists (r_1, \dots, r_n) \in \mathcal{R}, r_i = d_i$. Less formally, this means that its projection on each of its dimensions is total.

Let us give examples based on the KC map, with Ω the set of Boolean functions of real variables: $\Omega = \mathcal{B}_{\mathbb{R}}$. The operation associated with the “conditioning” transformation is the partial function ρ_{CD} from $\Omega \times (\mathbb{R})^*$ to Ω , defined as $\rho_{\text{CD}}: \langle f, \vec{x} \rangle \mapsto f|_{\vec{x}}$ (with ρ_{CD} only defined on pairs $\langle f, \vec{x} \rangle$ such that f is defined on $\vec{x}$). Here $\kappa = 1$, $\mathcal{P} = (\mathbb{R})^*$ is the set of all assignments, and $\mathcal{A} = \Omega$ is the set of Boolean functions. The operation associated with the “model extraction” query is not a function, but a relation ρ_{MX} associating any Boolean function f with *all* of its models; here $\kappa = 1$, $\mathcal{A} = (\mathbb{R})^*$, and $\mathcal{P}$ is ignored.⁹ The operation associated with the “conjunction” transformation is the mapping $\rho_{\wedge \text{C}}: \langle f_1, \dots, f_{n_\kappa} \rangle \mapsto \bigwedge_{i=1}^{n_\kappa} f_i$ for any $n_\kappa \in \mathbb{N}$. Here, the operation is unbounded: $\kappa = *$.

The complexity of a semantic operation $\rho \subseteq \Omega^\kappa \times \mathcal{P} \times \mathcal{A}$ depends on the representation languages considered. The most important one is the “input” language, that is, the language representing elements of Ω . Semantic operations can be much more generic than needed: the operations presented above apply to all Boolean functions of real variables, but they can also be used, in a more restricted fashion, to compare Boolean functions of Boolean variables. A semantic operation can indeed be *applied* to a subset Ω' of its universe: we denote $\rho|_{\Omega'}$ the largest (with respect to inclusion) semantic operation verifying $\rho|_{\Omega'} \subseteq \rho \cap (\Omega'^\kappa \times \mathcal{P} \times \mathcal{A})$, and $\mathcal{P}_{\Omega'}$ and $\mathcal{A}_{\Omega'}$ the sets of parameters and answers of $\rho|_{\Omega'}$, from which all “superfluous” elements have been removed. We will say that a semantic operation on $\Omega \subseteq \mathcal{U}$ is *applicable* to a language L when $\Omega_L \subseteq \Omega$ holds: for example, ρ_{CD} is applicable to CNF, and then $\mathcal{P}_{\Omega_{\text{CNF}}}$ is the set of Boolean assignments and $\mathcal{A}_{\Omega_{\text{CNF}}}$ the set of Boolean functions of Boolean variables. We use this notion of application to define the syntactic operations that are associated with semantic operations.

Definition 4.2 (Operation). A (*syntactic*) operation is a tuple $\mathcal{O} = \langle \rho, L, \text{PRM}, \text{ANS} \rangle$ such that (i) ρ is a semantic operation: $\rho \subseteq \Omega^\kappa \times \mathcal{P} \times \mathcal{A}$; (ii) L is a representation language to which ρ is applicable: $\Omega_L \subseteq \Omega$; (iii) PRM is a representation language covering all parameters compatible with L : $\mathcal{P}_{\Omega_L} \subseteq \Omega_{\text{PRM}}$; (iv) ANS is a representation language covering all answers compatible with L : $\mathcal{A}_{\Omega_L} \subseteq \Omega_{\text{ANS}}$.

Note how the choice of the main representation language in the syntactic operation restricts the semantic operation. When choosing L , we discard all elements in the domain of ρ that are not L -interpretations. Languages in an operation induce a syntactic version of its underlying semantic operation.

Definition 4.3 (Complexity of an operation). Let $\mathcal{O} = \langle \rho, L, \text{PRM}, \text{ANS} \rangle$ be a syntactic operation, with $\rho \subseteq \Omega^\kappa \times \mathcal{P} \times \mathcal{A}$. The problem associated with $\mathcal{O}$ is the following: for any tuple $\langle \varphi_1, \dots, \varphi_{n_\kappa}, \pi \rangle$ of formulae from $\Phi_L^\kappa \times \Phi_{\text{PRM}}$, compute a formula $\alpha \in \Phi_{\text{ANS}}$ such that $\langle \llbracket \varphi_1 \rrbracket_L, \dots, \llbracket \varphi_{n_\kappa} \rrbracket_L, \llbracket \pi \rrbracket_{\text{PRM}}, \llbracket \alpha \rrbracket_{\text{ANS}} \rangle \in \rho$, if there exists one. The complexity of $\mathcal{O}$ is that of its associated problem.

It is the complexity of this syntactic problem that defines the complexity of the corresponding operation. Thus, “operation $\mathcal{O}$ is in polynomial time” means that there exists a

⁹When a semantic operation needs no parameter, we omit $\mathcal{P}$ for simplicity (we implicitly consider that it is a singleton).

polynomial-time algorithm that solves the problem associated with $\mathcal{O}$. An operation is hence a way of specifying a problem; classical *decision problems* in computability theory constitute a special case of operations. A decision problem can be modeled as $\mathcal{O} = \langle \rho_\Phi, \Sigma^*, \text{bool} \rangle$, where ρ_Φ is the indicative function of a formal language Φ , and bool the representation language associating the symbol “0” with $\perp$ and the symbol “1” with $\top$. More generally, an operation allows one to express a *function problem* without losing the semantics.

4.2 Queries and Transformations

We recover queries and transformations as special operations by fixing some elements, namely, languages and complexity.

Definition 4.4 (Query, transformation). A *query* (resp. a *transformation*) is a tuple $\mathbf{Q} = \langle \rho, \text{PRM}, \text{ANS} \rangle$ (resp. $\mathbf{T} = \langle \rho, \text{PRM} \rangle$) such that for any representation language L to which ρ is applicable, $\mathcal{O}_{\mathbf{Q},L} = \langle \rho, L, \text{PRM}, \text{ANS} \rangle$ (resp. $\mathcal{O}_{\mathbf{T},L} = \langle \rho, L, \text{PRM}, L \rangle$) is an operation. Language L *supports* $\mathbf{Q}$ (resp. $\mathbf{T}$) if and only if $\mathcal{O}_{\mathbf{Q},L}$ is in input-output polynomial time (resp. $\mathcal{O}_{\mathbf{T},L}$ is in polynomial time).

For languages representing Boolean functions, denoting term the language of conjunctions of atoms of the form $[x = n]$ (where $x \in \Sigma$ is a variable and $n \in \mathbb{R}$), the model extraction query can be defined as $\mathbf{MX} = \langle \rho_{\text{MX}}, \text{term} \rangle$, the conditioning transformation can be defined as $\mathbf{CD} = \langle \rho_{\text{CD}}, \text{term} \rangle$, and the conjunction transformation as $\mathbf{\wedge C} = \langle \rho_{\wedge \text{C}} \rangle$. These operations apply to any language representing Boolean functions, be it restricted or not to a subset of this space, e.g., to functions of Boolean or integer variables.

Support of given queries and transformations is an indicator of the absolute and relative computational power of representation languages; with our definition, the concept applies universally to any representation language. However, it should be clear that the actual list of queries and transformations useful to compare a given hierarchy of languages generally *depends* on their interpretation space. Nevertheless, it is sometimes possible to infer results about queries and transformations on a hierarchy from results on another hierarchy; e.g., support of most queries and transformations for languages in the MDD family can be inferred from results on the BDD family, because these families are polynomially equivalent modulo a “suitable” translation. We will now present sufficient conditions for this kind of inference to be possible.

4.3 Operations and Polynomial Translation

In the general case, we are interested in deducing the tractability of an operation $\mathcal{O}$ from the tractability of another, into which $\mathcal{O}$ can be translated. We start by defining the notion of translation between two semantic operations.

Definition 4.5 (Semantic operation translation). Let ρ_1 and ρ_2 be two semantic operations: $\rho_1 \subseteq \Omega_1^\kappa \times \mathcal{P}_1 \times \mathcal{A}_1$ and $\rho_2 \subseteq \Omega_2^\kappa \times \mathcal{P}_2 \times \mathcal{A}_2$. A *semantic operation translation* from ρ_1 to ρ_2 is a triple $\langle T_\Omega, T_\mathcal{P}, T_\mathcal{A} \rangle$, where (i) $T_\Omega \subseteq \Omega_1 \times \Omega_2$, (ii) $T_\mathcal{P} \subseteq \mathcal{P}_1 \times \mathcal{P}_2$, and (iii) $T_\mathcal{A} \subseteq \mathcal{A}_1 \times \mathcal{A}_2$, that verifies $\rho_1 = T_\mathcal{A}^{-1} \circ \rho_2 \circ (T_\Omega^\kappa \cdot T_\mathcal{P})$, where $\circ$ (resp. $\cdot$) denotes the composition (resp. product) of relations.

We begin to see why some translations are more useful than others: for example, the trivial semantic correspondence

$\mathcal{T} = \Omega_1 \times \Omega_2$ is not likely to be used in a semantic operation translation. The interest of having $L_1 \geq_p^{\mathcal{T}} L_2$ actually depends on how this specific $\mathcal{T}$ relates to the operations considered. Moreover, nothing can be deduced if the translation of parameters and answers is not tractable: we need the notion of polynomial translation between syntactic operations.

Definition 4.6 (Polynomial translation between operations). Let $\mathcal{O}_1 = \langle \rho_1, L_1, \text{PRM}_1, \text{ANS}_1 \rangle$ and $\mathcal{O}_2 = \langle \rho_2, L_2, \text{PRM}_2, \text{ANS}_2 \rangle$ be two operations. We say that $\mathcal{O}_1$ is *polynomially translatable* into $\mathcal{O}_2$ if and only if there exists a semantic operation translation $\langle \mathcal{T}_\Omega, \mathcal{T}_\rho, \mathcal{T}_\mathcal{A} \rangle$ from ρ_1 to ρ_2 verifying (i) $L_1 \geq_p^{\mathcal{T}_\Omega} L_2$; (ii) $\text{PRM}_1 \geq_p^{\mathcal{T}_\rho} \text{PRM}_2$; and (iii) $\text{ANS}_1 \leq_p^{\mathcal{T}_\mathcal{A}} \text{ANS}_2$. If the polynomial translation between answer languages is stable, that is, if $\text{ANS}_1 \leq_p^{\mathcal{T}_\mathcal{A}} \text{ANS}_2$, then the translation between operations is said to be *answer-stable*.

This notion encompasses the familiar polynomial many-one reduction of computational complexity theory: a problem, i.e., an operation $\mathcal{O}_1 = \langle \rho_{\Phi_1}, \Sigma_1^*, \text{bool} \rangle$, is polynomial-time many-one reducible to another, i.e., to an operation $\mathcal{O}_2 = \langle \rho_{\Phi_2}, \Sigma_2^*, \text{bool} \rangle$, if and only if $\mathcal{O}_1$ is polynomially translatable into $\mathcal{O}_2$. In a fashion similar to reductions between problems, the tractability of an operation depends on that of operations into which it can be polynomially translated.

Theorem 4.7. *Let $\mathcal{O}_1$ and $\mathcal{O}_2$ be two operations such that $\mathcal{O}_1$ is polynomially translatable into $\mathcal{O}_2$. If $\mathcal{O}_2$ is in polynomial time, then $\mathcal{O}_1$ is in polynomial time. When the translation is answer-stable, it holds that if $\mathcal{O}_2$ is in input-output polynomial time, then $\mathcal{O}_1$ is in input-output polynomial time.*

By this theorem, $\langle \rho_{\text{MX}}, \text{MDD}, \text{term} \rangle$, $\langle \rho_{\text{CD}}, \text{MDD}, \text{term}, \text{MDD} \rangle$, and $\langle \rho_{\text{AC}}, \text{MDD}, \text{MDD} \rangle$ are tractable, relying on the tractability of the corresponding operations on languages of the BDD family. However, the fact that OBDD supports **SFO** (the *forgetting* of a single variable) does not imply that OMDD supports **SFO**, because forgetting a multivalued variable boils down to forgetting an unbounded number of Boolean variables, which is NP-hard on OBDD. Similarly, translations that do not maintain the number of models cannot be used to infer that OMDD supports **CT** from the fact that OBDD supports **CT**. Note that the condition of answer-stability is necessary for the second statement of the theorem to hold, because the time complexity of the overall procedure depends on the size of $\mathcal{O}_2$'s answer, which can be exponential in the input.

The deductions allowed by this theorem can pertain to very disparate operations. However, in the context of a KC map, the setting is generally more restricted: typically, one has found a semantic correspondence between two interpretation spaces, inducing a polynomial translation between languages in the two hierarchies, and would like to make deductions such as “if L supports this query, then L' also supports it”; i.e., one is interested in a *single* semantic operation applicable to both languages. This can actually be cast as a corollary of Theorem 4.7, as long as one considers only queries and transformations that are *suitable* to the given translation.

Definition 4.8 ($\mathcal{T}$ -suitability). Let $\mathcal{T}$ be a semantic correspondence between some $\Omega_1 \subseteq \mathcal{U}$ and some $\Omega_2 \subseteq \mathcal{U}$. A query $\mathbf{Q} = \langle \rho, \text{PRM}, \text{ANS} \rangle$ is $\mathcal{T}$ -*suitable* if and only if: (i) ρ is applicable to Ω_1 and Ω_2 ; (ii) there exist two seman-

tic correspondences $\mathcal{T}_\rho$ and $\mathcal{T}_\mathcal{A}$ such that $\langle \mathcal{T}, \mathcal{T}_\rho, \mathcal{T}_\mathcal{A} \rangle$ is a translation from $\rho|_{\Omega_1}$ to $\rho|_{\Omega_2}$; (iii) $\text{PRM}|_{\Omega_1} \geq_p^{\mathcal{T}_\rho} \text{PRM}|_{\Omega_2}$; (iv) $\text{ANS}|_{\Omega_1} \leq_p^{\mathcal{T}_\mathcal{A}} \text{ANS}|_{\Omega_2}$. A transformation $\mathbf{T} = \langle \rho, \text{PRM} \rangle$ is $\mathcal{T}$ -*suitable* if and only if it verifies conditions i–iii, with $\mathcal{T}_\mathcal{A} = \mathcal{T}$ (the last condition does not apply, since there is no ANS language in a transformation).

Most queries and transformations of the classical KC map are suitable to direct and log encoding, including **MX**, **CD**, **AC**, and those considered by Darwiche and Marquis [2002]—with the exception of **SFO**, for reasons stated earlier. This is particularly interesting with regard to the following result.

Theorem 4.9. *Let L_1 and L_2 be two representation languages, and $\mathcal{T}$ a translation from L_1 to L_2 . If $L_1 \geq_p^{\mathcal{T}} L_2$, then every $\mathcal{T}$ -suitable query supported by L_2 is also supported by L_1 . If $L_1 \sim_p^{\mathcal{T}} L_2$, then every $\mathcal{T}$ -suitable query or transformation supported by L_2 is also supported by L_1 .*

This allows to automatically extend most known results about BDDs to MDDs, and more generally, most results from the Boolean KC map to languages over non-Boolean variables. Note however that results are not exactly the same, as they depend on the suitability of each query and transformation to the translation considered; thus, OMDD does not support **SFO** [Amilhastre *et al.*, 2012], even though OBDD does.

5 Conclusion

In this paper, we have presented a framework for comparing representation languages. While taking as few hypotheses as possible about what constitutes an admissible representation language, we showed how the usual concepts of the knowledge compilation map could be adapted to this broader setting, allowing the comparison of heterogeneous languages with respect to their representation efficiency and their computational capabilities, and the exploitation of known results in the Boolean case to numerous languages, over discrete or continuous variables and non-Boolean valuation.

As an illustration of the latter point, let us take the simple example of the “bounded MDD” family: we define the k -MDD language as the restriction of MDD to variables with domains of cardinality k , and its fragments k -FMDD, k -OMDD, and k -OMDD_< as its restriction to read-once, ordered, and <-ordered diagrams [Darwiche and Marquis, 2002], respectively.

Proposition 5.1. *It holds that $k\text{-MDD} <_s k\text{-FMDD} <_s k\text{-OMDD} <_s k\text{-OMDD}_{<}$, and each of the queries and transformations considered by Darwiche and Marquis [2002] is supported by $k\text{-MDD}$ (resp. $k\text{-FMDD}$, $k\text{-OMDD}$, $k\text{-OMDD}_{<}$) if and only if it is supported by BDD (resp. FBDD, OBDD, OBDD_<).*

This proposition follows directly from the fact that $k\text{-MDD} \sim_p^{\mathcal{T}_k} \text{BDD}$, $k\text{-FMDD} \sim_p^{\mathcal{T}_k} \text{FBDD}$, $k\text{-OMDD} \sim_p^{\mathcal{T}_k} \text{OBDD}$, and $k\text{-OMDD}_{<} \sim_p^{\mathcal{T}_k} \text{OBDD}_{<}$, denoting $\mathcal{T}_k$ the restriction of $\mathcal{T}_{\text{dir}}$ to variables with domains of cardinality k , by applying Corollary 3.6 ($\mathcal{T}_k$ is bijective) and Theorem 4.9 (all queries and transformations defined by Darwiche and Marquis [2002] are $\mathcal{T}_k$ -suitable), respectively.

This is a first step towards a generalized knowledge compilation map, in which heterogeneous language hierarchies could be presented in a unified way.

References

- [Amilhastre *et al.*, 2012] Jérôme Amilhastre, Hélène Fargier, Alexandre Niveau, and Cédric Pralet. Compiling CSPs: A complexity map of (non-deterministic) multivalued decision diagrams. In *Proceedings of the International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 1–8, 2012.
- [Bahar *et al.*, 1997] R. Iris Bahar, Erica A. Frohm, Charles M. Gaona, Gary D. Hachtel, Enrico Macii, Abelardo Pardo, and Fabio Somenzi. Algebraic decision diagrams and their applications. *Formal Methods in System Design*, 10(2/3):171–206, 1997.
- [Beckmann *et al.*, 1990] Norbert Beckmann, Hans-Peter Kriegel, Ralf Schneider, and Bernhard Seeger. The R*-tree: An efficient and robust access method for points and rectangles. In *SIGMOD Conference*, pages 322–331, 1990.
- [Brachman and Levesque, 2004] Ronald J. Brachman and Hector J. Levesque. *Knowledge Representation and Reasoning*. Elsevier, 2004.
- [Bryant, 1986] Randall E. Bryant. Graph-based algorithms for Boolean function manipulation. *IEEE Transactions on Computers*, 35(8):677–691, 1986.
- [Darwiche and Marquis, 2002] Adnan Darwiche and Pierre Marquis. A knowledge compilation map. *Journal of Artificial Intelligence Research (JAIR)*, 17:229–264, 2002.
- [Fargier and Marquis, 2007] Hélène Fargier and Pierre Marquis. On valued negation normal form formulas. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 360–365, 2007.
- [Fargier and Marquis, 2008] Hélène Fargier and Pierre Marquis. Extending the knowledge compilation map: Krom, Horn, affine and beyond. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 442–447, 2008.
- [Fargier and Marquis, 2009] Hélène Fargier and Pierre Marquis. Knowledge compilation properties of trees-of-BDDs, revisited. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 772–777, 2009.
- [Finkel and Bentley, 1974] Raphael A. Finkel and Jon Louis Bentley. Quad trees: A data structure for retrieval on composite keys. *Acta Inf.*, 4:1–9, 1974.
- [Gogic *et al.*, 1995] Goran Gogic, Henry A. Kautz, Christos H. Papadimitriou, and Bart Selman. The comparative linguistics of knowledge representation. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 862–869, 1995.
- [Gottlob, 1995] Georg Gottlob. Translating default logic into standard autoepistemic logic. *Journal of the ACM*, 42(4):711–740, 1995.
- [Levesque and Brachman, 1985] Hector J. Levesque and Ronald J. Brachman. A fundamental tradeoff in knowledge representation and reasoning (revised version). In R. J. Brachman and H. J. Levesque, editors, *Readings in Knowledge Representation*, pages 41–70. Kaufmann, Los Altos, CA, 1985.
- [Niveau *et al.*, 2010] Alexandre Niveau, Hélène Fargier, Cédric Pralet, and Gérard Verfaillie. Knowledge compilation using interval automata and applications to planning. In *Proceedings of the European Conference on Artificial Intelligence (ECAI)*, pages 459–464, 2010.
- [Rossi *et al.*, 2006] Francesca Rossi, Peter van Beek, and Toby Walsh, editors. *Handbook of Constraint Programming*. Elsevier, 2006.
- [Russell and Norvig, 2010] Stuart J. Russell and Peter Norvig. *Artificial Intelligence — A Modern Approach (3. internat. ed.)*. Pearson Education, 2010.
- [Srinivasan *et al.*, 1990] Arvind Srinivasan, Timothy Kam, Sharad Malik, and Robert K. Brayton. Algorithms for discrete function manipulation. In *Proceedings of the International Conference on Computer Aided Design (ICCAD)*, pages 92–95, November 1990.
- [van Harmelen *et al.*, 2008] Frank van Harmelen, Vladimir Lifschitz, and Bruce Porter, editors. *Handbook of Knowledge Representation*, volume 1. Elsevier Science, 2008.
- [Walsh, 2000] Toby Walsh. SAT v CSP. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*, pages 441–456, 2000.