

Evonne: Interactive Proof Visualization for Description Logics (System Description)

Christian Alrabbaa¹ , Franz Baader¹ , Stefan Borgwardt¹ ,
Raimund Dachzelt² , Patrick Koopmann¹ , and Julián Méndez² 

¹ Institute of Theoretical Computer Science, TU Dresden, Germany

² Interactive Media Lab Dresden, TU Dresden, Germany

`firstname.lastname@tu-dresden.de`, `julian.mendez2@tu-dresden.de`

Abstract. Explanations for description logic (DL) entailments provide important support for the maintenance of large ontologies. The “justifications” usually employed for this purpose in ontology editors pinpoint the parts of the ontology responsible for a given entailment. Proofs for entailments make the intermediate reasoning steps explicit, and thus explain how a consequence can actually be derived. We present an interactive system for exploring description logic proofs, called EVONNE, which visualizes proofs of consequences for ontologies written in expressive DLs. We describe the methods used for computing those proofs, together with a feature called *signature-based proof condensation*. Moreover, we evaluate the quality of generated proofs using real ontologies.

1 Introduction

Proofs generated by Automated Reasoning (AR) systems are sometimes presented to humans in textual form to convince them of the correctness of a theorem [9, 11], but more often employed as certificates that can automatically be checked [20]. In contrast to the AR setting, where very long proofs may be needed to derive a deep mathematical theorem from very few axioms, DL-based ontologies are often very large, but proofs of a single consequence are usually of a more manageable size. For this reason, the standard method of explanation in description logic [8] has long been to compute so-called *justifications*, which point out a minimal set of source statements responsible for an entailment of interest. For example, the ontology editor Protégé³ supports the computation of justifications since 2008 [12], which is very useful when working with large DL ontologies. Nevertheless, it is often not obvious why a given consequence actually follows from such a justification [13]. Recently, this explanation capability has been extended towards showing full *proofs* with intermediate reasoning steps, but this is restricted to ontologies written in the lightweight DLs supported by the ELK reasoner [15, 16], and the graphical presentation of proofs is very basic.

In this paper, we present EVONNE as an interactive system, for exploring DL proofs for description logic entailments, using the methods for computing small

³ <https://protege.stanford.edu/>

proofs presented in [3,5]. Initial prototypes of EVONNE were presented in [6,10], but since then, many improvements were implemented. While EVONNE does more than just visualizing proofs, this paper focuses on the proof component of EVONNE: specifically, we give a brief overview of the interface for exploring proofs, describe the proof generation methods implemented in the back-end, and present an experimental evaluation of these proofs generation methods in terms of proof size and run time. The improved back-end uses Java libraries that extract proofs using various methods, such as from the ELK calculus, or *forgetting-based proofs* [3] using the forgetting tools LETHE [17] and FAME [21] in a black-box fashion. The new front-end is visually more appealing than the prototypes presented in [6,10], and allows to inspect and explore proofs using various interaction techniques, such as zooming and panning, collapsing and expanding, text manipulation, and compactness adjustments. Additional features include the minimization of the generated proofs according to various measures and the possibility to select a *known signature* that is used to automatically hide parts of the proofs that are assumed to be obvious for users with certain previous knowledge. Our evaluation shows that proof sizes can be significantly reduced in this way, making the proofs more user-friendly. EVONNE can be tried and downloaded at <https://imld.de/evonne>. The version of EVONNE described here, as well as the data and scripts used in our experiments, can be found at [2].

2 Preliminaries

We recall some relevant notions for DLs; for a detailed introduction, see [8]. DLs are decidable fragments of first-order logic (FOL) with a special, variable-free syntax, and that use only unary and binary predicates, called *concept names* and *role names*, respectively. These can be used to build complex *concepts*, which correspond to first-order formulas with one free variable, and *axioms* corresponding to first-order sentences. Which kinds of concepts and axioms can be built depends on the expressivity of the used DL. Here we mainly consider the light-weight DL \mathcal{ELH} and the more expressive \mathcal{ALCH} . We have the usual notion of FOL *entailment* $\mathcal{O} \models \alpha$ of an axiom α from a finite set of axioms \mathcal{O} , called an ontology. Of special interest are entailments of *atomic CIs* (concept inclusions) of the form $A \sqsubseteq B$, where A and B are concept names. Following [3], we define *proofs* of $\mathcal{O} \models \alpha$ as finite, acyclic, directed hypergraphs, where vertices v are labeled with axioms $\ell(v)$ and hyperedges are of the form (S, d) , with S a set of vertices and d a vertex such that $\{\ell(v) \mid v \in S\} \models \ell(d)$; the leaves of a proof must be labeled by elements of \mathcal{O} and the root by α . In this paper, all proofs are *trees*, i.e. no vertex can appear in the first component of multiple hyperedges (see Fig. 1).

3 The Graphical User Interface

The user interface of EVONNE is implemented as a web application. To support users in understanding large proofs, they are offered various layout options and

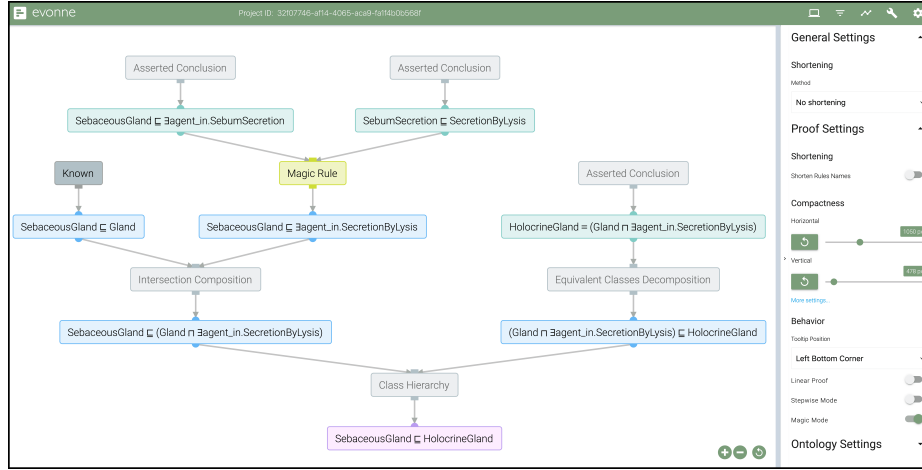


Fig. 1: Overview of EVONNE - a condensed proof in the bidirectional layout

interaction components. The proof visualization is linked to a second view showing the context of the proof in a relevant subset of the ontology. In this ontology view, interactions between axioms are visualized, so that users can understand the context of axioms occurring in the proof. The user can also examine possible ways to eliminate unwanted entailments in the ontology view. The focus of this system description, however, is on the proof component: we describe how the proofs are generated and how users can interact with the proof visualization. For details on the ontology view, we refer the reader to the workshop paper [6], where we also describe how EVONNE supports ontology repair.

Initialization. After starting EVONNE for the first time, users create a new project, for which they specify an ontology file. They can then select an entailed atomic CI to be explained. The user can choose between different proof methods, and optionally select a signature of *known terms* (cf. Sec. 4), which can be generated using the term selection tool Protégé-TS [14].

Layout. Proofs are shown as graphs with two kinds of vertices: colored vertices for axioms, gray ones for inference steps. By default, proofs are shown using a *tree layout*. To take advantage of the width of the display when dealing with long axioms, it is possible to show proofs in a *vertical layout*, placing axioms linearly below each other, with inferences represented through edges on the side (without the inference vertices). It is possible to automatically re-order vertices to minimize the distance between conclusion and premises in each step. The third layout option is the *bidirectional layout* (see Fig. 1), a tree layout where, initially, the entire proof is collapsed into a *magic vertex* that links the conclusion directly to its justification, and from which individual inference steps can be pulled out and pushed back from both directions.

Exploration. In all views, each vertex is equipped with multiple functionalities for exploring a proof. For proofs generated with ELK, clicking on an inference vertex shows the inference rule used, and the particular inference with relevant sub-elements highlighted in different colors. Axiom vertices show different button (\uparrow , \downarrow , \leftarrow , \odot) when hovered over. In the standard tree layout, users can hide sub-proofs under an axiom (\downarrow). They can also reveal the previous inference step (\leftarrow) or the entire sub-proof (\uparrow). In the vertical layout, the button (\odot) highlights and explains the inference of the current axiom. In the bidirectional layout, the arrow buttons are used for pulling inference steps out of the magic vertex, as well as pushing them back in.

Presentation. A *minimap* allows users to keep track of the overall structure of the proof, thus enriching the zooming and panning functionality. Users can adjust width and height of proofs through the options side-bar. Long axiom labels can be *shortened* in two ways: either by setting a fixed size to all vertices, or by abbreviating names based on capital letters. Afterwards, it is possible to restore the original labels individually.

4 Proof Generation

To obtain the proofs that are shown to the user, we implemented different proof generation techniques, some of which were initially described in [3]. For \mathcal{ELH} ontologies, proofs can be generated natively by the DL reasoner ELK [16]. These proofs use rules from the calculus described in [16]. We apply the Dijkstra-like algorithm introduced in [4, 5] to compute a *minimized proof* from the ELK output. This minimization can be done w.r.t. different measures, such as the size, depth, or weighted sum (where each axiom is weighted by its size), as long as they are *monotone* and *recursive* [5]. For ontologies outside of the \mathcal{ELH} fragment, we use the forgetting-based approach originally described in [3], for which we now implemented two alternative algorithms for computing more compact proofs (Sec. 4.1). Finally, independently of the proof generation method, one can specify a signature of known terms. This signature contains terminology that the user is familiar with, so that entailments using only those terms do not need to be explained. The condensation of proofs w.r.t. signatures is described in Sec. 4.2.

4.1 Forgetting-Based Proofs

In a forgetting-based proof, proof steps represent inferences on concept or role names using a *forgetting* operation. Given an ontology \mathcal{O} and a predicate name x , the result \mathcal{O}^{-x} of forgetting x in \mathcal{O} does not contain any occurrences of x , while still capturing all entailments of \mathcal{O} that do not use x [18]. In a forgetting-based proof, an inference takes as premises a set \mathcal{P} of axioms and has as conclusion some axiom $\alpha \in \mathcal{P}^{-x}$ (where a particular forgetting operation is used to compute \mathcal{P}^{-x}). Intuitively, α is obtained from \mathcal{P} by performing inferences on x . To compute a forgetting-based proof, we have to forget the names occurring in the

ontology one after the other, until only the names occurring in the statement to be proved are left. For the forgetting operation, the user can select between two implementations: **LETHE** [17] (using the method supporting \mathcal{ALCH}) and **FAME** [21] (using the method supporting \mathcal{ALCOT}). Since the space of possible inference steps is exponentially large, it is not feasible to minimize proofs after their computation, as we do for \mathcal{EL} entailments, which is why we rely on heuristics and search algorithms to generate small proofs. Specifically, we implemented three methods for computing forgetting-based proofs: **HEUR** tries to find proofs fast, **SYMB** tries to minimize the number of predicates forgotten in a proof, with the aim of obtaining proofs of small depth, and **SIZE** tries to optimize the size of the proof. The heuristic method **HEUR** is described in [3], and its implementation has not been changed since then. The search methods **SYMB** and **SIZE** are new (details can be found in the extended version [1]).

4.2 Signature-Based Proof Condensation

When inspecting a proof over a real-world ontology, different parts of the proof will be more or less familiar to the user, depending on their knowledge about the involved concepts or their experience with similar inference steps in the past. For CIs between concepts for which a user has application knowledge, they may not need to see a proof, and consequently, sub-proofs for such axioms can be automatically hidden. We assume that the user’s knowledge is given in the form of a *known signature* Σ and that axioms that contain only symbols from Σ do not need to be explained. The effect can be seen in Fig. 1 through the “known”-inference on the left, where Σ contains **SebaceousGland** and **Gland**. The known signature is taken into consideration when minimizing the proofs, so that proofs are selected for which more of the known information can be used if convenient. This can be easily integrated into the Dijkstra approach described in [3], by initially assigning to each axiom covered by Σ a proof with a single vertex.

5 Evaluation

For EVONNE to be usable in practice, it is vital that proofs are computed efficiently and that they are not too large. An experimental evaluation of minimized proofs for \mathcal{EL} and forgetting-based proofs obtained with **FAME** and **LETHE** is provided in [3]. We here present an evaluation of additional aspects: 1) a comparison of the three methods for computing forgetting-based proofs, and 2) an evaluation on the impact of signature-based proof condensation. All experiments were performed on Debian Linux (Intel Core i5-4590, 3.30 GHz, 23 GB Java heap size).

5.1 Minimal Forgetting-Based Proofs

To evaluate forgetting-based proofs, we extracted \mathcal{ALCH} “proof tasks” from the ontologies in the 2017 snapshot of BioPortal [19]. We restricted all ontologies to \mathcal{ALCH} and collected all entailed atomic CIs α , for each of which we computed

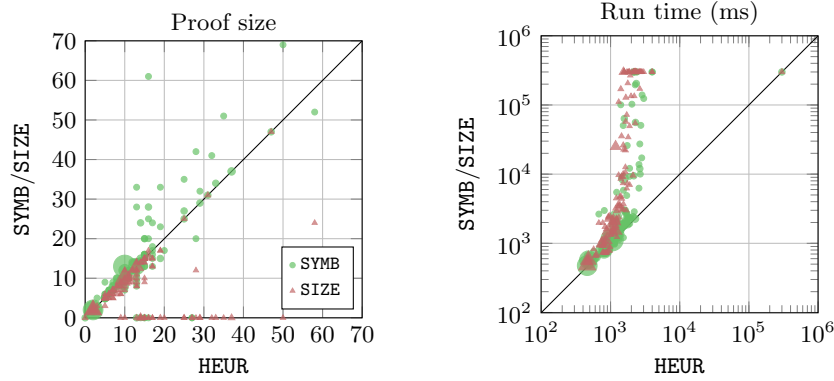


Fig. 2: Run times and proof sizes for different forgetting-based proof methods. Marker size indicates how often each pattern occurred in the BioPortal snapshot. Instances that timed out were assigned size 0.

the union \mathcal{U} of all their justifications. We identified pairs (α, \mathcal{U}) that were isomorphic modulo renaming of predicates, and kept only those patterns (α, \mathcal{U}) that contained at least one axiom not expressible in \mathcal{ELH} . This was successful in 373 of the ontologies⁴ and resulted in 138 distinct *justification patterns* (α, \mathcal{U}) , representing 327 different entailments in the BioPortal snapshot. We then computed forgetting-based proofs for $\mathcal{U} \models \alpha$ with our three methods using LETHE, with a 5-minute timeout. This was successful for 325/327 entailments for the heuristic method (HEUR), 317 for the symbol-minimizing method (SYMB), and 279 for the size-minimizing method (SIZE). In Fig. 2 we compare the resulting *proof sizes* (left) and the *run times* (right), using HEUR as baseline (x-axis). HEUR is indeed faster in most cases, but SIZE reduces proof size by 5% on average compared to HEUR, which is not the case for SYMB. Regarding *proof depth* (not shown in the figure), SYMB did not outperform HEUR on average, while SIZE surprisingly yielded an average reduction of 4% compared to HEUR. Despite this good performance of SIZE for proof size and depth, for entailments that depend on many or complex axioms, computation times for both SYMB and SIZE become unacceptable, while proof generation with HEUR mostly stays in the area of seconds.

5.2 Signature-Based Proof Condensation

To evaluate how much hiding proof steps in a known signature decreases proof size in practice, we ran experiments on the large medical ontology SNOMED CT (International Edition, July 2020) that is mostly formulated in \mathcal{ELH} .⁵ As signatures we used SNOMED CT *Reference Sets*,⁶ which are restricted vocabularies

⁴ The other ontologies could not be processed in this way within the memory limit.

⁵ <https://www.snomed.org/>

⁶ <https://confluence.ihtsdotools.org/display/DOCRFSPG/2.3.+Reference+Set>

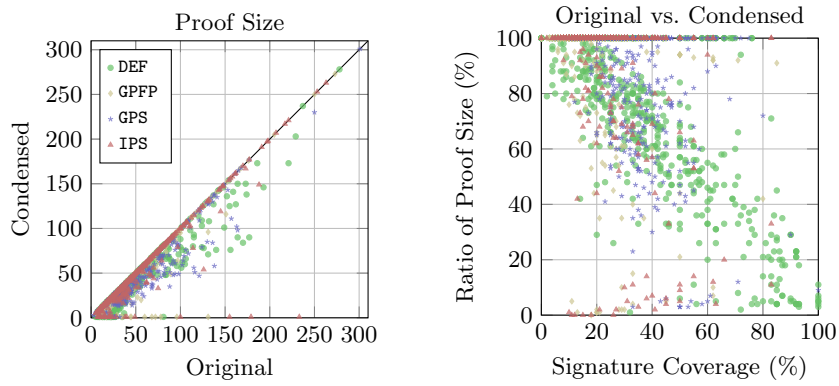


Fig. 3: Size of original and condensed proofs (left). Ratio of proof size depending on the signature coverage (right).

for specific use cases. We extracted justifications similarly to the previous experiment, but did not rename predicates and considered only proof tasks that use at least 5 symbols from the signature, since otherwise no improvement can be expected by using the signatures. For each signature, we randomly selected 500 out of 6.689.452 *proof tasks* (if at least 500 existed). This left the 4 reference sets *General Practitioner/Family Practitioner* (GPFP), *Global Patient Set* (GPS), *International Patient Summary* (IPS), and the one included in the SNOMED CT distribution (DEF). For each of the resulting 2.000 proof tasks, we used ELK [16] and our proof minimization approach to obtain (a) a proof of minimal size and (b) a proof of minimal size after hiding the selected signature. The distribution of proof sizes can be seen in Fig. 3. In 770/2.000 cases, a smaller proof was generated when using the signature. In 91 of these cases, the size was even be reduced to 1, i.e. the target axiom used only the given signature and therefore nothing else needed to be shown. In the other 679 cases with reduced size, the average *ratio* of reduced size to original size was 0.68–0.93 (depending on the signature). One can see that this ratio is correlated with the *signature coverage* of the original proof (i.e. the ratio of signature symbols to total symbols in the proof), with a weak or strong correlation depending on the signature (r between -0.26 and -0.74). However, a substantial number of proofs with relatively high signature coverage could still not be reduced in size at all (see the top right of the right diagram). In summary, we can see that signature-based condensation can be useful, but this depends on the proof task and the signature. We also conducted experiments on the Galen ontology,⁷ with comparable results (see the extended version of this paper [1]).

⁷ <https://biportal.bioontology.org/ontologies/GALEN>

6 Conclusion

We have presented and compared the proof generation and presentation methods used in EVONNE, a visual tool for explaining entailments of DL ontologies. While these methods produce smaller or less deep proofs, which are thus easier to present, there is still room for improvements. Specifically, as the forgetting-based proofs do not provide the same degree of detail as the ELK proofs, it would be desirable to also support methods for more expressive DLs that generate proofs with smaller inference steps. Moreover, our current evaluation focuses on proof size and depth—to understand how well EVONNE helps users to understand DL entailments, we would also need a qualitative evaluation of the tool with potential end-users. We are also working on explanations for non-entailments using countermodels [7] and a plugin for the ontology editor Protégé that is compatible with the PULi library and Proof Explanation plugin presented in [15], which will support all proof generation methods discussed here and more.⁸

Acknowledgements This work was supported by the German Research Foundation (DFG) in Germany’s Excellence Strategy: EXC-2068, 390729961 – Cluster of Excellence “Physics of Life” and EXC 2050/1, 390696704 – Cluster of Excellence “Centre for Tactile Internet” (CeTI) of TU Dresden, by DFG grant 389792660 as part of TRR 248 – CPEC, by the AI competence center ScaDS.AI Dresden/Leipzig, and the DFG Research Training Group QuantLA, GRK 1763.

References

1. Alrabbaa, C., Baader, F., Borgwardt, S., Dachzelt, R., Koopmann, P., Méndez, J.: Evonne: Interactive Proof Visualization for Description Logics (System Description) – Extended Version (2022). <https://doi.org/10.48550/ARXIV.2205.09583>
2. Alrabbaa, C., Baader, F., Borgwardt, S., Dachzelt, R., Koopmann, P., Méndez, J.: Evonne: Interactive Proof Visualization for Description Logics (System Description) - IJCAR22 - Resources (May 2022). <https://doi.org/10.5281/zenodo.6560603>
3. Alrabbaa, C., Baader, F., Borgwardt, S., Koopmann, P., Kovtunova, A.: Finding small proofs for description logic entailments: Theory and practice. In: Albert, E., Kovács, L. (eds.) Proceedings of the 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR 2020). EPiC Series in Computing, vol. 73, pp. 32–67. EasyChair (2020). <https://doi.org/10.29007/nhpp>
4. Alrabbaa, C., Baader, F., Borgwardt, S., Koopmann, P., Kovtunova, A.: On the complexity of finding good proofs for description logic entailments. In: Borgwardt, S., Meyer, T. (eds.) Proceedings of the 33rd International Workshop on Description Logics (DL 2020). CEUR Workshop Proceedings, vol. 2663. CEUR-WS.org (2020), <http://ceur-ws.org/Vol-2663/paper-1.pdf>
5. Alrabbaa, C., Baader, F., Borgwardt, S., Koopmann, P., Kovtunova, A.: Finding good proofs for description logic entailments using recursive quality measures. In: Platzer, A., Sutcliffe, G. (eds.) Proceedings of the 28th International Conference on Automated Deduction (CADE-28). Lecture Notes in Computer Science, vol. 12699, pp. 291–308. Springer (2021). https://doi.org/10.1007/978-3-030-79876-5_17

⁸ <https://github.com/de-tu-dresden-inf-lat/evee>

6. Alrabbaa, C., Baader, F., Dachzelt, R., Flemisch, T., Koopmann, P.: Visualising proofs and the modular structure of ontologies to support ontology repair. In: Borgwardt, S., Meyer, T. (eds.) Proceedings of the 33rd International Workshop on Description Logics (DL 2020). CEUR Workshop Proceedings, vol. 2663. CEUR-WS.org (2020), <http://ceur-ws.org/Vol-2663/paper-2.pdf>
7. Alrabbaa, C., Hieke, W., Turhan, A.: Counter model transformation for explaining non-subsumption in \mathcal{EL} . In: Beierle, C., Ragni, M., Stolzenburg, F., Thimm, M. (eds.) Proceedings of the 7th Workshop on Formal and Cognitive Reasoning. CEUR Workshop Proceedings, vol. 2961, pp. 9–22. CEUR-WS.org (2021), <http://ceur-ws.org/Vol-2961/paper-2.pdf>
8. Baader, F., Horrocks, I., Lutz, C., Sattler, U.: An Introduction to Description Logic. Cambridge University Press (2017). <https://doi.org/10.1017/9781139025355>
9. Fiedler, A.: Natural language proof explanation. In: Mechanizing Mathematical Reasoning, Essays in Honor of Jörg H. Siekmann on the Occasion of His 60th Birthday. pp. 342–363. Springer (2005). https://doi.org/10.1007/978-3-540-32254-2_20
10. Flemisch, T., Langner, R., Alrabbaa, C., Dachzelt, R.: Towards designing a tool for understanding proofs in ontologies through combined node-link diagrams. In: Ivanova, V., Lambrix, P., Pesquita, C., Wiens, V. (eds.) Proceedings of the Fifth International Workshop on Visualization and Interaction for Ontologies and Linked Data (VOILA 2020). CEUR Workshop Proceedings, vol. 2778, pp. 28–40. CEUR-WS.org (2020), <http://ceur-ws.org/Vol-2778/paper3.pdf>
11. Horacek, H.: Presenting proofs in a human-oriented way. In: Ganzinger, H. (ed.) 16th International Conference on Automated Deduction (CADE-16). Lecture Notes in Computer Science, vol. 1632, pp. 142–156. Springer (1999). https://doi.org/10.1007/3-540-48660-7_10
12. Horridge, M., Parsia, B., Sattler, U.: Explanation of OWL entailments in Protege 4. In: Bizer, C., Joshi, A. (eds.) Proceedings of the Poster and Demonstration Session at the 7th International Semantic Web Conference (ISWC 2008). CEUR Workshop Proceedings, vol. 401. CEUR-WS.org (2008), http://ceur-ws.org/Vol-401/iswc2008pd_submission_47.pdf
13. Horridge, M., Parsia, B., Sattler, U.: Justification oriented proofs in OWL. In: Patel-Schneider, P.F., Pan, Y., Hitzler, P., Mika, P., Zhang, L., Pan, J.Z., Horrocks, I., Glimm, B. (eds.) Proc. of the 9th International Semantic Web Conference (ISWC 2010). Lecture Notes in Computer Science, vol. 6496, pp. 354–369. Springer (2010). https://doi.org/10.1007/978-3-642-17746-0_23
14. Hyland, I., Schmidt, R.A.: Protégé-TS: An OWL ontology term selection tool. In: Borgwardt, S., Meyer, T. (eds.) Proceedings of the 33rd International Workshop on Description Logics (DL 2020). CEUR Workshop Proceedings, vol. 2663. CEUR-WS.org (2020), <http://ceur-ws.org/Vol-2663/paper-12.pdf>
15. Kazakov, Y., Klinov, P., Stupnikov, A.: Towards reusable explanation services in protege. In: Artale, A., Glimm, B., Kontchakov, R. (eds.) Proceedings of the 30th International Workshop on Description Logics (DL 2017). CEUR Workshop Proceedings, vol. 1879. CEUR-WS.org (2017), <http://ceur-ws.org/Vol-1879/paper31.pdf>
16. Kazakov, Y., Krötzsch, M., Simancik, F.: The incredible ELK - from polynomial procedures to efficient reasoning with \mathcal{EL} ontologies. J. Autom. Reason. **53**(1), 1–61 (2014). <https://doi.org/10.1007/s10817-013-9296-3>
17. Koopmann, P.: LETHE: Forgetting and uniform interpolation for expressive description logics. Künstliche Intell. **34**(3), 381–387 (2020). <https://doi.org/10.1007/s13218-020-00655-w>

18. Koopmann, P., Schmidt, R.A.: Forgetting concept and role symbols in \mathcal{ALCH} -ontologies. In: McMillan, K.L., Middeldorp, A., Voronkov, A. (eds.) Proc. of the 19th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR-19). Lecture Notes in Computer Science, vol. 8312, pp. 552–567. Springer (2013). https://doi.org/10.1007/978-3-642-45221-5_37
19. Matentzoglou, N., Parsia, B.: Bioportal snapshot 30.03.2017 (Mar 2017). <https://doi.org/10.5281/zenodo.439510>
20. Reger, G., Suda, M.: Checkable proofs for first-order theorem proving. In: Reger, G., Traytel, D. (eds.) 1st International Workshop on Automated Reasoning: Challenges, Applications, Directions, Exemplary Achievements (ARCADE 2017). EPiC Series in Computing, vol. 51, pp. 55–63. EasyChair (2017). <https://doi.org/10.29007/s6d1>
21. Zhao, Y., Schmidt, R.A.: FAME: an automated tool for semantic forgetting in expressive description logics. In: Galmiche, D., Schulz, S., Sebastiani, R. (eds.) Proceedings of the 9th International Joint Conference on Automated Reasoning (IJCAR 2018). Lecture Notes in Computer Science, vol. 10900, pp. 19–27. Springer (2018). https://doi.org/10.1007/978-3-319-94205-6_2