On the Power and Limitations of Examples for Description Logic Concepts (Extended Abstract)

Balder ten Cate[1], Raoul Koudijs[2] and Ana Ozaki[3,2]

[1]University of Amsterdam, ILLC - Postbus 94242, 1090 GE Amsterdam, The Netherlands
[2]University of Bergen, HIB - Thormøhlens gate 55 5006 Bergen, Norway
[3]University of Oslo, Gaustadalléen 23B 0373 Oslo, Norway

Abstract

Labeled examples (i.e., positive and negative examples) are an attractive medium for communicating complex concepts. They are useful for deriving concept expressions (such as in concept learning, interactive concept specification, and concept refinement) as well as for illustrating concept expressions to a user or domain expert. We investigate the power of labeled examples for describing description logic concepts. Specifically, we systematically study the existence and efficient computability of *finite characterizations*, i.e., finite sets of labeled examples that uniquely characterize a single concept, for a wide variety of description logics between $\mathcal{EL}$ and $\mathcal{ALCQI}$, both without an ontology and in the presence of a DL-Lite ontology. Finite characterizations are relevant for debugging purposes, and their existence is a necessary condition for exact learnability with membership queries.

Keywords

Finite Characterisations, Concept Languages, Concept Learning, Membership Queries

1. Introduction

Labeled examples (i.e., positive and negative examples) are an attractive medium for communicating complex concepts. They are useful as data for deriving concept expressions (such as in concept learning, interactive concept specification, and example-driven concept refinement) as well as for illustrating concept expressions to a user or domain expert [1, 2, 3, 4, 5, 6, 7]. In this extended abstract we report on a recent study [8] into the utility of labeled examples for describing description logic concepts.

Example 1. *In the description logic $\mathcal{EL}$, we may define the concept of an* e-bike *by means of the concept expression $C := \mathit{Bicycle} \sqcap \exists \mathit{Contains}.\mathit{Battery}$. Suppose we wish to illustrate C by a collection of positive and negative examples. What would be a good choice of examples? Take the interpretation $\mathcal{I}$ consisting of the following facts.*

$$\mathit{Bicycle}\ \mathit{soltera2} \xrightarrow{\mathit{Contains}} \mathit{li360Wh}\ \mathit{Battery}$$
$$\mathit{Bicycle}\ \mathit{px10} \xrightarrow{\mathit{Contains}} \mathit{b12}\ \mathit{Basket}$$
$$\mathit{Car}\ \mathit{teslaY} \xrightarrow{\mathit{Contains}} \mathit{li81kWh}\ \mathit{Battery}$$

In the context of this interpretation $\mathcal{I}$, it is clear that soltera2 *is a positive example for C, and* px10 *and* teslaY *are negative examples for C. In fact, as it turns out, C is the* only *$\mathcal{EL}$-concept (up to equivalence) that is consistent with, or 'fits' these three labeled examples. In other words, these three labeled examples "uniquely characterize" C within the class of all $\mathcal{EL}$-concepts. This shows that the above three labeled examples are a good choice of examples. Further examples would be redundant. Note, however, that this depends on the choice of the description logic. For instance, the richer concept language $\mathcal{ALC}$ allows for concept expressions such as $\mathit{Bicycle} \sqcap \neg\exists \mathit{Contains}.\mathit{Basket}$ that also fit.*

DL 2024: 37th International Workshop on Description Logics, June 18–21, 2024, Bergen, Norway
anaoz@uio.no (A. Ozaki)
https://www.mn.uio.no/ifi/personer/vit/anaoz/ (A. Ozaki)
0000-0002-3889-6207 (A. Ozaki)

Example 2. *This example involves $\mathcal{ALC}$-concepts over a signature (Σ_C, Σ_R) consisting of the concept names $\Sigma_R = \{Animal, Cat, Dog, Red\}$ and no role names, i.e., $\Sigma_R = \emptyset$. Let $\mathcal{O}$ furthermore be the DL-Lite ontology $\{Cat \sqsubseteq Animal, Dog \sqsubseteq Animal, Cat \sqsubseteq \neg Dog\}$ expressing that Cat and Dog are disjoint subconcepts of Animal. Consider the concept $C := Cat \sqcap Red$. If we wish to illustrate C by a collection of positive and negative examples, what would be a good choice of examples? Take the interpretation $\mathcal{I}$ consisting of the following facts:*

Animal, Red a	a' *Animal*
Red b	b'
Cat, Animal , Red c	c' *Cat, Animal*
Dog, Animal, Red d	d' *Dog, Animal*

Note that $\mathcal{I}$ satisfies the ontology $\mathcal{O}$. In the context of this interpretation, it is clear that

- *c is a positive example for C.*
- *All the other elements (that is, a, a', b, b', c', d, and d') are negative examples for C.*

As it turns out, every $\mathcal{ALC}$-concept (over the specified signature) that fits these labeled examples is equivalent to C under the ontology $\mathcal{O}$. In other words, the above labeled examples uniquely characterize C relative to $\mathcal{O}$ and the signature (Σ_C, Σ_R). On the other hand, these labeled examples do not *uniquely characterize C in the absence of an ontology, because, for instance, $Cat \sqcap Red \sqcap \neg Dog$ and $Cat \sqcap Red \sqcap Animal$ also fit but are not equivalent to C in the absence of the ontology.*

An ontology can help reduce the number of examples needed to uniquely characterize C. Moreover, it can help avoid unnatural examples. For instance, without an ontology, a unique characterization for C would need to include negative examples satisfying $Cat \sqcap Red \sqcap Dog$ and $Cat \sqcap Red \sqcap \neg Animal$.

Motivated by the above examples, we investigate the existence and efficient computability of *finite characterizations*, i.e., finite sets of labeled examples that uniquely characterize a single concept. Finite characterizations are relevant not only for *illustrating* a complex concept to a user (e.g., to verify the correctness of a concept expression obtained using machine learning techniques). Their existence is a necessary condition for *exact learnability with membership queries* [9]. Furthermore, from a more fundamental point of view, by studying the existence of finite characterizations, we gain insight into the power and limitations of labeled examples as a medium for describing concepts.

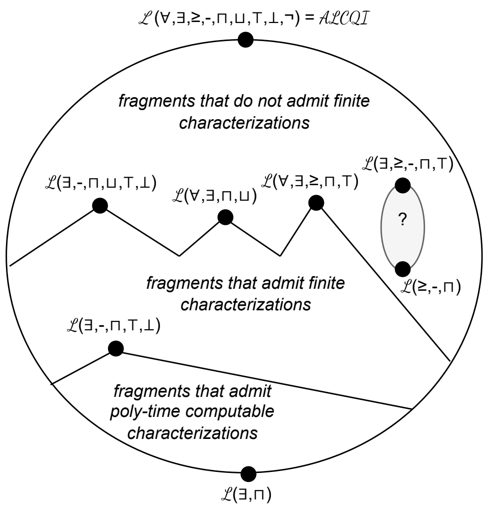

Figure 1: Summary of Thm. 1 and 2.

Finite characterisations were first studied in the computational learning theory literature under the name of teaching sets, with a corresponding notion of teaching dimension, measuring the maximal size of minimal teaching sets of some class of concepts [10]. Several recent works study finite characterisations for *description logic concepts* ([11, 12] for $\mathcal{ELI}$; [13] for temporal instance queries); a systematic study of finite characterizations for syntactic fragments of *modal logic* was carried out in [14].

In this prior literature, two types of examples have been considered, namely *open-world examples* (cf. [12, 15]) and *closed-world examples* (e.g. [13, 16]). An *open-world example* is a pair $(\mathcal{A}, a)$, where $\mathcal{A}$ is an ABox and a is an individual name. Such a pair is a *positive example* for a concept expression C relative to an ontology $\mathcal{O}$ if a belongs to the certain answers of C w.r.t. $\mathcal{A}$ and it is a *negative example* otherwise.

In contrast, a *closed-world example* is a pair $(\mathcal{I}, d)$ where $\mathcal{I}$ is a finite interpretation satisfying the ontology $\mathcal{O}$, and d is an element in the domain of $\mathcal{I}$. Such a pair is a *positive example* for a concept expression C if $d \in C^{\mathcal{I}}$, and it is a *negative example* otherwise.

2. Contributions

First contribution: We study the difference in descriptive power between closed-world examples and open-world examples for description logic concepts. It was shown in [12] that $\mathcal{ELI}$ admits finite characterizations with open-world examples under DL-Lite ontologies. We show that the same does not hold with closed-world examples. On the other hand, for non-monotone concept languages, e.g. with $\forall$ or $\neg$, it is necessary to work with closed-world rather than open-world examples. For instance, the concepts $\forall R.A$ and $\forall R.B$ do not have *any* positive open-world examples (relative to an empty ontology), and hence neither can be uniquely characterized by open-world examples.

Second contribution: We systematically study the existence of, and the complexity of computing, finite characterisations for concept expressions in a wide range of description logics, using *closed-world examples*. Specifically, we look at concept languages $\mathcal{L}(\mathbb{O})$ generated by a set of connectives $\mathbb{O}$, where $\{\exists, \sqcap\} \subseteq \mathbb{O} \subseteq \{\forall, \exists, \geq, -, \sqcap, \sqcup, \top, \bot, \neg\}$. In other words, we look at fragments of the description logic $\mathcal{ALCQI}$ that contain at least the $\exists$ and $\sqcap$ constructors from $\mathcal{EL}$. Within this framework, we obtain an almost complete classification of the concept languages that admit finite characterizations.

Theorem 1. *Let* $\{\exists, \sqcap\} \subseteq \mathbb{O} \subseteq \{\forall, \exists, \geq, -, \sqcap, \sqcup, \top, \bot, \neg\}$.

1. *If* $\mathbb{O}$ *is a subset of* $\{\exists, -, \sqcap, \sqcup, \top, \bot\}$, $\{\forall, \exists, \sqcup, \sqcap\}$ *or* $\{\forall, \exists, \geq, \sqcap, \top\}$ *then* $\mathcal{L}(\mathbb{O})$ *admits finite characterizations with closed-world examples.*
2. *Otherwise* $\mathcal{L}(\mathbb{O})$ *does not admit finite characterizations with closed-world examples, except possibly if* $\{\geq, -, \sqcap\} \subseteq \mathbb{O} \subseteq \{\exists, \geq, -, \sqcap, \top\}$.

The above theorem identifies three maximal fragments of $\mathcal{ALCQI}$ that admit finite characterisations. It leaves open the status of essentially only two concept languages, namely $\mathcal{L}(\exists, \geq, -, \sqcap)$ and $\mathcal{L}(\exists, \geq, -, \sqcap, \top)$ (note that $\exists$ is definable in terms of $\geq$). The proof builds on prior results from [11] and [14]. Our main novel technical contributions are a construction of finite characterisations for $\mathcal{L}(\forall, \exists, \geq, \sqcap, \top)$ and complementary negative results for $\mathcal{L}(\geq, \bot)$, $\mathcal{L}(\geq, \sqcup)$ and $\mathcal{L}(\forall, \exists, -, \sqcap)$. The construction of finite characterisations for $\mathcal{L}(\forall, \exists, \geq, \sqcap, \top)$ is non-elementary. We give an elementary (doubly exponential) construction for $\mathcal{L}(\exists, \geq, \sqcap, \top)$, via a novel polynomial-time algorithm for constructing *frontiers* (i.e., complete sets of minimal weakenings) of $\mathcal{L}(\exists, \geq, \sqcap, \top)$-concepts. It also follows that subsumptions between such concepts can be checked in polynomial time.

Next, we study which concept languages $\mathcal{L}(\mathbb{O})$ *admit polynomial-time computable characterisations*, i.e., have a polynomial-time algorithm to compute finite characterisations.

Theorem 2. *Let* $\{\exists, \sqcap\} \subseteq \mathbb{O} \subseteq \{\forall, \exists, \geq, -, \sqcap, \sqcup, \top, \bot, \neg\}$.

1. *If* $\mathbb{O}$ *is a subset of* $\{\exists, -, \sqcap, \top, \bot\}$, *then* $\mathcal{L}(\mathbb{O})$ *admits polynomial-time computable characterizations with closed-world examples.*
2. *Otherwise,* $\mathcal{L}(\mathbb{O})$ *does not admit polynomial-time computable characterizations with closed-world examples, assuming P≠NP.*

The first item follows from known results [11]; we prove the second by establishing an exponential lower bound on characterisations for $\mathcal{L}(\exists, \geq, \sqcap)$ and NP-hardness of computing characterisations for $\mathcal{L}(\forall, \exists, \sqcap)$. Theorems 1 and 2 are summarized in Figure 1. Theorems 1 and 2 still holds true when 'admits finite (polynomial) time computable characterisations' is replaced with 'is exact learnable from membership queries in finite (polynomial) time'. On the one hand, existence finite (polynomial) characterisations is a necessary condition for (polynomial) membership query learnability, so the negative results transfer immediately. On the other hand, finite characterisations always enable a brute force membership query algorithm, and [11] shows that $\mathcal{L}(\exists, -, \sqcap, \top)$ is polynomial-time learnable from membership queries.

Finally, we investigate finite characterizations relative to an ontology $\mathcal{O}$, where we require all the (positive and negative) examples to be interpretations satisfying the ontology $\mathcal{O}$, and require that every fitting concept is equivalent to the input concept C relative to $\mathcal{O}$ (cf. Example 2).

Theorem 3. *Let* $\{\exists, \sqcap\} \subseteq \mathbb{O} \subseteq \{\forall, \exists, \geq, -, \sqcap, \sqcup, \top, \bot, \neg\}$.

1. *If* $\mathbb{O}$ *is a subset of* $\{\exists, \sqcap, \top, \bot\}$, *then* $\mathcal{L}(\mathbb{O})$ *admits finite characterizations with closed-world examples w.r.t. all DL-Lite ontologies.*
2. *Otherwise,* $\mathcal{L}(\mathbb{O})$ *does not admit finite characterizations with closed-world examples w.r.t. all DL-Lite ontologies, except possibly if* $\{\exists, \sqcap, \sqcup\} \subseteq \mathbb{O} \subseteq \{\exists, \sqcap, \sqcup, \top, \bot\}$.

In fact, for $\mathcal{L}(\exists, \sqcap, \top, \bot)$*-concepts* C *and DL-Lite ontologies* $\mathcal{O}$ *such that* C *is satisfiable w.r.t.* $\mathcal{O}$*, a finite characterisation can be computed in polynomial time.*

Acknowledgements

Balder ten Cate is supported by the European Union's Horizon 2020 research and innovation programme under grant MSCA-101031081. Raoul Koudijs and Ana Ozaki are supported by the Research Council of Norway, project number 316022, led by Ana Ozaki.

References

[1] J. Lehmann, Dl-learner: Learning concepts in description logics, Journal of Machine Learning Research 10 (2009) 2639–2642.

[2] M. Funk, J. C. Jung, C. Lutz, H. Pulcini, F. Wolter, Learning description logic concepts: When can positive and negative examples be separated?, in: IJCAI, 2019, pp. 1682–1688.

[3] M. Funk, J. C. Jung, C. Lutz, Actively learning concept and conjunctive queries under $\mathcal{EL}^r$-ontologies, in: IJCAI, 2021, pp. 1887–1893.

[4] A. Ozaki, Learning description logic ontologies: Five approaches. where do they stand?, Künstliche Intell. 34 (2020) 317–327. URL: https://doi.org/10.1007/s13218-020-00656-9. doi:10.1007/S13218-020-00656-9.

[5] N. Fanizzi, C. d'Amato, F. Esposito, DL-FOIL concept learning in description logics, in: Proc. of ILP, 2008, pp. 107–121.

[6] L. Iannone, I. Palmisano, N. Fanizzi, An algorithm based on counterfactuals for concept learning in the semantic web, Appl. Intell. 26 (2007) 139–159.

[7] J. Lehmann, P. Hitzler, Concept learning in description logics using refinement operators, Mach. Learn. 78 (2010) 203–250.

[8] B. ten Cate, R. Koudijs, A. Ozaki, On the power and limitations of examples for description logic concepts, in: Proc. of IJCAI, 2024, pp. 739–745.

[9] D. Angluin, Queries and concept learning, Mach. Learn. 2 (1987) 319–342.

[10] S. Goldman, M. Kearns, On the complexity of teaching, Journal of Computer and System Sciences 50 (1995) 20–31. URL: https://www.sciencedirect.com/science/article/pii/S0022000085710033. doi:https://doi.org/10.1006/jcss.1995.1003.

[11] B. ten Cate, V. Dalmau, Conjunctive queries: Unique characterizations and exact learnability, ACM Trans. Database Syst. 47 (2022). URL: https://doi.org/10.1145/3559756. doi:10.1145/3559756.

[12] M. Funk, J. C. Jung, C. Lutz, Frontiers and exact learning of ELI queries under dl-lite ontologies, in: IJCAI, 2022, pp. 2627–2633.

[13] M. Fortin, B. Konev, V. Ryzhikov, Y. Savateev, F. Wolter, M. Zakharyaschev, Unique Characterisability and Learnability of Temporal Instance Queries, in: Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning, 2022, pp. 163–173. URL: https://doi.org/10.24963/kr.2022/17. doi:10.24963/kr.2022/17.

[14] B. t. Cate, R. Koudijs, Characterising modal formulas with examples, ACM Trans. Comput. Logic (2024). URL: https://doi.org/10.1145/3649461. doi:10.1145/3649461, just Accepted.

[15] J. C. Jung, V. Ryzhikov, F. Wolter, M. Zakharyaschev, Temporalising unique characterisability and learnability of ontology-mediated queries (extended abstract), in: O. Kutz, C. Lutz, A. Ozaki (Eds.), Proceedings of the 36th International Workshop on Description Logics (DL 2023) co-located with the 20th International Conference on Principles of Knowledge Representation and Reasoning and the 21st International Workshop on Non-Monotonic Reasoning (KR 2023 and NMR 2023)., Rhodes, Greece, September 2-4, 2023, volume 3515 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023. URL: https://ceur-ws.org/Vol-3515/abstract-13.pdf.

[16] L. De Raedt, Logical settings for concept-learning, Artificial Intelligence 95 (1997) 187–201. URL: https://www.sciencedirect.com/science/article/pii/S0004370297000416. doi:https://doi.org/10.1016/S0004-3702(97)00041-6.