Towards a Thorough Understanding of ABox Abduction Under Repair Semantics

Anselm Haak[1], Patrick Koopmann[2], Yasir Mahmood[3] and Anni-Yasmin Turhan[1]

[1]Knowledge Representation Group, Paderborn University, Germany
[2]Knowledge in Artificial Intelligence, Vrije Universiteit Amsterdam, The Netherlands
[3]Data Science Group, Paderborn University, Germany

Abstract

Abduction is the task of computing a sufficient extension of a knowledge base (KB) that entails a conclusion not entailed by the original KB. It serves to compute explanations, or *hypotheses*, for such missing entailments. Little is known about abduction when erroneous data results in inconsistent KBs. In this paper we present abduction under repair semantics, and discuss combined complexity results on deciding existence of and verifying abductive solutions fulfilling additional criteria. In particular, we consider different repair semantics and the description logics DL-Lite and $\mathcal{EL}_\bot$.

Keywords

Description Logics, abduction, repair semantics

1. Introduction

Abduction is an important reasoning problem that can be formalized as follows: given two sentences $\mathcal{K}$ and Φ s.t. $\mathcal{K} \not\models \Phi$, find an appropriate *hypothesis* $\mathcal{H}$ such that $\mathcal{K}, \mathcal{H} \models \Phi$. Depending on the context, there may be different definitions of what "appropriate" means: for instance, one would commonly require that $\Phi \notin \mathcal{H}$ (non-triviality) and that $\mathcal{K}, \mathcal{H} \not\models \bot$ (consistency). In the context of knowledge representation, $\mathcal{K}$ is a knowledge base and $\mathcal{H}$ is a set of facts that can be added to $\mathcal{K}$ to explain or to fix a missing entailment. We call Φ the *observation*.

In *description logics*, a knowledge base has two components: a *TBox*, which provides knowledge about the terminology, and an *ABox*, which contains assertions about concrete individual objects. We find different variants of abduction in the literature, depending on which component we want to apply abduction on: *TBox* abduction requires Φ and $\mathcal{H}$ to consist of TBox axioms [1, 2, 3, 4], while *ABox* abduction requires them to consist of ABox assertions [5, 6, 7, 8, 9]. Knowledge base abduction generalizes both and allows Φ and $\mathcal{H}$ to be mixed [10, 11]. In this paper, we are interested in ABox abduction with observations that are single assertions, which is relevant for diagnosis tasks.

Example 1. Assume a TBox $\mathcal{T}_1$ representing a medical ontology, which includes the following axioms:

$$\mathit{High} \sqcap \mathit{Low} \sqsubseteq \bot \quad (1) \qquad \exists \mathit{glucoseLevel}.\mathit{High} \sqcap \mathit{OverdosedInsulin} \sqsubseteq \mathit{DiabeticComa} \quad (2)$$
$$\exists \mathit{glucoseLevel}.\mathit{High} \sqsubseteq \mathit{GlycemicCrisis} \quad (3) \qquad \exists \mathit{glucoseLevel}.\mathit{Low} \sqsubseteq \mathit{GlycemicCrisis} \quad (4)$$
$$\mathit{GlycemicCrisis} \sqcap \mathit{Ketoacidosis} \sqsubseteq \mathit{DiabeticComa} \quad (5)$$

These axioms state that 1. nothing can be high and low at the same time, 2. a high glucose level and an overdose of insulin leads to a diabetic coma, 3. a high glucose level causes a glycemic crisis, as does 4. a low glucose level, and 5. a glycemic crisis together with a ketoacidosis leads to a diabetic coma.

Logical approaches to handling INconsistent DAta - Linda 2026, July 24, 2026, Lisbon, Portugal

anselm.haak@uni-paderborn.de (A. Haak); p.k.koopmann@vu.nl (P. Koopmann); yasir.mahmood@uni-paderborn.de (Y. Mahmood); turhan@uni-paderborn.de (A. Turhan)
0000-0003-1031-5922 (A. Haak); 0000-0001-5999-2583 (P. Koopmann); 0000-0002-5651-5391 (Y. Mahmood); 0000-0001-6336-335X (A. Turhan)

Our knowledge contains additionally an ABox $\mathcal{A}_1$ with medical evidence about a patient that is obtained through a glucose monitor. It is expressed by the following two assertions:

$$\mathit{glucoseLevel}(\mathit{patient}, l) \qquad \mathit{High}(l)$$

We observe that the patient passed out. Based on his glucose reading, a reasonable explanation is that he took too much insulin. Indeed, using ABox abduction for the observation $\mathit{DiabeticComa}(\mathit{patient})$, a possible hypothesis would be $\{ \mathit{OverdosedInsulin}(\mathit{patient}) \}$. An alternative hypothesis w.r.t. $\mathcal{T}_1$ and $\mathcal{A}_1$ would be $\{ \mathit{Ketoacidosis}(\mathit{patient}) \}$.

ABox abduction can be used to solve a variety of tasks in knowledge representation and knowledge engineering: in addition to diagnosis tasks [10, 12, 8, 3], it can also be used to explain missing entailments [7, 13, 3] and to suggest repairs to incomplete knowledge bases [1, 2, 4]. Abduction has even been suggested to explain machine learning results [14].

In many realistic settings, it is unreasonable to assume that the ABox is consistent—especially very large ABoxes are prone to errors that lead to logical inconsistencies once combined with the TBox. A prominent approach to cope with this issue is to use *repair-based semantics* [15]. These semantics are based on the notion of a *repair*, which is a maximal consistent subset of the ABox. Two prominent edge-cases of repair-based semantics are brave semantics and AR semantics. Under *brave semantics*, an assertion is entailed if it is entailed by at least one repair. Under *AR semantics*, it must be entailed by all repairs. If we replace the classical entailment relation with a repair-based one, we obtain ABox abduction under repair semantics.

Example 2. To obtain further evidence, the doctor uses a finger stick sensor that measures the glucose level of the patient. This leads to the additional ABox fact $\mathit{Low}(l)$. The new ABox $\mathcal{A}_2$ is inconsistent with $\mathcal{T}_1$ and has two repairs:

$$\mathcal{R}_1 = \{ \mathit{glucoseLevel}(\mathit{patient}, l), \quad \mathit{High}(l) \} \qquad \mathcal{R}_2 = \{ \mathit{glucoseLevel}(\mathit{patient}, l), \quad \mathit{Low}(l) \}$$

Under brave semantics, $\mathit{High}(l)$, $\mathit{Low}(l)$ and $\mathit{GlycemicCrisis}(\mathit{patient})$ are all entailed, while under *AR* semantics, only $\mathit{GlycemicCrisis}(\mathit{patient})$ is entailed. If we use ABox abduction under repair semantics on our observation $\mathit{DiabeticComa}(\mathit{patient})$, both $\{ \mathit{OverdosedInsulin}(\mathit{patient}) \}$ and $\{ \mathit{Ketoacidosis}(\mathit{patient}) \}$ are hypotheses under brave semantics, while only $\{ \mathit{Ketoacidosis}(\mathit{patient}) \}$ is a hypothesis under AR semantics. Indeed, a sceptical doctor should investigate this hypothesis first, as it has the bigger support.

Inconsistency of the initial knowledge base and the use of repair semantics affect what properties a desirable hypothesis should have. Indeed, the typical requirement that hypotheses are consistent with the knowledge base is of little use if the knowledge base is already inconsistent. On the other hand, if we ignore conflicts between the hypothesis and the knowledge base, our hypotheses may not align well with our knowledge. In [16], the authors proposed requiring hypotheses to be *conflict-confining*, which means that adding them to the knowledge base does not introduce new conflicts. We complement this notion with a corresponding optimality criterion, which requires the set of introduced conflicts to be minimal. The authors in [16] consider one repair semantics, one description logic (DL-Lite) and one set of properties the hypothesis should satisfy. We provide a more detailed analysis considering also the DL $\mathcal{EL}_\bot$, three desirable properties of hypotheses, and a total of four optimality criteria. In particular, we studied the complexity of the existence problem—does there exist a hypothesis for a given abduction problem—as well as the verification problem—is a given set of assertions a hypothesis for a given abduction problem. The results presented here are taken from a recent conference publication [17].

2. ABox Abduction Under Repair-Based Semantics

We define repair-based semantics and our abduction problem formally. We consider knowledge bases in the description logics DL-Lite$_{\text{core}}$, DL-Lite$_{\mathcal{R}}$ and $\mathcal{EL}_\bot$. The syntax and semantics is not important for this paper—the interested reader is referred to [18]. A *knowledge base* is a tuple $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$, where the *TBox* $\mathcal{T}$ consists of terminological axioms and the *ABox* $\mathcal{A}$ of flat assertions (i.e. not using complex concepts). Classical entailment is defined as usual using the semantics of first-order logics. We assume that the TBox is consistent, that is, $\mathcal{T} \not\models \bot$. Then, given an ABox $\mathcal{A}$ s.t. $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle \models \bot$, a *repair* of $\mathcal{K}$ is a subset $\mathcal{R} \subseteq \mathcal{A}$ s.t. $\langle \mathcal{T}, \mathcal{R} \rangle \not\models \bot$ and for every $\mathcal{R}'$ s.t. $\mathcal{R} \subsetneq \mathcal{R}' \subseteq \mathcal{A}$, we have $\langle \mathcal{T}, \mathcal{R}' \rangle \models \bot$. The dual notion to repair is that of a *conflict*, which is a subset $\mathcal{C} \subseteq \mathcal{A}$ s.t. $\langle \mathcal{T}, \mathcal{C} \rangle \models \bot$ and for every $\mathcal{C}' \subsetneq \mathcal{C}$, we have $\langle \mathcal{T}, \mathcal{C}' \rangle \not\models \bot$. Intuitively, a conflict is a minimal explanation for inconsistency of a knowledge base. We use the notation $\mathrm{Conf}(\mathcal{K})$ to refer to the set of all conflicts of a knowledge base $\mathcal{K}$.

Based on the notion of repair, we define two entailment relations. Given an assertion α,

1. α is entailed in $\mathcal{K}$ under *brave semantics*, in symbols $\mathcal{K} \models_{\text{brave}} \alpha$, if *there exists* a repair $\mathcal{R}$ of $\mathcal{K}$ s.t. $\langle \mathcal{T}, \mathcal{R} \rangle \models \alpha$,
2. α is entailed in $\mathcal{K}$ under *AR semantics*, in symbols $\mathcal{K} \models_{\text{AR}} \alpha$, if *for every* repair $\mathcal{R}$ of $\mathcal{K}$ we have $\langle \mathcal{T}, \mathcal{R} \rangle \models \alpha$.

Using these notions, we can define our abduction problem.

Definition 3. Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ be an inconsistent knowledge base, Φ an assertion called the *observation*, and $\mathcal{S} \in \{\text{brave}, \text{AR}\}$ a repair semantics such that $\mathcal{K} \not\models_{\mathcal{S}} \Phi$. Then, the pair $\langle \mathcal{K}, \Phi \rangle$ is called an *$\mathcal{S}$-abduction problem*. A solution for such a problem, called *$\mathcal{S}$-hypothesis*, is an ABox $\mathcal{H}$ using only individuals already occurring in $\mathcal{K}$ and Φ s.t. $\langle \mathcal{T}, \mathcal{A} \cup \mathcal{H} \rangle \models_{\mathcal{S}} \Phi$.

If additionally a signature Σ of predicate names is given, we call a triple $\langle \mathcal{K}, \Phi, \Sigma \rangle$ a *Σ-restricted $\mathcal{S}$-abduction problem*. A solution to this problem is an ABox $\mathcal{H}$ that uses only names from Σ and that is an $\mathcal{S}$-hypothesis for $\langle \mathcal{K}, \Phi \rangle$.

Signature-restrictions can have several motivations: one is that in many applications, hypotheses constructed from a specific vocabulary are more useful. For example, in case of a diagnosis use-case as in the introduction, we might want to use a signature that only contains names used to describe causes. Another motivation are ABoxes that correspond to database tables with a restricted schema [19, 7].

In addition to the signature-restriction, we consider the following properties of hypotheses:

Definition 4. Let $\mathcal{S} \in \{\text{brave}, \text{AR}\}$, $\langle \mathcal{K}, \alpha \rangle$ be an $\mathcal{S}$-abduction problem, and $\preceq \in \{\subseteq, \leq\}$. Then, a hypothesis $\mathcal{H}$ is

1. *conflict-confining for* $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ if $\mathrm{Conf}(\langle \mathcal{T}, \mathcal{A} \cup \mathcal{H} \rangle) = \mathrm{Conf}(\mathcal{K})$,
2. *non-trivial* if $\alpha \notin \mathcal{H}$,
3. *$\preceq$-minimal* if there is no $\mathcal{S}$-hypothesis $\mathcal{H}'$ for $\langle \mathcal{K}, \alpha \rangle$ such that $\mathcal{H}' \prec \mathcal{H}$, and
4. *$\preceq_c$-minimal* if there is no $\mathcal{S}$-hypothesis $\mathcal{H}'$ for $\langle \mathcal{K}, \alpha \rangle$ such that $\mathrm{Conf}(\langle \mathcal{T}, \mathcal{A} \cup \mathcal{H}' \rangle) \prec \mathrm{Conf}(\langle \mathcal{T}, \mathcal{A} \cup \mathcal{H} \rangle)$.

Non-triviality (Item 2) as well as subset minimality and size minimality (Item 3) of hypotheses are standard for abduction, and do not deserve much explanation. The property conflict-confining (Item 1) and conflict-minimality (Item 4) are tailored to the setting of inconsistent knowledge bases. Specifically, the requirement of conflict-confining generalizes the requirement of consistency typically applied in the classical setting. Even if the knowledge base is already inconsistent to begin with, we might still want to avoid hypotheses that contradict our knowledge base additionally. In case a conflict-confining hypothesis does not exist, a $\preceq_c$-minimal hypothesis could be a sufficient approximation.

3. First Results and Outlook

We analysed the *combined complexity* of the existence and the verification problem in different settings for the tractable description logics DL-Lite and $\mathcal{EL}_\bot$. Specifically, we considered the dialects $\text{DL-Lite}_{\text{core}}$ and $\text{DL-Lite}_{\mathcal{R}}$, for which we obtained the same completeness results. Here, *combined complexity* refers to the assumption that both TBox and ABox are given as input, and is in the literature often contrasted with *data complexity*, which assumes the TBox to be fixed. Table 1 summarizes our results from [17].

		Existence Problem				Verification Problem				
		general	signature	conflict-conf.	non-trivial	$\leq$-min	$\subseteq$-min	$\subseteq_c$-min	$\leq_c$-min	conflict-conf.
DL-Lite	brave	NL	NL	NL	NL	NL	NL	NL	NL	NL
	AR	NL	coNP	NL	coNP	coNP	DP	NL	NL	NL
$\mathcal{EL}_\bot$	brave	P	NP	Σ_2^P	NP	NP	DP	Π_2^P	–	DP
	AR	coNP	Σ_2^P	coNP	Σ_2^P	coNP	Π_2^P	coNP	coNP	coNP

Table 1
Complexity for existence and verification of hypotheses with different properties and repair semantics. Verifying general hypotheses has the same complexity as with $\leq$-minimality.

We recall that testing entailment for DL-Lite under brave semantics is NL-complete, while for $\mathcal{EL}_\bot$ it is NP-complete for combined complexity. Under AR semantics, the complexity of testing entailment is coNP-complete for both description logics [15]. Considering the results in Table 1, we observe that for DL-Lite, the complexity of both existence and verification is not higher than that of entailment in the respective semantics, with the exception of verification of $\subseteq$-minimal hypotheses. In fact, the complexity is often even lower. For $\mathcal{EL}_\bot$, the picture is quite different: we have a case where deciding existence of a hypothesis (including trivial ones) is tractable (under brave semantics). In all the other considered settings, the complexity remains the same or even goes up, often one level in the polynomial hierarchy. We observe an interesting behaviour in the case of conflict-confining hypotheses in $\mathcal{EL}_\bot$: determining their existence is Σ_2^P-complete under brave semantics, and only coNP-complete under AR-semantics. If we drop the requirement of being conflict-confining, and instead require hypotheses to be signature-restricted or non-trivial, the picture reverts: now it is easier under Brave semantics (NP-complete), and one level up in the polynomial hierarchy for AR semantics.

While we have already quite a detailed picture of the complexity landscape, there are a few relevant questions we plan to investigate in the future. First of all, our current study considers the different properties and optimality criteria in isolation. In realistic scenarios, one would usually want to combine them and e.g. ask for a signature-restricted hypothesis that is also conflict-confining and non-trivial. We have already some initial results on combinations of up to three properties in [20].

Another direction is moving towards abduction with queries, as it was investigated in [7] for abduction under classical semantics. Indeed, query answering is a relevant application of ABox reasoning, and a central application of DL-Lite ontologies especially. Specifically, instead of a single assertion, we would consider a set of assertions that may also include quantified variables (i.e. conjunctive queries), or even formulas in the positive existential fragment of first-order logic, which corresponds to relevant and well-studied query languages for description logic knowledge bases. In the context of query answering, it is also typical to not only consider *combined* complexity, as we did in our study, but also *data complexity*, for which one assumes the query and TBox to be fixed. To complete the complexity landscape in that direction is also future work.

References

[1] F. Wei-Kleiner, Z. Dragisic, P. Lambrix, Abduction framework for repairing incomplete $\mathcal{EL}$ ontologies: Complexity results and algorithms, in: Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, AAAI Press, 2014, pp. 1120–1127. URL: http://www.aaai.org/ocs/index.php/AAAI/AAAI14/paper/view/8239.

[2] J. Du, H. Wan, H. Ma, Practical TBox abduction based on justification patterns, in: Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, 2017, pp. 1100–1106. URL: http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14402.

[3] P. Koopmann, Signature-based abduction with fresh individuals and complex concepts for description logics, in: Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, ijcai.org, 2021, pp. 1929–1935. URL: https://doi.org/10.24963/ijcai.2021/266.

[4] F. Haifani, P. Koopmann, S. Tourret, C. Weidenbach, Connection-minimal abduction in $\mathcal{EL}$ via

translation to FOL, in: J. Blanchette, L. Kovács, D. Pattinson (Eds.), Proceedings of the 11th International Joint Conference on Automated Reasoning IJCAR 2022, volume 13385 of *LNCS*, Springer, 2022, pp. 188–207. URL: https://doi.org/10.1007/978-3-031-10769-6_12.

[5] S. Klarman, U. Endriss, S. Schlobach, ABox abduction in the description logic $\mathcal{ALC}$, Journal of Automated Reasoning 46 (2011) 43–80.

[6] K. Halland, K. Britz, ABox abduction in $\mathcal{ALC}$ using a DL tableau, in: Proceedings of the South African Institute for Computer Scientists and Information Technologists Conference — SAICSIT '12, Association for Computing Machinery, 2012, p. 51–58. doi:10.1145/2389836.2389843.

[7] D. Calvanese, M. Ortiz, M. Simkus, G. Stefanoni, Reasoning about explanations for negative query answers in DL-Lite, J. Artif. Intell. Res. 48 (2013) 635–669. doi:10.1613/jair.3870.

[8] İ. İ. Ceylan, T. Lukasiewicz, E. Malizia, C. Molinaro, A. Vaicenavicius, Explanations for negative query answers under existential rules, in: Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2020, AAAI Press, 2020, pp. 223–232. URL: https://doi.org/10.24963/kr.2020/23.

[9] M. Homola, J. Pukancová, J. Boborová, I. Balintová, Merge, explain, iterate: A combination of MHS and MXP in an ABox abduction solver, in: Proceedings of JELIA 2023, volume 14281 of *LNCS*, Springer, 2023, pp. 338–352. URL: https://doi.org/10.1007/978-3-031-43619-2_24.

[10] C. Elsenbroich, O. Kutz, U. Sattler, A case for abductive reasoning over ontologies, in: Proceedings of the OWLED'06 Workshop on OWL: Experiences and Directions, 2006. URL: http://ceur-ws.org/Vol-216/submission_25.pdf.

[11] P. Koopmann, W. Del-Pinto, S. Tourret, R. A. Schmidt, Signature-based abduction for expressive description logics, in: Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2020, AAAI Press, 2020, pp. 592–602. URL: https://doi.org/10.24963/kr.2020/59.

[12] M. Obeid, Z. Obeid, A. Moubaiddin, N. Obeid, Using description logic and ABox abduction to capture medical diagnosis, in: 32nd International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems, IEA/AIE 2019, LNCS, Springer, 2019, pp. 376–388. doi:10.1007/978-3-030-22999-3_33.

[13] C. Alrabbaa, S. Borgwardt, T. Friese, A. Hirsch, N. Knieriemen, P. Koopmann, A. Kovtunova, A. Krüger, A. Popovic, I. S. R. Siahaan, Explaining reasoning results for OWL ontologies with Evee, in: Proceedings of KR'24, 2024. doi:10.24963/KR.2024/67.

[14] A. Ignatiev, N. Narodytska, J. Marques-Silva, Abduction-based explanations for machine learning models, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 33, 2019, pp. 1511–1519.

[15] M. Bienvenu, C. Bourgaux, Inconsistency-tolerant querying of description logic knowledge bases, in: Reasoning Web: Logical Foundation of Knowledge Graph Construction and Query Answering — 12th International Summer School 2016, volume 9885 of *LNCS*, Springer, 2016, pp. 156–202. doi:10.1007/978-3-319-49493-7_5.

[16] J. Du, K. Wang, Y. Shen, Towards tractable and practical ABox abduction over inconsistent description logic ontologies, in: Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI Press, 2015, pp. 1489–1495. URL: https://doi.org/10.1609/aaai.v29i1.9393.

[17] A. Haak, P. Koopmann, Y. Mahmood, A.-Y. Turhan, ABox abduction for inconsistent knowledge bases under repair semantics, in: R. Wassermann, M.-L. Mugnier, F. Baader (Eds.), Proceedings of KR'26, 2026.

[18] F. Baader, I. Horrocks, C. Lutz, U. Sattler, An Introduction to Description Logic, Cambridge University Press, 2017.

[19] G. Xiao, D. Calvanese, R. Kontchakov, D. Lembo, A. Poggi, R. Rosati, M. Zakharyaschev, Ontology-based data access: A survey, in: J. Lang (Ed.), Proceedings of IJCAI 2018, ijcai.org, 2018, pp. 5511–5519. doi:10.24963/IJCAI.2018/777.

[20] A. Haak, P. Koopmann, Y. Mahmood, A.-Y. Turhan, The more the merrier: Combining properties for ABox abduction under repair semantics for $\mathcal{EL}_\bot$, in: Proceedings of the 39th International Workshop on Description Logics - DL'26, 2026.