

Patient Selection for Clinical Trials Using Temporalized Ontology-Mediated Query Answering

Franz Baader
Institute for Theoretical Computer
Science, TU Dresden
franz.baader@tu-dresden.de

Stefan Borgwardt
Institute for Theoretical Computer
Science, TU Dresden
stefan.borgwardt@tu-dresden.de

Walter Forkel
Institute for Theoretical Computer
Science, TU Dresden
walter.forkel@tu-dresden.de

ABSTRACT

Finding suitable candidates for clinical trials is a labor-intensive task that requires expert medical knowledge. Our goal is to design (semi-)automated techniques that can support clinical researchers in this task. We investigate the issues involved in designing formal query languages for selecting patients that are eligible for a given clinical trial, leveraging existing ontology-based query answering techniques. In particular, we propose to use a temporal extension of existing approaches for accessing data through ontologies written in Description Logics. We sketch how such a query answering system could work and show that eligibility criteria and patient data can be adequately modeled in our formalism.

CCS CONCEPTS

• **Computing methodologies** → **Temporal reasoning**; *Description logics*; • **Applied computing** → **Health care information systems**;

KEYWORDS

clinical trials; patient cohort recruitment; temporal description logic

ACM Reference Format:

Franz Baader, Stefan Borgwardt, and Walter Forkel. 2018. Patient Selection for Clinical Trials Using Temporalized Ontology-Mediated Query Answering. In *WWW '18 Companion: The 2018 Web Conference Companion, April 23–27, 2018, Lyon, France*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3184558.3191538>

1 INTRODUCTION

Clinical trials play an important role in the evaluation of new medications and treatments. After designing the study, the first main task is to find patients that are eligible for a given study, i.e., that satisfy all *inclusion criteria* and do not satisfy any *exclusion criteria*. Unfortunately, in practice it is often a resource-intensive task to recruit enough patients to get statistically meaningful results. However, the increased usage of Electronic Health Records (EHRs) in hospitals offers a promising opportunity to improve the recruitment process by automating parts of it.

Even though data are recorded more and more in semi-structured ways, most of them are not sufficiently structured to be processed directly using a logic-based formalism. This means that all the

complications of natural language processing (NLP) apply when the computer is supposed to find semantically matching patients.

Apart from the NLP nature of the task, many other challenges need to be overcome to automate patient selection. As already emphasized in [23], a major challenge lies in the fact that criteria are described on different levels of granularity, which range from quite specific to very general. This can often be bridged by using medical background knowledge that links broad categories (“*lung cancer*”) to more specific ones (“*adenocarcinoma*”) or even to more detailed descriptions (“*malignant neoplasm was found in the left lower lobe*”). Fortunately, a large amount of medical knowledge is nowadays readily available in a machine-processable form in medical ontologies such as SNOMED CT,¹ which is easily convertible into a Description Logic representation.

There exist many different approaches to the problem of automatic patient matching. Most are based on classic information retrieval techniques [13, 21], while others are logic-based [12, 23]. In this paper we consider the problems encountered when designing an ontology-based query formalism that automatically maps eligible patients to clinical trials. We put a special emphasis on analyzing which are the conditions that the query formalism must be able to express. For the moment, we mostly ignore the problems caused by the fact that natural language descriptions of eligibility criteria need to be translated automatically into our formalism. This will, however, be an important topic for future research.

2 BACKGROUND

Our observations are based on real world datasets which we will briefly introduce in the following. The data consist of three main parts: the patient data containing the EHRs, the clinical trials, and the medical ontologies that can provide the necessary background knowledge to the computer.

2.1 Patient Data

Our observations about patient data are based on the real-world MIMIC-III² dataset. It is a de-identified dataset of hospital admissions collected in two different hospitals in the US. In total it includes data associated with over 40.000 patients who stayed in intensive care units (ICUs). Each admission is associated with a number of diagnoses and procedures specified in the ICD-9 coding system.³ Additionally, reports, care notes, and discharge summaries collected during the admission are available in unstructured text form. In contrast to the diagnosis codes, each note has an exact

This paper is published under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '18 Companion, April 23–27, 2018, Lyon, France

© 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY 4.0 License.

ACM ISBN 978-1-4503-5640-4/18/04.

<https://doi.org/10.1145/3184558.3191538>

¹<https://www.snomed.org/snomed-ct>

²<https://mimic.physionet.org>

³<http://www.who.int/classifications/icd/en/>

time stamp, which allows more fine granular insights into the order of events occurring during a patient’s stay in the hospital.

Since MIMIC-III contains EHRs of patients that stayed in ICUs, the patients mostly suffer from severe diseases like cancer, or need an acute surgery or treatment, e.g., due to a stroke or a serious accident. Because of this it is much easier to find information about, e.g., cancer patients in MIMIC-III, than it is to find information about a patient with e.g., chronic lower back pain.

2.2 Clinical Trial Specifications

A clinical trial specifies a collection of eligibility criteria, each of which is either an *inclusion criterion*, which needs to be satisfied by an eligible patient, or an *exclusion criterion*, which must not be satisfied by the patient. More than 250.000 specifications of real-world trials can be found at ClinicalTrials.gov.⁴ Criteria are given in text form and range from very specific ones (“Serum creatinine $\leq 2.5 \frac{mg}{dL}$ ”⁵) to very general ones (“Uncontrolled chronic disease”⁶). Building a system that can deal with such a broad range of granularities is one of the challenges in automatic patient recruitment.

2.3 Medical Ontologies

The classification systems offered by medical ontologies have a great potential to bridge the gap between broad selection criteria and very fine granular patient data [23]. In recent years, numerous medical ontologies have been constructed and enjoy a growth in popularity. The modeling paradigms range from relatively high-level hierarchical codings like ICD-9 to very detailed formal ontologies modeled in Description Logics, like SNOMED CT. Mapping concepts from one system to another can be problematic, but is facilitated by the UMLS Metathesaurus,⁷ which provides mappings between equivalent concepts in different ontologies.

We will present our examples mainly from the point of view of ICD-9 and SNOMED CT, the latter of which can be formulated in the Description Logic \mathcal{EL} . We translate all ICD-9 codes from the patient database into SNOMED CT, which allows a more fine-grained classification of diagnoses. Since SNOMED CT and ICD-9 have been developed for different purposes, it is often not possible to find an exact match for each code and concept. For example the ICD-9 code 174.9 “Malignant neoplasm of breast (female), unspecified” is mapped to 41 possible concepts in SNOMED CT.

Originally, most ontologies were designed to provide a reference for doctors and other domain experts about the relations of symptoms and diseases. Modeling patient data in such ontologies can be problematic, since they provide only limited support to representing a patient’s history. Especially the support for temporal sequences, which play an important role when modeling patient records [13], is not included in a directly applicable way [12].

2.4 Description Logic

Previous work [16, 25] has proposed custom ontologies that define temporal concepts, which can be used to model clinical narratives. Unfortunately, this approach is not easily amenable to *reasoning*

over these temporal representations, i.e., checking whether a patient’s history actually matches a temporal criterion. In this paper, we base our approach on temporal ontology and query languages that include a dedicated temporal semantics [2]. Since existing medical ontologies are non-temporal and should ideally be used without modifications, we in particular focus on temporal *query* languages to express inclusion and exclusion criteria. Our query language is similar to the one proposed in [14], but we extend it in several ways to accommodate the needs of clinical trial criteria.

The Description Logic \mathcal{EL} [4, 9] is a popular formalism for reasoning over medical ontologies, such as SNOMED CT, the Gene Ontology,⁸ and the NCI Thesaurus.⁹ It is based on the notions of *concepts* (“cancer” or “liver”) and *roles* (“found in”), which express relations between concepts (“liver cancer” is “found in” the “liver”) as well as concrete patient data (patient 143 is “diagnosed with” the disease “liver cancer”¹⁰). As the most basic queries that can be stated over patient data and \mathcal{EL} ontologies, we consider *conjunctive queries* (CQs), which form a subset of the database query language SQL; for example, the CQ

$$\exists y. \text{diagnosed_with}(x, y) \wedge \text{Cancer}(y) \quad (1)$$

asks for all patients x that are diagnosed with a disease y that is classified as a cancer. Note that “diagnosed with” is not a relation from SNOMED CT, but a relation that is available in the patient database. The variable y is existentially quantified since we are not interested in the precise cancer diagnosis for each patient, but only need to ensure that there *exists* such a diagnosis. In the presence of the background knowledge (e.g., from SNOMED CT) that “liver cancer” is a subconcept of “cancer”, the answers to (1) include patient 143 from above.

3 CASE ANALYSIS

In the following, we highlight specific issues involved when trying to model clinical trials and patient data. Throughout this section, we use CQs, and later temporal CQs [6], as query language. For each example, we discuss how this query language needs to be extended in order to handle patient data and eligibility criteria.

In many cases, the exact set of patients that match the trial criteria can only be approximated, and the final decision needs to be taken by the doctor in charge. Since the goal is to reduce the workload as much as possible, the list of candidates returned by our querying approach should not contain too many false positives. It is even more important, however, to eliminate the presence of false negatives, i.e., not to discard any eligible patients; otherwise, the system would be nearly worthless for the doctors since in the end they still have to look at all patient records themselves. When designing the queries, we should thus determine whether they over- or under-approximate the actual set of eligible patients. *Under-approximating* the answers to an *inclusion criterion* may generate false negatives, and the same can happen when *over-approximating* the answers to an *exclusion criterion*.

⁴<https://clinicaltrials.gov/>

⁵From study NCT01064557; see <https://clinicaltrials.gov/ct2/show/NCT01064557>

⁶NCT01693861

⁷https://www.nlm.nih.gov/research/umls/knowledge_sources/metathesaurus/

⁸<http://www.geneontology.org/>

⁹<https://ncit.nci.nih.gov/>

¹⁰Patient 143 from MIMIC-III

3.1 Ambiguity

Apart from ambiguity in natural language, which is exacerbated by the extensive use of abbreviations in clinical texts (for example, “TBI” could mean either “total body irradiation” or “traumatic brain injury”), ambiguity also appears in the structured data. For example, patient 6470 from the MIMIC-III database has two recorded admissions; the first admission lists the ICD-9 diagnosis code 250.61 “Diabetes with neurological manifestations, type I, not stated as uncontrolled”, while in the second admission another doctor used 250.01 “Diabetes mellitus without mention of complications, type I, not stated as uncontrolled”. This may have been intended, i.e., the “neurological manifestations” disappeared in the meantime, or they were simply not relevant for the second admission.

Now imagine that we are searching for patients that are diagnosed with the SNOMED CT concept “Diabetes mellitus type 1”, using the CQ

$$\exists y.\text{diagnosed_with}(x, y) \wedge \text{Diabetes_mellitus_type_1}(y) \quad (2)$$

This query would only find the second admission, leading to an underestimation of the duration of the diabetes in the example patient. The reason is that, while the ICD-9 code 250.01 is correctly matched by the UMLS Metathesaurus to subconcepts of “Diabetes mellitus type 1” in SNOMED CT, this is not the case for 250.61. Rather, 250.61 is matched to “Neurological disorder associated with type I diabetes mellitus (disorder)”, which is defined as a “Neurological disorder” that is “associated with” “Diabetes mellitus type 1” (where “associated with” is a role in SNOMED CT). Reasoning over SNOMED CT does not allow us to draw the conclusion that the patient had “Diabetes mellitus type 1” at the time of the first admission.

One possibility to resolve this is to add an ontological axiom to SNOMED CT to state that every patient “diagnosed with” a disease that is “associated with” a second disease is automatically “diagnosed with” the second disease as well. However, it is not clear how to practically answer CQs in the presence of such axioms, although theoretically optimal algorithms are known [22]. Another way would be to under-approximate the effect of this axiom by supplementing the query (2) with the additional CQ

$$\exists y, z.\text{diagnosed_with}(x, y) \wedge \text{associated_with}(y, z) \wedge \text{Diabetes_mellitus_type_1}(z) \quad (3)$$

and taking the union of the results. This is a sound and nearly complete approximation of the proposed axiom, because SNOMED CT contains only chains of “associated with” relations of length at most 2. That is, by adding one more intermediate variable to this query, we can actually capture *all* associated diseases that may occur in SNOMED CT.

3.2 Negation

To obtain the candidates for a single clinical trial, we need to query the database for all patients that satisfy all inclusion criteria, while *not* satisfying any exclusion criteria of that trial. Formally, we can express this as

$$I_1(x) \wedge \dots \wedge I_n(x) \wedge \neg E_1(x) \wedge \dots \wedge \neg E_m(x), \quad (4)$$

where we combine the CQs $I_i(x)$ for the inclusion criteria and $E_j(x)$ for the exclusion criteria using the logical connectives *and* (\wedge) and *not* (\neg).

This can lead to problems, since sometimes a diagnosis is not stated in a patient’s admission simply because it is irrelevant for the primary reason of the admission. For example, consider looking for patients with “Post-traumatic stress disorder (PTSD)”. In contrast to cancer or diabetes, PTSD is not relevant for many hospital admissions and is therefore likely to be omitted in most cases. Indeed, patient 9248 (among many others) from MIMIC-III has five recorded admissions over a period of three years, and PTSD is listed as diagnosis only in the 2nd and 4th admissions, but not in the 1st, 3rd or 5th. This is problematic when evaluating *exclusion criteria*, since we can never be certain about the *absence* of a disease. From a formal point of view, there is a similar problem: SNOMED CT is modeled in the DL \mathcal{EL} , which does not support negative statements. In other words, the query (4) will in general not yield any results over the (incomplete) patient data, even when augmented with SNOMED CT.

An obvious solution is to ask only for patients for which the data actually indicates that they satisfy an exclusion criterion (which is an under-approximation), and then remove these patients from the query results (which yields an over-approximation). This can be modeled by adding the so-called epistemic operator \mathbf{K} (for “knowledge”), as investigated in [15, 19], to our query language. We can then ask for patients that satisfy all inclusion criteria, and for which *we do not know* for certain that they satisfy any exclusion criteria:

$$I_1(x) \wedge \dots \wedge I_n(x) \wedge \neg \mathbf{K}E_1(x) \wedge \dots \wedge \neg \mathbf{K}E_m(x) \quad (5)$$

In addition, negation also appears *inside* single criteria, for example in “cancer other than skin cancer,”¹¹ which can only be expressed as a combination of positive and negative statements that refer to the same existentially quantified variable:

$$\exists y.\text{diagnosed_with}(x, y) \wedge \text{Cancer}(y) \wedge \neg \text{Skin_Cancer}(y) \quad (6)$$

However, this is not a CQ anymore, and adding negation inside of CQs quickly leads to undecidability [18]. It should be possible to extend CQs by *epistemic negation* $\neg \mathbf{K}$ similar to (5) while staying decidable (but this particular combination of CQs and epistemic negation has not been investigated so far). In this way, we can obtain all patients that have a diagnosis y that is a cancer diagnosis, for which it is not explicitly stated in the EHR that this cancer is a skin cancer.

3.3 Time

To access temporal patient data, where all diagnoses are annotated with time stamps, we propose to use *temporal conjunctive queries* (TCQs) [6, 14], in which CQs are combined with linear-temporal operators. As a simple example, with a TCQ we can ask for all patients that have a “history of cancer,”¹² i.e., those that were diagnosed with cancer *at some point in the past*, using the temporal operator \diamond :

$$\diamond_{(-\infty, 0]}(\exists y.\text{diagnosed_with}(x, y) \wedge \text{Cancer}(y)), \quad (7)$$

where the time interval $(-\infty, 0]$ refers to the whole history (relative to the current time point 0). More complex queries include “type 1 diabetes with duration at least 12 months,”¹³ utilizing the temporal

¹¹NCT02873052

¹²NCT00064766

¹³NCT02280564

always operator \square :

$$\diamond_{(-\infty, 0]} \square_{[-12, 0]} (\exists y. \text{diagnosed_with}(x, y) \wedge \text{Diabetes}(y)). \quad (8)$$

In this case, we are looking for some time point t in the past ($\diamond_{(-\infty, 0]}$) for which during the whole time interval $[-12, 0]$ relative to t (that is, $[t - 12, t]$) the patient had a diagnosis of diabetes (such specific time intervals are not supported by TCQs yet, but have been investigated in other settings [1, 5]). Here, we used months as the basic time unit; in practice, statements with a different temporal granularity like *days* or *years* have to be converted into a common time unit before combining them into a logical query.

In general we can distinguish two different kinds of temporal queries: (7) asks whether there exists a *single* point in time at which something is true, while (8) asks whether something was true for the *whole* duration of a given interval. Since the data is restricted to the times where the patient was in the hospital, for the latter kind of queries we need to make additional assumptions about what happens in between admissions. Did the patient still suffer from diabetes? Maybe the disease was healed and after some time reappeared? In the specific case of (8), medical knowledge tells us that diabetes cannot (yet) be cured, which means that it suffices to find patients diagnosed with diabetes at least 12 months ago. In principle, such knowledge can be encoded in the setting proposed in [5, 6, 14]. Unfortunately, currently there does not exist a formalized ontology containing information on the usual duration and the curability of diseases, which means that for now such information would have to be hand-coded for each clinical trial.

Similar to (8), consider the following criterion: “*Temperature > 38.5°C for seven consecutive days within the 14 days prior to first study dose.*”¹⁴

$$\diamond_{[-7, 0]} \square_{[-7, 0]} (\exists y. \text{measured_temperature}(x, y) \wedge (y > 38.5^\circ\text{C})). \quad (9)$$

The statement $y > 38.5^\circ\text{C}$ is discussed in more detail in the next section. For now, we focus on the problem of finding valid continuous intervals when given only single measurements. The availability of body temperature measurements varies from patient to patient: When a patient stays in an ICU, the temperature is measured frequently, whereas outside the ICU it may not be measured at all.

One strategy to interpolate the values in between measurements could be to assume them to be constant until the next measurement. This might be a good approach if the measurements are a few hours apart, but certainly not if they are years apart. It is difficult to define intervals in which values stay the same, since it strongly depends on the context. Moreover, this would again require medical background knowledge that is currently not available in a logical form.

An approximation that does not require background knowledge is based on the assumption that abnormal values, such as a raised temperature or decreased red blood cell count, will be measured more frequently. Therefore, we can over-approximate (9) by looking for an interval of 7 days during which

- at least one supporting measurement was taken and
- no contradicting measurement exists:

¹⁴NCT02355184

$$\diamond_{[-7, 0]} (\exists y. \text{measured_temperature}(x, y) \wedge (y > 38.5^\circ\text{C})) \wedge \square_{[-7, 0]} \neg \mathbf{K} (\exists y. \text{measured_temperature}(x, y) \wedge (y \leq 38.5^\circ\text{C})).$$

A similar idea can be used to approximate queries like (8).

3.4 Measurements

EHRs include a lot of numerical data, like age, doses of medications, or blood pressure. In addition, eligibility criteria often require a value to be smaller or greater than some threshold, e.g., “*serum creatinine <= 2.5 mg/dL.*”¹⁵ A high serum creatinine value indicates that the kidney is not working properly. The measurement can be obtained by analyzing a small blood sample. A system for patient selection should therefore be able to deal with values and value comparisons to some extent. DLs and CQs can be extended by so called *concrete domains* to support reasoning with numbers [3, 7, 8, 11, 24]. A query could look similar to the following:

$$\exists y. \text{measured_serum_creatinine}(x, y) \wedge (y \leq 2.5 \text{mg/dL}) \quad (10)$$

Here, the temporal component is also important: We do not consider an arbitrary measurement, but rather we want to find the most recent one that can still be considered valid (see the discussion in the previous section).

Sometimes value comparisons are explicitly combined with temporal expressions, e.g., when talking about aggregated values over a given interval. As an example, consider “*any previous treatment with sodium fluoride at daily doses ≥ 5 mg/day for a period exceeding 1 month.*”¹⁶ To answer the query, it is necessary to find a prescription of sodium fluoride with a daily sum greater or equal to 5 mg, which also includes taking 2.5 mg sodium fluoride twice a day. Unfortunately, adding arithmetic expressions beyond simple value comparisons to a DL quickly results in high computational complexity or even undecidability of the logic. To work around this limitation, we need to employ a preprocessor that normalizes all numbers in the patient data.

3.5 Limits of Automated Systems

There are a number of criteria that simply cannot be answered from the available patient data. Among other factors, this can be because the data do not include certain kinds of information (“*Positive pregnancy test*”),¹⁷ the criteria ask for subjective things (“*Willingness to replace the missing tooth/teeth with dental implants*”),¹⁸ or refer to time points that potentially lie in the future (“*Vaccination with a live vaccine within 28 days prior to randomization*”).¹⁹ For such criteria, we cannot hope to find answers automatically. Instead, we should detect and ignore them, because in the best case they do not contribute anything to the query, while in the worst case they could introduce errors. In a working system, these criteria should be marked to indicate to the doctor that they are still open and need to be checked manually.

¹⁵NCT01064557

¹⁶NCT00171704

¹⁷NCT01444170

¹⁸NCT02884401

¹⁹NCT01332968

4 RELATED WORK

Previous work has already considered using ontologies for patient selection for clinical trials before. Patel et al. [23] worked with patient records from Columbia University Medical Center that were recorded using the MED ontology. They mapped MED to SNOMED CT using a semi-automated approach that was guided by domain experts. The patient records were then integrated using a pattern matching rule-based approach. They showed impressively that it is actually possible to find patient matches using an ontology, and were able to scale their approach to one year of patient data.

Besana et al. [12] focused on 200 trials about prostate cancer and annotated them manually with UMLS concepts. As formal basis, they use OWL (which is based on DLs) together with SWRL rules, which allows them to add rules for temporal relations. They then load one patient at a time into the ontology and query the studies that the patient is eligible for. Their approach allows traceability of the results, which is a very desirable property. While they demonstrate that patients can be selected using their formal framework, they assume that the data are already formalized.

Tao et al. [25] further analyzed and modeled the temporal patterns that occur in patient data. To represent them they introduce the OWL-based CNTRO 2.0 ontology for clinical narratives. Later, Crowe and Tao [16] classified most temporal statements occurring in descriptions of clinical trials and clinical guidelines into 16 basic temporal patterns that are expressible in CNTRO 2.0. Unfortunately, CNTRO 2.0 is not suitable for temporal query answering, since it can express temporal statements, but does not provide a temporal semantics and allows only rudimentary temporal inferences.

Other approaches to model temporal medical data use graph- or constraint-based formalisms to representing and reasoning with temporal statements [13, 20].

For a survey regarding also non-temporal, non-logical proposals for automated processing of EHRs and other medical data, see [21].

5 CONCLUSION

Automated patient recruitment is an important and at the same time challenging task. In this paper we have presented the issues involved when representing EHRs and eligibility criteria for clinical trials in a formal knowledge representation system that supports effective reasoning. We have seen that the support for temporal operators plays a crucial role in representing the data and the criteria in a concise way. In addition, epistemic negation, bounded temporal intervals like $[0, 12]$, and concrete domains are desirable means of expressiveness in this setting, but also in many other real-world scenarios with incomplete data.

In the future, we aim to develop effective procedures that enable automated reasoning over knowledge expressed in this way. While answering TCQs has been addressed in our previous work (e.g., [6, 14]), query answering for extensions of TCQs that support epistemic negation, bounded temporal intervals, and concrete domains have not been investigated until now. On the theoretical side, we intend to investigate the decidability and complexity of the corresponding decision problems. On the practical side, our aim is to build a prototype system that demonstrates the applicability of this kind of temporal reasoning for selecting patients for clinical trials. Additional important features of such a system include, firstly,

perspicuity of the automatically computed results: for each patient that is marked as eligible or not eligible, the system should be able to explain why this is the case. Secondly, in cases with no exact matches, we suggested to return patients that only partially satisfy the criteria. To support the recruiters, the returned list of possibly eligible patients should then be ranked by its degree of certainty or significance [10, 17]. In addition, the properties missing for an exact match should be explicated.

Another important topic for future research is tackling the problem of (semi-)automatically translating natural language descriptions of eligibility criteria into formal queries w.r.t. a given medical ontology.

6 ACKNOWLEDGMENTS

This work was supported by DFG in the joint DFG-ANR project BA 1122/19-1 (GOASQ). The first author is also partially supported by DFG within the collaborative research center SFB 912 (HAEC), where TCQs with the mentioned extensions are also relevant.

REFERENCES

- [1] Rajeev Alur and Thomas A. Henzinger. 1993. Real-Time Logics: Complexity and Expressiveness. *Information and Computation* 104, 1 (1993), 35–77. <https://doi.org/10.1006/inco.1993.1025>
- [2] Alessandro Artale, Roman Kontchakov, Alisa Kovtunova, Vladislav Ryzhikov, Frank Wolter, and Michael Zakharyashev. 2017. Ontology-Mediated Query Answering over Temporal Data: A Survey (Invited Talk). In *Proc. of the 24th Int. Symp. on Temporal Representation and Reasoning (TIME'17) (Leibniz International Proceedings in Informatics)*, Sven Schewe, Thomas Schneider, and Jef Wijsen (Eds.), Vol. 90. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 1:1–1:37. <https://doi.org/10.4230/LIPIcs.TIME.2017.1>
- [3] Alessandro Artale, Vladislav Ryzhikov, and Roman Kontchakov. 2012. DL-Lite with Attributes and Datatypes. In *Proc. of the 20th Eur. Conf. on Artificial Intelligence (ECAI'12) (Frontiers in Artificial Intelligence and Applications)*, Luc De Raedt, Christian Bessiere, Didier Dubois, Patrick Doherty, Paolo Frasconi, Fredrik Heintz, and Peter Lucas (Eds.), Vol. 242. IOS Press, 61–66. <https://doi.org/10.3233/978-1-61499-098-7-61>
- [4] Franz Baader. 2003. Terminological Cycles in a Description Logic with Existential Restrictions. In *Proc. of the 18th Int. Joint Conf. on Artificial Intelligence (IJCAI'03)*, Georg Gottlob and Toby Walsh (Eds.), Morgan Kaufmann, 325–330. <http://ijcai.org/Proceedings/03/Papers/048.pdf>
- [5] Franz Baader, Stefan Borgwardt, Patrick Koopmann, Ana Ozaki, and Veronika Thost. 2017. Metric Temporal Description Logics with Interval-Rigid Names. In *Proc. of the 11th Int. Symp. on Frontiers of Combining Systems (FroCoS'17) (Lecture Notes in Computer Science)*, Clare Dixon and Marcelo Finger (Eds.), Vol. 10483. Springer, 60–76. https://doi.org/10.1007/978-3-319-66167-4_4
- [6] Franz Baader, Stefan Borgwardt, and Marcel Lippmann. 2013. Temporalizing Ontology-Based Data Access. In *Proc. of the 24th Int. Conf. on Automated Deduction (CADE'13) (Lecture Notes in Computer Science)*, Maria Paola Bonacina (Ed.), Vol. 7898. Springer, 330–344. https://doi.org/10.1007/978-3-642-38574-2_23
- [7] Franz Baader, Stefan Borgwardt, and Marcel Lippmann. 2017. Query Rewriting for DL-Lite with n -ary Concrete Domains. In *Proc. of the 26th Int. Joint Conf. on Artificial Intelligence (IJCAI'17)*, Carles Sierra (Ed.), IJCAI, 786–792. <https://doi.org/10.24963/ijcai.2017/109>
- [8] Franz Baader, Sebastian Brandt, and Carsten Lutz. 2005. Pushing the \mathcal{EL} Envelope. In *Proc. of the 19th Int. Joint Conf. on Artificial Intelligence (IJCAI'05)*, Leslie Pack Kaelbling and Alessandro Saffioti (Eds.), Professional Book Center, 364–369. <http://ijcai.org/Proceedings/09/Papers/053.pdf>
- [9] Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider (Eds.). 2007. *The Description Logic Handbook: Theory, Implementation, and Applications* (2 ed.). Cambridge University Press.
- [10] Franz Baader and Oliver Fernández Gil. 2017. Decidability and Complexity of Threshold Description Logics Induced by Concept Similarity Measures. In *Proc. of the 32nd Annual ACM Symp. on Applied Computing*, Sung Y. Shin, Dongwan Shin, and Maria Lencastre (Eds.), ACM, 983–988. <https://doi.org/10.1145/3019612.3019715>
- [11] Franz Baader and Philipp Hanschke. 1991. A Scheme for Integrating Concrete Domains into Concept Languages. In *Proc. of the 12th Int. Joint Conf. on Artificial Intelligence (IJCAI'91)*, John Mylopoulos and Raymond Reiter (Eds.), Morgan Kaufmann, 452–457. <http://ijcai.org/Proceedings/91-1/Papers/070.pdf>

- [12] Paolo Besana, Marc Cuggia, Oussama Zekri, Annabel Bourde, and Anita Burgun. 2010. Using Semantic Web Technologies for Clinical Trial Recruitment. In *Proc. of the 9th Int. Semantic Web Conf. (ISWC'10) (Lecture Notes in Computer Science)*, Peter F. Patel-Schneider, Yue Pan, Pascal Hitzler, Peter Mika, Lei Zhang, Jeff Z. Pan, Ian Horrocks, and Birte Glimm (Eds.), Vol. 6497. Springer, 34–49. https://doi.org/10.1007/978-3-642-17749-1_3
- [13] Luca Bonomi and Xiaoqian Jiang. 2018. Patient Ranking with Temporally Annotated Data. *Journal of Biomedical Informatics* 78 (2018), 43–53. <https://doi.org/10.1016/j.jbi.2017.12.007>
- [14] Stefan Borgwardt, Marcel Lippmann, and Veronika Thost. 2015. Temporalizing Rewritable Query Languages over Knowledge Bases. *Journal of Web Semantics* 33 (2015), 50–70. <https://doi.org/10.1016/j.websem.2014.11.007>
- [15] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. 2007. *EQL-Lite: Effective First-Order Query Processing in Description Logics*. In *Proc. of the 20th Int. Joint Conf. on Artificial Intelligence (IJCAI'07)*, Manuela M. Veloso (Ed.), 274–279. <http://www.ijcai.org/Proceedings/07/Papers/042.pdf>
- [16] Christopher L. Crowe and Cui Tao. 2015. Designing Ontology-based Patterns for the Representation of the Time-Relevant Eligibility Criteria of Clinical Protocols. *AMIA Joint Summits on Translational Science Proc.* 2015 (2015), 173–177. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4525239/>
- [17] Andreas Ecke, Rafael Peñaloza, and Anni-Yasmin Turhan. 2015. Similarity-Based Relaxed Instance Queries. *Journal of Applied Logic* 13, 4, Part 1 (2015), 480–508. <https://doi.org/10.1016/j.jal.2015.01.002>
- [18] Víctor Gutiérrez-Basulto, Yazmín Ibáñez-García, Roman Kontchakov, and Egor V. Kostylev. 2015. Queries with Negation and Inequalities over Lightweight Ontologies. *Journal of Web Semantics* 35 (2015), 184–202. <https://doi.org/10.1016/j.websem.2015.06.002>
- [19] Víctor Gutiérrez-Basulto and Szymon Klarman. 2012. Towards a Unifying Approach to Representing and Querying Temporal Data in Description Logics. In *Proc. of the 6th Int. Conf. on Web Reasoning and Rule Systems (RR'12) (Lecture Notes in Computer Science)*, Markus Krötzsch and Umberto Straccia (Eds.), Vol. 7497. Springer, 90–105. https://doi.org/10.1007/978-3-642-33203-6_8
- [20] George Hripcsak, Li Zhou, Simon Parsons, Amar K. Das, and Stephen B. Johnson. 2005. Modeling Electronic Discharge Summaries as a Simple Temporal Constraint Satisfaction Problem. *Journal of the American Medical Informatics Association* 12, 1 (2005), 55–63. <https://doi.org/10.1197/jamia.m1623>
- [21] Felix Köpcke and Hans-Ulrich Prokosch. 2014. Employing Computers for the Recruitment into Clinical Trials: A Comprehensive Systematic Review. *Journal of Medical Internet Research* 16, 7 (2014), e161. <https://doi.org/10.2196/jmir.3446>
- [22] Markus Krötzsch and Sebastian Rudolph. 2007. Conjunctive Queries for $\mathcal{E}\mathcal{L}$ with Role Composition. In *Proc. of the 20th Int. Workshop on Description Logics (DL'07) (CEUR Workshop Proceedings)*, Diego Calvanese, Enrico Franconi, Volker Haarslev, Domenico Lembo, Boris Motik, Anni-Yasmin Turhan, and Sergio Tessaris (Eds.), Vol. 250. 355–362. http://ceur-ws.org/Vol-250/paper_58.pdf
- [23] Chintan Patel, James Cimino, Julian Dolby, Achille Fokoue, Aditya Kalyanpur, Aaron Kershenbaum, Li Ma, Edith Schonberg, and Kavitha Srinivas. 2007. Matching Patient Records to Clinical Trials Using Ontologies. In *Proc. of the 6th Int. Semantic Web Conf. (ISWC'07) (Lecture Notes in Computer Science)*, Karl Aberer, Key-Sun Choi, Natasha Noy, Dean Allemang, Kyung-Il Lee, Lyndon Nixon, Jennifer Goldbeck, Peter Mika, Diana Maynard, Riichiro Mizoguchi, Guus Schreiber, and Philippe Cudré-Mauroux (Eds.), Vol. 4825. Springer, 816–829. https://doi.org/10.1007/978-3-540-76298-0_59
- [24] Ognjen Savković and Diego Calvanese. 2012. Introducing Datatypes in *DL-Lite*. In *Proc. of the 20th Eur. Conf. on Artificial Intelligence (ECAI'12) (Frontiers in Artificial Intelligence and Applications)*, Luc De Raedt, Christian Bessière, Didier Dubois, Patrick Doherty, Paolo Frasconi, Fredrik Heintz, and Peter Lucas (Eds.), Vol. 242. IOS Press, 720–725. <https://doi.org/10.3233/978-1-61499-098-7-720>
- [25] Cui Tao, Harold R. Solbrig, and Christopher G. Chute. 2011. CNTRO 2.0: A Harmonized Semantic Web Ontology for Temporal Relation Inferencing in Clinical Narratives. *AMIA Joint Summits on Translational Science Proc.* 2011 (2011), 64–68. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3248753/>