

Temporalized Ontology-Mediated Query Answering under Minimal-World Semantics

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

Alisa Kovtunova
Institute for Theoretical Computer
Science, TU Dresden
alisa.kovtunova@tu-dresden.de

Chao Xu
Department of Philosophy, Peking
University
c.xu@pku.edu.cn

Beihai Zhou
Department of Philosophy, Peking
University
zhoubh@phil.pku.edu.cn

ABSTRACT

Selecting patients for clinical trials is very labor-intensive. Our goal is to design (semi-)automated techniques that can support clinical researchers in this task. In this paper we summarize our recent advances towards such a system: First, we present the challenges involved when representing electronic health records and eligibility criteria for clinical trials in a formal language. Second, we introduce temporal conjunctive queries with negation as a formal language suitable to represent clinical trials. Third, we describe our methodology for automatic translation of clinical trial eligibility criteria from natural language into our query language. The evaluation of our prototypical implementation shows promising results. Finally, we talk about the parts we are currently working on and the challenges involved.

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, Walter Forkel, Alisa Kovtunova, Chao Xu, and Beihai Zhou. 2019. Temporalized Ontology-Mediated Query Answering under Minimal-World Semantics. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Ontology-mediated query answering (OMQA) allows using background knowledge for answering user queries, supporting data-focused applications offering search, analytics, or data integration

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference'17, July 2017, Washington, DC, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

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

functionality. An *ontology* is a logical theory formulated in a decidable fragment of first-order logic, with a trade-off between the expressivity of the ontology and the efficiency of query answering. *Rewritability* is a popular topic of research, the idea being to reformulate ontological queries into database queries that can be answered by traditional database management systems [9, 15, 20, 24, 31].

The biomedical domain is a fruitful area for OMQA, due to the availability of large ontologies¹ and the demand for managing large amounts of patient data, in the form of *electronic health records (EHRs)* [17]. For example, for the preparation of clinical trials² a large number of patients need to be screened for eligibility, and an important area of current research is how to automate this process [8, 25, 32, 33].³ Last year, we proposed [5] a research program to (i) translate the eligibility criteria of clinical trials into logical queries, (ii) extract structured data from semi-structured EHRs, and (iii) answer the queries over the data with the help of an ontology (SNOMED CT⁴) to select the patients for the trials. In this paper, we report on progress made on aspects (i) and (iii), whereas (ii) and an integration of the different parts into a whole system remain future work.

Regarding (iii), we have formalized the expressivity requirements of clinical trials into the query language of *metric temporal conjunctive queries with negation (MTNCQs)* [12, 13], which includes all features proposed in [5] except for *concrete domains*, which allow to talk about measurements and dosages. However, we believe it is feasible to extend our language with concrete domains, as described in [6, 16]. We have also developed a practical algorithm based on query rewriting techniques. Towards (i), we have developed a prototype implementation based on state-of-the-art natural language tools that can translate a selection of criteria from the clinicaltrials.gov website into MTNCQs.

2 BACKGROUND

Our work is grounded in real-world datasets which we will briefly introduce in the following: the patient data containing the EHRs, the clinical trials, and the medical ontologies that can provide the necessary background knowledge to the computer.

¹<https://bioportal.bioontology.org>

²<https://clinicaltrials.gov>

³<https://n2c2.dbmi.hms.harvard.edu>

⁴<https://www.snomed.org/>

2.1 Patient Data

Our observations about patient data are based on the real-world MIMIC-III⁵ 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.⁶ 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 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.⁷ Criteria are given in text form and range from very specific ones (“Serum creatinine $\leq 2.5 \frac{mg}{dL}$ ”⁸) to very general ones (“Uncontrolled chronic disease”⁹). 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 [33]. 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,¹⁰ 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 [11], is not included in a directly applicable way [8].

2.4 Description Logic

Previous work [18, 34] 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 [3]. 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, 7] is a popular formalism for reasoning over medical ontologies, such as SNOMED CT, the Gene Ontology,¹¹ and the NCI Thesaurus.¹² 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”¹³). 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

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

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 METRIC TEMPORAL CONJUNCTIVE QUERIES WITH NEGATION

As argued in [5], evaluating clinical trial criteria over patient records requires both negated and temporal queries. For example, the eligibility criterion “history of migraine of at least 1 year prior to enrollment [...], but not a history of chronic migraine”¹⁴ contains both temporal and negated requirements. In order to represent the temporal dimension and adequately treat negative knowledge, the query language in [13] extends conjunctive queries by metric temporal operators and negation. *Metric temporal logic* (MTL) was

⁵<https://mimic.physionet.org>

⁶<http://www.who.int/classifications/icd/en/>

⁷<https://clinicaltrials.gov/>

⁸From study NCT01064557; see <https://clinicaltrials.gov/ct2/show/NCT01064557>

⁹<https://clinicaltrials.gov/ct2/show/NCT01693861>

¹⁰https://www.nlm.nih.gov/research/umls/knowledge_sources/metathesaurus/

¹¹<http://www.geneontology.org/>

¹²<https://ncit.nci.nih.gov/>

¹³Patient 143 from MIMIC-III

¹⁴<https://clinicaltrials.gov/ct2/show/NCT01625988>

originally designed for modelling and reasoning about real-time systems [1].

Unfortunately, ontologies and EHRs mostly contain positive information, and moreover employ open-world semantics, which means that a “history of chronic migraine” cannot be ruled out unless it is explicitly stated in the patient record that “the patient does not have a history of chronic migraine”. As a countermeasure for this behavior, we developed a *minimal-world semantics* [12] that allows for interpreting patient information as the most specific description. In other words, if a patient has not been diagnosed explicitly with a history of chronic migraine, she is eligible for the trial.

The temporal elements of clinical trial criteria are represented by operators of metric temporal logic. For example, the diamond operator $\Diamond_{[a,b]}$ can express that some event happened at some point between times a and b (relative to the current time). Thus, $\Diamond_{(-\infty,1 \text{ year}]}$ means “at least 1 year prior”. In total, the criterion from above could be written as the *metric temporal conjunctive query with negation (MTNCQ)*

$$\begin{aligned} & \Diamond_{(-\infty,1 \text{ year}]} (\exists y. \text{diagnosedWith}(x, y) \wedge \text{Migraine}(y)) \wedge \\ & \neg \Diamond_{(-\infty,0]} (\exists y. \text{diagnosedWith}(x, y) \wedge \text{ChronicMigraine}(y)), \end{aligned}$$

where \neg denotes negation.

A problem with EHRs is that they are usually *sparse*, i.e., data is only available for individual time points, since patients usually visit the doctor only sporadically. It is known that highly sparse medical datasets due to missing data are problematic statistical analysis [36]. Therefore, in [13] for handling missing data we proposed adding statements of the form

$$\Diamond_{1 \text{ year}} \text{ChronicMigraine} \sqsubseteq \text{ChronicMigraine} \quad (2)$$

to express that two diagnoses of chronic migraine that are one year apart mean that the patient suffered from migraine during the whole year. This can be used to bridge temporal gaps, thus alleviating the problem of sparse data.

Finally, we developed an algorithm to answer MTNCQs over sparse datasets via a *rewriting* technique, in which the query is reformulated into a database query that can be executed efficiently [12, 13].

4 TRANSLATING CLINICAL TRIAL CRITERIA INTO MTNCQS

In [37], we developed a methodology to translate clinical trial criteria into MTNCQs. Instead of a direct translation, we first annotate each part of a criterion with a semantic type, and then construct a formal query from this annotation.

We analyzed the criteria from clinicaltrials.gov to find the most important semantic types, which include *clinical findings*, *products*, *procedures* etc. from SNOMED CT, and additionally the types *age*, *time*, *number*, *comparison sign*, *negation*, and *conjunction*. For example, consider the criterion “history of lung disease other than asthma”.¹⁵ Its semantic annotation is shown in Figure 1. The concepts from SNOMED CT are annotated with the help of MetaMap tagger [2], a tool for recognizing concepts from the UMLS Metathesaurus, which subsumes SNOMED CT. Given a phrase or sentence,

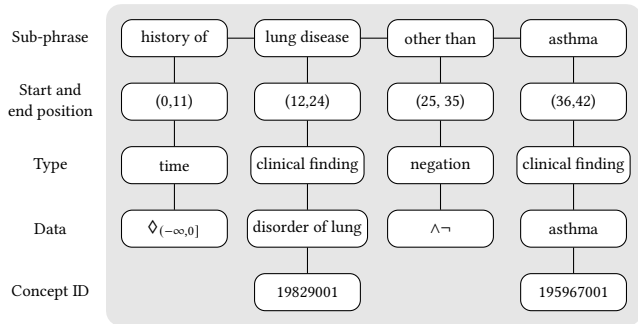


Figure 1: The semantic annotation for our example.

it returns the most likely phrase-concept matches. After refining these matches with the help of syntactic and semantic analysis using the Stanford NLP tools [27] and Word2Vec [28], the remaining parts are annotated using dedicated regular expressions for each semantic type, e.g. “history of” is recognized as a temporal expression that can be expressed as a diamond operator of MTL.

In the second step, the semantic annotations are used to recursively generate an MTNCQ. The SNOMED CT concepts are translated into basic CQs like

$$\exists y. \text{diagnosedWith}(x, y) \wedge \text{DisorderOfLung}(y),$$

which are then combined (when compatible) with conjunctions and negation into larger queries. Finally, temporal expressions are added. In our example, we obtain the MTNCQ

$$\begin{aligned} & \Diamond_{(-\infty,0]} (\exists y. \text{diagnosedWith}(x, y) \wedge \\ & \text{DisorderOfLung}(y) \wedge \neg \text{Asthma}(y)). \end{aligned}$$

We have evaluated this approach on a random selection of 401 natural language criteria from clinicaltrials.gov. In a manual evaluation, about 237–308 of them were found to be unanswerable by an automated system (according to the opinions of different evaluators). Such criteria include ones that refer to the future (“during study phase”) or ask for subjective information (“in the opinion of the investigator”, “willingness to”). The 93 criteria that were considered answerable by all evaluators were then translated by our system, and 58–70% of the translations were manually evaluated to be “good”, and about 11% were wrong, the remainder providing only a partial translation.

This means that much of the patient screening still has to be done by doctors, but an automated system can at least take into account basic criteria for pre-selecting eligible patients.

5 RELATED WORK

Previous work has considered using ontologies for patient selection for clinical trials before. Patel et al. [33] 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.

¹⁵<https://clinicaltrials.gov/ct2/show/NCT02548598>

Besana et al. [8] 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. [34] 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 [18] 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 [11, 22].

For a survey regarding also non-temporal, non-logical proposals for automated processing of EHRs and other medical data, see [25].

Regarding the automatic translation of eligibility criteria, Tu et al. [35] proposed a practical translation method based on the ERGO annotation, which is an intermediate representation for criteria. However, ERGO annotation can only be done manually or semi-automatically. Milian et al. [29, 30] focused on breast-cancer trials and summarized 165 patterns, and used these patterns and concept recognition tools to structure criteria. After that, they generated a formal representation by projecting the concepts in criteria to the predefined query template. There is also some work about extraction and representation of *partial* knowledge in eligibility criteria. Zhou et al. [38], Luo et al. [26] and Boland et al. [10] focused on the recognition and representation of temporal knowledge. Huang et al. [23] and Enger et al. [21] proposed several methods for detecting negated expressions.

In the field of natural language processing, automatic translation from natural language into formal language, e.g., first-order logic formulas, is also known as *automatic semantic parsing*. Dong et al. [19] proposed an automatic semantic parsing method based on machine learning, different from traditional rule-based or template-based methods.

6 CONCLUSION

Automated patient recruitment is an important and at the same time challenging task. In this paper we have reported on recent progress regarding an approach based on logic-based query answering. Our proposal supports (metric) temporal as well as negated queries, which is crucial to express as many eligibility criteria as possible. To support information about measurements and dosages, an extension with concrete domains [6, 16] remains as future work, as well as an automatic extraction of EHRs into description logic ABoxes. The latter presents a different set of challenges than the translation of clinical trial criteria, since EHRs often consist of long texts instead of only single sentences, and while they do not contain complicated

temporal constructions, resolving different mentions of the same entity among different sentences is a major challenge.

7 ACKNOWLEDGMENTS

This work was partially supported by DFG grant 389792660 as part of TRR 248 (see <https://perspicuous-computing.science>) and the joint DFG-ANR project BA 1122/19-1 (GOASQ).

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] Alan R Aronson. 2001. Effective mapping of biomedical text to the UMLS Metathesaurus: The MetaMap program.. In *Proc. AMIA Symposium*. American Medical Informatics Association, 17.
- [3] 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>
- [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, and Walter Forkel. 2018. Patient Selection for Clinical Trials Using Temporalized Ontology-Mediated Query Answering. In *Proc. HQA Workshop*, Franz Baader, Brigitte Grau, and Yue Ma (Eds.). ACM, 1069–1074. <https://doi.org/10.1145/3184558.3191538>
- [6] 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>
- [7] 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.
- [8] 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
- [9] Meghyn Bienvenu and Magdalena Ortiz. 2015. Ontology-mediated query answering with data-tractable description logics. In *Reasoning Web 11th International Summer School (Lecture Notes in Computer Science)*, Wolfgang Faber and Adrian Paschke (Eds.), Vol. 9203. Springer, 218–307. https://doi.org/10.1007/978-3-319-21768-0_9
- [10] Mary Regina Boland, Samson W Tu, Simona Carini, Ida Sim, and Chunhua Weng. 2012. EliXR-TIME: A temporal knowledge representation for clinical research eligibility criteria. *AMIA Transl. Sci. Proc.* 2012 (2012), 71.
- [11] 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>
- [12] Stefan Borgwardt and Walter Forkel. 2019. Closed-World Semantics for Conjunctive Queries with Negation over \mathcal{ELH}_\perp Ontologies. In *Proc. of the 16th Eur. Conf. on Logics in Artificial Intelligence (JELIA'19) (Lecture Notes in Artificial Intelligence)*, Francesco Calimeri, Nicola Leone, and Marco Manna (Eds.), Vol. 11468. Springer-Verlag, 371–386. https://doi.org/10.1007/978-3-030-19570-0_24
- [13] Stefan Borgwardt, Walter Forkel, and Alisa Kovtunova. 2019. Finding New Diamonds: Temporal Minimal-World Query Answering over Sparse ABoxes. In *Proc. of the 3rd Int. Joint Conf. on Rules and Reasoning (RuleML+RR'19) (Lecture Notes in Computer Science)*, Paul Foder, Marco Montali, Diego Calvanese, and Dumitru Roman (Eds.), Vol. 11784. Springer-Verlag, 3–18. https://doi.org/10.1007/978-3-030-31095-0_1
- [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, Benjamin Cogrel, Sarah Komla-Ebri, Roman Kontchakov, Davide Lanti, Martin Rezk, Mariano Rodriguez-Muro, and Guohui Xiao. 2017. Ontop: Answering SPARQL queries over relational databases. *Semantic Web* 8 (2017), 471–487. Issue 3.

- [16] Patrick Koopmann Christian Alrabbaa and Anni-Yasmin Turhan. 2019. Practical Query Rewriting for DL-Lite with Numerical Predicates. In *Proc. of the 5th Global Conf. on Artificial Intelligence (GCAI'19) (EPIC)*, Diego Calvanese and Luca Iocchi (Eds.). EasyChair.
- [17] Kathrin M. Cresswell and Aziz Sheikh. 2017. Inpatient Clinical Information Systems. In *Key Advances in Clinical Informatics*, Aziz Sheikh, Kathrin M. Cresswell, Adam Wright, and David W. Bates (Eds.). Academic Press, Chapter 2, 13–29. <https://doi.org/10.1016/B978-0-12-809523-2.00002-9>
- [18] Christopher L Crowe and Cui Tao. 2015. Designing ontology-based patterns for the representation of the time-relevant eligibility criteria of clinical protocols. *AMIA Transl. Sci. Proc.* 2015 (2015), 173. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4525239/>
- [19] Li Dong and Mirella Lapata. 2016. Language to logical form with neural attention. In *Proc. Annual Meeting of the ACL*.
- [20] Thomas Eiter, Magdalena Ortiz, Mantas Šimkus, Trung-Kien Tran, and Guohui Xiao. 2012. Query Rewriting for Horn-*SHIQ* Plus Rules. In *Proc. of the 26th AAAI Conf. on Artificial Intelligence (AAAI'12)*, Jörg Hoffmann and Bart Selman (Eds.). AAAI Press, 726–733. <http://www.aaai.org/ocs/index.php/AAAI/AAAI12/paper/view/4931>
- [21] Martine Enger, Erik Velldal, and Lilja Øvrelid. 2017. An open-source tool for negation detection: A maximum-margin approach. In *Proc. SemBEaR Workshop*, 64–69.
- [22] 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>
- [23] Yang Huang and Henry J Lowe. 2007. A novel hybrid approach to automated negation detection in clinical radiology reports. *J. Am. Med. Inform. Assn.* 14, 3 (2007), 304–311.
- [24] Evgeny Kharlamov, Dag Hovland, Martin G. Skjæveland, Dimitris Bilidas, Ernesto Jiménez-Ruiz, Guohui Xiao, Ahmet Soylu, Davide Lanti, Martin Rezk, Dmitriy Zheleznyakov, Martin Giese, Hallstein Lie, Yannis Ioannidis, Yannis Kotidis, Manolis Koubarakis, and Arild Waaler. 2017. Ontology Based Data Access in Statioil. *J. Web Semantics* 44 (2017), 3–36. <https://doi.org/10.1016/j.websem.2017.05.005>
- [25] 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.
- [26] Zhihui Luo, Meliha Yetisgen-Yildiz, and Chunhua Weng. 2011. Dynamic categorization of clinical research eligibility criteria by hierarchical clustering. *J. Biomed. Inform.* 44, 6 (2011), 927–935.
- [27] Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Proc. ALC Meeting*, 55–60.
- [28] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- [29] Krystyna Milian, Anca Bucur, and Annette Ten Teije. 2012. Formalization of clinical trial eligibility criteria: Evaluation of a pattern-based approach. In *Proc. BIBM Conference*. IEEE, 1–4.
- [30] Krystyna Milian and Annette ten Teije. 2013. Towards automatic patient eligibility assessment: From free-text criteria to queries. In *Proc. AIME Conference*. Springer, 78–83.
- [31] Marie-Laure Mugnier and Michaël Thomazo. 2014. An Introduction to Ontology-Based Query Answering with Existential Rules. In *Reasoning Web International Summer School*. 245–278. https://doi.org/10.1007/978-3-319-10587-1_6
- [32] Yizhao Ni, Jordan Wright, John Perentesis, Todd Lingren, Louise Deleger, Megan Kaiser, Isaac Kohane, and Imre Solti. 2015. Increasing the efficiency of trial-patient matching: Automated clinical trial eligibility Pre-screening for pediatric oncology patients. *BMC Medical Informatics and Decision Making* 15 (2015), 1–10. Issue 28. <https://doi.org/10.1186/s12911-015-0149-3>
- [33] 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
- [34] 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/>
- [35] Samson W Tu, Mor Peleg, Simona Carini, Michael Bobak, Jessica Ross, Daniel Rubin, and Ida Sim. 2011. A practical method for transforming free-text eligibility criteria into computable criteria. *J. Biomed. Inform.* 44, 2 (2011), 239–250.
- [36] Angela Wood, Ian White, and Simon Thompson. 2004. Are missing outcome data adequately handled? A review of published randomized controlled trials in major medical journals. *Clinical trials (London, England)* 1 (2 2004), 368–76. <https://doi.org/10.1191/1740774504cn0320a>
- [37] Chao Xu, Walter Forkel, Stefan Borgwardt, Franz Baader, and Beihai Zhou. 2019. Automatic Translation of Clinical Trial Eligibility Criteria into Formal Queries. In *Proc. of the 9th Workshop on Ontologies and Data in Life Sciences (ODLS'19), part of The Joint Ontology Workshops (JOWO'19) (CEUR Workshop Proceedings)*, Martin Boeker, Ludger Jansen, Frank Loebe, and Stefan Schulz (Eds.).
- [38] Li Zhou, Genevieve B Melton, Simon Parsons, and George Hripcsak. 2006. A temporal constraint structure for extracting temporal information from clinical narrative. *J. Biomed. Inform.* 39, 4 (2006), 424–439.