

A Case for Abductive Reasoning over Ontologies

Corinna Elsenbroich¹, Oliver Kutz², and Ulrike Sattler²

¹ Department of Computer Science, King's College London,
Strand, London WC2R 2LS, UK

`corinna@dcs.kcl.ac.uk`

² School of Computer Science, The University of Manchester,
Kilburn Building, Oxford Road, Manchester M13 9PL, UK

`{kutz, sattler}@cs.man.ac.uk`

Abstract. We argue for the usefulness of abductive reasoning in the context of ontologies. We discuss several application scenarios in which various forms of abduction would be useful, introduce corresponding abductive reasoning tasks, give examples, and begin to develop the formal apparatus needed to employ abductive inference in expressive description logics.

1 Introduction and Motivation

Although the Description Logic and Non-Monotonic Logic communities are still vastly disjoint, there is a growing consensus in the DL and Semantic Web communities that many interesting application areas emerging from the usage of DLs as ontology languages will also require, in addition to the mostly deductive and monotonic reasoning techniques of current systems and the development of ever more expressive DLs such as the DL *SR_QIQ* (11) underlying OWL 1.1, the adoption of various forms of *non-monotonic* reasoning techniques, as well as so called non-standard inferences (3). For instance, recent work trying to integrate DLs with non-monotonicity includes non-monotonic reasoning for DLs based on circumscription (4), knowledge integration for DLs using techniques from propositional inconsistency management with belief update (18), and the relationship between DL and Logic Programming (19).

A related area of reasoning that is essential for common sense reasoning is the ability to reason from *observations* to *explanations*, and is a fundamental source of new knowledge, i.e., *learning*. This mode of reasoning, introduced originally by Charles Sanders Peirce (21), is traditionally called **abduction**.

In general, abduction is often understood as a form of *backward reasoning* from a set of observations back to a cause. For example, from a set of clues to a murderer (criminology), a set of symptoms to a disease (medical diagnosis), or from a malfunction in a system to a faulty part in a system (model-based diagnosis). I.e., from the observation that *B* holds, and the fact that *A* ‘implies’ *B* (*A* ‘is a reason for *B*’/‘is a cause for *B*’ etc.), infer *A*.

Within classical logic, this kind of reasoning is a *non sequitur* inference, called *affirming the consequent*. Thus we have to constrain this reasoning in some ways

and cannot just add this rule to a deductive calculus. There are different constraints for abduction, the most common go by the following names, where Γ is some knowledge base, and A, B are formulae, cf. (2):

1. Consistency: $\Gamma + A \not\models \perp$;
2. Minimality: A is a ‘*minimal explanation*’ for B ;
3. Relevance: $A \not\models B$;
4. Explanatoriness: $\Gamma \not\models B, A \not\models B$.

These criteria can be taken in a pick and choose fashion. For example, we might want consistent and minimal abduction, but leave the case where B is added as its own explanation as a limiting case, and thus forsake both, relevance and explanatoriness. The most fundamental constraint of the above list is consistency. There is basically no definition of abduction in the literature that does not have the requirement of consistency, unless the inference is left completely unconstrained as $\Gamma + A \models B$, often referred to as *plain abduction*. The reason for this requirement is clear. If our formal language is based on classical logic, any formula A inconsistent with the knowledge base Γ counts as an ‘explanation’ for every other formula. Thus, giving up consistency leads to a trivialisation of the notion of ‘explanation’ unless one adopts some form of paraconsistent reasoning.

Abduction as a fundamental form of reasoning next to induction and deduction is not per se non-monotonic. It often exhibits, however, non-monotonic features. For instance, if we only seek for consistent explanations, an explanation A for B w.r.t the knowledge base Γ ceases to count as an explanation for B once the KB Γ is extended by information that is inconsistent with A . Unlike deduction, abduction and induction are *ampliative* in the sense that they provide more ‘knowledge’ than can be obtained deductively from a given knowledge base.

The relevance of forms of abductive reasoning in the context of the World Wide Web was noted early on, before the idea of a *Semantic Web* in the current sense was even envisioned. Shank and Cunningham write in 1996 (24):

There is an embarrassing lack of research on learning from the World Wide Web. [...] The learner is an information seeker, not the end point of a communicative act. As we move into the information age and are inundated by increasing volumes of information, the need for reasoning skills – rather than mastery of a subject matter – is ever more evident. [...]

Traditional models of inductive and deductive inference are simply inadequate in conceptualizing the skills necessary to utilize the WWW. On the web we are seeking omens and clues, diagnosing symptoms and scenarios, etc. In other words, the inferential basis of learning from the web is largely abductive (although induction and deduction may also come to the fore at various points in the information exploration process).

Curiously, however, while abductive reasoning techniques have been intensively studied in the context of classical, mostly propositional logic, we are aware only of very little work on abductive reasoning that addresses non-classical logics in general, and description logics in particular, although, as we argue, the potential applications are legion.

We outline the potential use-cases for abduction in the context of ontologies by discussing several abductive reasoning tasks and examples, and begin to develop the formal apparatus needed to employ abductive inference in expressive description logics.

2 Preliminaries

Since we intend to capture general abductive reasoning tasks, we will usually not specify the logic we are dealing with, but just refer to it generically as \mathcal{L} , denoting, for instance, DLs such as \mathcal{ALC} or \mathcal{SHIQ} . Furthermore, since the concept language etc. is not fixed, we will avoid terms such as *general concept inclusion* (GCI) and will not specify the syntactic form of Abox assertions. Instead, we will rather use the neutral terms Tbox and Abox assertion. Thus, for instance, by $\phi(\bar{a})$ we shall denote an Abox assertion that uses individuals from the tuple $\bar{a} = \langle a_1, \dots, a_n \rangle$. Moreover, we will identify a logic \mathcal{L} with its language, i.e., with its set of concept, role, and individual names, and its set of concept constructors, e.g., conjunction, existential restrictions, etc. Also, we will distinguish between the **source logic** \mathcal{L} of an abduction problem, and the **target logic** \mathcal{L}' in which we seek for solutions, which may or may not be the same. Given a set of assertions Δ , we shall denote by $\text{Sig}(\Delta)$, the **signature** of Δ , the set of non-logical symbols (i.e., concept, role, and individual names) used in Δ . Given a logic \mathcal{L} , by \models we always mean the global consequence relation of \mathcal{L} , i.e., $\Gamma \models C \sqsubseteq D$ means that the concept D subsumes C w.r.t. the ontology Γ . Also, the notions of consistency of a knowledge base, satisfiability of a concept, etc., are standard.

3 Abductive Reasoning Tasks in DL

3.1 Concept Abduction

We start by introducing the most basic abductive reasoning task, that of finding abductively a concept H that is subsumed by a given (satisfiable) concept C .

Definition 1 (Simple Concept Abduction). *Let \mathcal{L} be a DL, C a concept in \mathcal{L} , Γ a knowledge base in \mathcal{L} , and suppose that C is satisfiable w.r.t. Γ . A solution to the **simple concept abduction problem** for $\langle \Gamma, C \rangle$ is any concept H in \mathcal{L}' such that*

$$\Gamma \models H \sqsubseteq C.$$

The set of all such solutions is denoted by $\mathcal{S}_{SCA}(\Gamma, C)$.

Note that if \mathcal{L}' as well as C are restricted to *concept names*, this is just the problem of building the class tree and picking the concept names being subsumed by a given concept C , a task readily accomplished by current DL implementations even for very expressive description logics. However, we might be interested in finding concepts that are subsumed by a *complex concept* C .

Example 1 (Geographical Information System). Assume Γ is an ontology about geography in Europe, comprising concepts `Country` etc., nominals `France` etc. for specific countries, a role `is_member_of`, etc. Suppose your Tbox contains

$$\{\text{France}\} \sqsubseteq \exists \text{is_member_of}.\text{Schengen_Treaty}$$

etc., and the Abox contains assertions `France: Country` etc.

Now, given the concept `country` $\sqcap \exists \text{isMemberOf}.\text{SchengenTreaty}$, a solution to the simple concept abduction problem, maximal w.r.t. \sqsubseteq , is the disjunction `France` $\sqcup \dots \sqcup \text{Germany}$, i.e., the set of countries that happen to be members of the Schengen Treaty.

A variant of plain concept abduction that we call *conditionalised concept abduction*, introduced in (5) where it is called *concept abduction*, is motivated by the *matchmaking problem*: in a scenario where we want to match a ‘demand’ D with a ‘supply’ C , we need to find a concept H such that $\Gamma \models C \sqcap H \sqsubseteq D$, i.e., where H gives the additional assumptions needed to meet the demand.

Definition 2 (Conditionalised Concept Abduction). *Let \mathcal{L} be a DL, C and D concepts in \mathcal{L} , Γ a knowledge base in \mathcal{L} , and suppose that C, D are satisfiable concepts w.r.t. Γ . A solution to the **conditionalised concept abduction problem** for $\langle \Gamma, C, D \rangle$ is any concept H in \mathcal{L}' such that*

$$\Gamma \not\models C \sqcap H \equiv \perp \quad \text{and} \quad \Gamma \models C \sqcap H \sqsubseteq D.$$

The set of all such solutions is denoted by $\mathcal{S}_{CCA}(\Gamma, C, D)$.

Clearly, simple concept abduction is the same as conditionalised concept abduction if we choose \top for the input concept C and require that solutions H are consistent with Γ . Put another way, conditionalised concept abduction seeks (consistent) solutions to an SCA-problem of a specific syntactic form, namely conjunctions $C \sqcap H$, where C is fixed in advance. Thus, algorithms that can find solutions to the conditionalised concept abduction problem also solve the simple concept abduction problem.

In (6), a tableaux algorithm to compute solutions to the conditionalised concept abduction problem is proposed, and various minimality principles are discussed. However, only the rather inexpressive DL \mathcal{ALN} (which allows only atomic negation and no disjunctions or existential restrictions) is considered, and it is not clear how far their techniques extend to more expressive DLs.

Another area where the use of concept abduction has been discussed is the problem of querying for similar workflow fragments (10), where fragments that are largely relevant to a user may happen to fall outside a strict subsumption relationship, and thus a ‘matching of similar concepts’ is sought.

3.2 Abox Abduction

Abox abduction can be understood as a new query answering service, retrieving abductively instances of concepts (or roles) that would entail a desired Abox assertion.

Definition 3 (Abox Abduction). Let \mathcal{L} be a DL, Γ a knowledge base in \mathcal{L} , and $\phi(\bar{a})$ an Abox assertion in \mathcal{L} such that $\Gamma \cup \phi(\bar{a})$ is consistent. A solution to the **Abox abduction problem** for $\langle \Gamma, \bar{a}, \phi \rangle$ is any finite set $\mathcal{S}_A = \{\psi_i(b_i) \mid i \leq n\}$ of Abox assertions of \mathcal{L}' such that

$$\Gamma \cup \mathcal{S}_A \models \phi(\bar{a}).$$

If \mathcal{S}_A contains only assertions of the form $\psi(\bar{b})$ with $\bar{b} \subseteq \bar{a}$, it is called a **solipsistic solution**. The set of all solutions is denoted by $\mathcal{S}_A(\Gamma, \bar{a}, \phi)$, and the set of all solipsistic solutions by $\mathcal{S}_{AS}(\Gamma, \bar{a}, \phi)$.

Example 2 (Medical Diagnosis). Consider the problem of *diagnosis* in medical ontologies. Suppose there is a disease, called the Shake-Hands-Disease (SHD), that always develops when you shake hands with someone who carries the Shake-Hands-Disease-Virus (SHDV). Suppose further your medical ontology Γ contains roles *has_symptom*, *carries_virus*, *has_disease*, etc., concepts SHD, SHDV Laziness, Pizza_Appetite, Google_Lover, etc., and individual names Peter, Paul, Mary, etc. Further, assume your Tbox contains axioms

$$\begin{aligned} \exists has_disease.SHD \sqsubseteq \exists has_symptom.(Laziness \sqcap Pizza_Appetite) \\ \text{Researcher} \sqsubseteq \exists has_symptom.(Laziness \sqcap Pizza_Appetite \sqcap Google_Lover) \\ \exists shake_hands.\exists carries_virus.SHDV \sqsubseteq \exists has_disease.SHD. \end{aligned}$$

and your Abox contains $\text{Mary} : \exists carries_virus.SHDV$.

Now, suppose you observe the fact that

$$(\dagger) \quad \text{Paul} : \exists has_symptom.(Laziness \sqcap Pizza_Appetite).$$

Then a solipsistic solution to the Abox abduction problem given by (\dagger) is

$$\{\text{Paul} : \text{Researcher}\}.$$

However, a non-solipsistic solution is also given by

$$\{(\text{Paul}, \text{Mary}) : shake_hands\},$$

which would suggest that Paul has the Shake-Hands-Disease.

3.3 Tbox Abduction

Tbox abduction can be used to repair unwanted non-subsumptions.

Definition 4 (Tbox Abduction). Let \mathcal{L} be a DL, Γ a knowledge base in \mathcal{L} , and C, D concepts that are satisfiable w.r.t. Γ and such that $\Gamma \cup \{C \sqsubseteq D\}$ is consistent. A solution to the **Tbox abduction problem** for $\langle \Gamma, C, D \rangle$ is any finite set $\mathcal{S}_T = \{G_i \sqsubseteq H_i \mid i \leq n\}$ of Tbox assertions in \mathcal{L}' such that

$$\Gamma \cup \mathcal{S}_T \models C \sqsubseteq D.$$

The set of all such solutions is denoted by $\mathcal{S}_T(\Gamma, C, D)$.

Example 3 (Debugging). In (23; 12), it is shown how to debug and repair unsatisfiable concepts C , respectively, unwanted subsumptions $C \sqsubseteq D$. By computing all the minimal unsatisfiability preserving sub-TBoxes (MUPS) of an unsatisfiable concept, various repair plans are suggested that remove certain GCIs from a knowledge base Γ resulting in a knowledge base Γ' such that $\Gamma' \not\models C \sqsubseteq D$. Tbox abduction provides a solution to the *dual* problem: to propose a repair plan for an unwanted non-subsumption, i.e., the case where a subsumption $C \sqsubseteq D$ is *expected* by an ontology engineer, but does not follow from Γ as developed so far. Every (explanatory) solution \mathcal{S} to the Tbox abduction problem $\langle \Gamma, C, D \rangle$ provides a finite set of Tbox assertions such that, when added to Γ to obtain Γ' , $\Gamma' \models C \sqsubseteq D$.

3.4 Knowledge Base Abduction

Knowledge base abduction generalises both Tbox and Abox abduction, thus interweaves statements about individuals with general axioms, a form of abduction that is closely related to so-called ‘explanatory induction’ (13).

Definition 5 (Knowledge Base Abduction). *Let \mathcal{L} be a DL, Γ a knowledge base in \mathcal{L} , and ϕ an Abox or Tbox assertion such that $\Gamma \cup \{\phi\}$ is consistent. A solution to the **knowledge base abduction problem** for $\langle \Gamma, \phi \rangle$ is any finite set $\mathcal{S} = \{\phi_i \mid i \leq n\}$ of Tbox and Abox assertions such that*

$$\Gamma \cup \mathcal{S} \models \phi.$$

The set of all such solutions is denoted by $\mathcal{S}_K(\phi)$.

Obviously, if ϕ is an Abox assertion, any solution to the Abox abduction problem is also a solution to the knowledge base abduction problem, and similarly if ϕ is a Tbox assertion, i.e., the following inclusions hold:

$$\mathcal{S}_A(\phi) \subseteq \mathcal{S}_K(\phi) \quad \text{and} \quad \mathcal{S}_T(\phi) \subseteq \mathcal{S}_K(\phi).$$

In fact, we can sometimes rewrite solutions to an Abox abduction problem into solutions to an equivalent Tbox abduction problem, and conversely. For instance, suppose we work in the logic \mathcal{ALCO} , comprising nominals. Then an Abox assertion $a : C$ can be equivalently rewritten as the Tbox assertion $\{a\} \sqsubseteq C$.

Example 4 (More about Mary). Suppose we have the following knowledge base (with \circ denoting the composition of roles), using concepts BOS (‘burnt out syndrome’), BOTeacher (‘burnt out teacher’), etc., consisting of TBox assertions:

$$\begin{aligned} \text{BOS} & \equiv \text{Syndrom} \sqcap \exists \text{has_symptom} . (\text{Headache} \sqcup \text{Depressed}) \\ \text{BOTeacher} & \sqsupseteq \text{Teacher} \sqcap \exists \text{shows} . \text{BOS} \\ \text{Depression} & \equiv \exists \text{has_symptom} . \text{Depressed} \\ \text{shows} \circ \text{has_symptom} & \sqsubseteq \text{has_symptom} \end{aligned}$$

and the Abox contains: Mary: Teacher $\sqcap \forall \text{has_symptom} . \neg \text{Headache}$.

Now suppose you observe that: (‡) Mary: $\exists has_symptom.Depressed$. Then, a simple solipsistic Abox solution to (‡) is given by:

Mary: Depression

This is, however, not very interesting, since it is equivalent to the observation itself and does not take into account other knowledge that we have of Mary. The following solution involving both Tbox and Abox statements is much more informative:

Mary: BOTeacher and BOTeacher \sqsubseteq Teacher \sqcap $\exists shows.BOS$

This adds a definition of the concept BOTeacher to the KB, and makes Mary an instance of this subclass of Teacher.

4 Finding and Selecting Solutions

We now address the two main problem areas associated with abductive reasoning, namely the *selection* of ‘good’ solutions, and the algorithmic problem of *finding* solutions. We begin with the selection problem.

4.1 Selecting Solutions

Clearly, the definitions given so far are very general in that they even allow for inconsistent and trivial solutions, and do not make any explicit restrictions on the syntactic form of solutions. These are in fact the most common restrictions, i.e., restrictions concerning the deductive properties of solutions, and restrictions concerning the syntactic form of solutions. Apart from the definition of *creative solutions*, the following is a standard classification often used in connection with solutions to abductive problems (20; 2):

Definition 6. Let \mathcal{L} be a DL, Γ a knowledge base in \mathcal{L} , $\mathbb{P}(\Gamma, \phi)$ an abductive problem, and $\mathcal{S}(\mathbb{P}(\Gamma, \phi))$ the set of its solutions in language \mathcal{L}' . We call solutions:

1. $\mathcal{S}(\mathbb{P}(\Gamma, \phi))$ **plain**
2. $\mathcal{S}_{Con}(\mathbb{P}(\Gamma, \phi)) := \{S \in \mathcal{S}(\mathbb{P}(\Gamma, \phi)) \mid S \cup \Gamma \not\models \perp\}$ **consistent**
3. $\mathcal{S}_{Rel}(\mathbb{P}(\Gamma, \phi)) := \{S \in \mathcal{S}(\mathbb{P}(\Gamma, \phi)) \mid S \not\models \phi\}$ **relevant**
4. $\mathcal{S}_{Exp}(\mathbb{P}(\Gamma, \phi)) := \{S \in \mathcal{S}_{Rel}(\mathbb{P}(\Gamma, \phi)) \mid \Gamma \not\models \phi\}$ **explanatory**
5. $\mathcal{S}_{Cre}(\mathbb{P}(\Gamma, \phi)) := \{S \in \mathcal{S}(\mathbb{P}(\Gamma, \phi)) \mid \text{Sig}(S) \not\subseteq \text{Sig}(\Gamma, \phi)\}$ **creative**

If we restrict our attention to consistent solutions, abduction is often seen as a special form of belief expansion (used here as a technical term in line with the AGM framework of belief change (1)). This means that a set $\mathcal{S}(\mathbb{P}(\Gamma, \phi))$ of solutions is found and an element $\psi \in \mathcal{S}(\mathbb{P}(\Gamma, \phi))$ is selected and then added to the knowledge base. If, however, no consistent solution is found, either ϕ itself is added (expansion as the limiting case) or the abduction is declined (failure as the limiting case).

Clearly, this is a restrictive view of abduction. Firstly, consistent abduction cannot challenge our background theory Γ as it never triggers a revision (14). This, however, is often desirable, for instance in the case where a solution S for $\langle \Gamma, \phi \rangle$ is inconsistent but explanatory, and where it is considered important to have ϕ as a consequence of the knowledge base. In such a case, we want to revise Γ such that ϕ becomes a consequence of it, S is kept, and Γ is modified to Γ' such that $\Gamma' \cup \{S\}$ is consistent. Such a revision can be performed, for instance, by using abduction together with the techniques of (12). Secondly, our KB Γ might be inconsistent from the outset. At this point, we might not want an abduction problem $\mathbb{P}(\Gamma, \phi)$ to trigger a revision of Γ because, depending on the way revisions for Γ are performed, this might yield an updated Γ' whose modification has nothing to do with $\mathbb{P}(\Gamma, \phi)$, in which case we might prefer to perform paraconsistent abductive reasoning, as proposed in (9). Hence, although we can see applications of abduction challenging Γ and abduction not being restricted to the lucky, consistent cases, we think that it is dangerous to treat abductively inferred beliefs equal to knowledge. After all, ‘explanations’ are hypothetical and abductive inference can be superseded by new information.

Relevant and explanatory solutions impose further conditions on the deductive properties of solutions. If a solution S is relevant for $\mathbb{P}(\Gamma, \phi)$, ϕ is not already a logical consequence of S . If S is explanatory, it is guaranteed that ϕ does not already follow from Γ , but is only a consequence of Γ *together* with S .

Another interesting class of solutions are the *creative* ones—these are explanations that provide a proper extension to the vocabulary of the knowledge base, and thus add a genuinely new part to the ontology in order to explain, for example, a given subsumption. Obviously, such extensions are difficult to obtain automatically, but could be obtained semi-automatically via an interactive ontology revision and refinement process using abductive reasoning.

Further conditions that can be imposed on the class of solutions are:

Syntactic The language \mathcal{L}' in which we seek solutions to an abductive problem can be restricted in various ways. As in propositional abduction, we can insist, for example, that solutions have to be conjunctions of concept names or their negations, or that only a certain subset of the concept constructors of \mathcal{L} be used in solutions. The vocabulary of \mathcal{L}' might also be restricted in such a way that we can effectively list all finitely many non-equivalent possible solutions, in which case any abduction problem defined is decidable and solvable by brute force.

Minimality Since in general abductive problems can have infinitely many solutions, the definition of notions of *minimality* is an important part of abductive reasoning, allowing to *prefer* one solution over another. These can be *syntactic* notions, such as preferring solutions of a specific form of minimal length, or *semantic* notions, such as maximality w.r.t. subsumption, which, for the case of simple concept abduction, is defined as follows: H is a \sqsubseteq -**maximal** solution for $\mathbb{P}(\Gamma, C)$ if for all concepts E : $\Gamma \models E \sqsubseteq C$ and $\Gamma \models H \sqsubseteq E$ implies $\Gamma \models E \sqsubseteq H$ (note that different solutions can be incomparable w.r.t. \sqsubseteq).

In (15), it is shown that, for the case of abduction in first-order logic, \sqsubseteq -maximal solutions need not exist for a given problem. The same can be shown,

for instance, for the modal logic **S4** (17), and is also true for most DLs. However, (17) also shows that, when using the local consequence relation, every modal logic that enjoys the finite model property admits minimal solutions in this sense. We believe that a similar result can be shown in the context of DLs.

Conservativeness In the context of reasoning with ontologies, a number of further restrictions on what constitutes a ‘good solution’ to an abduction problem can be identified. For example, we might want our solutions to be ‘conservative’ in the sense of ensuring a certain ‘stability’ of the class tree. For instance, if our ontology comprises both a ‘foundational’ part and a domain ontology, we might want to require that the foundational part of the class tree remains unchanged whenever we add new assertions obtained by an abductive process, cf. (7).

4.2 Finding Solutions

The most developed techniques for finding abductive explanations, used mostly in the area of propositional logic, are based on abductive logic programming (ALP) and resolution, compare (20; 2; 8) and references therein. The third algorithmic technique employed for finding solutions to abductive problems is based on semantic tableaux, and builds on the simple idea that, in order to obtain a solution S for a problem $\mathbb{P}(T, \phi)$, every “tableau for $T \cup \{-\phi\}$ ” needs to be “closed by S ”. In the context of ontology reasoning with its highly developed and optimised tableaux based reasoning methods, the integration of abductive reasoning techniques into existing algorithms seems the most promising approach.

The work that is most directly relevant to pursue this line of research further is (16), which studies the use of tableaux for abduction in the context of first-order logic using ‘reversed skolemization’, (17), which adapts the techniques developed in (16) to propositional modal logics, and (22), which restricts the notion of satisfaction in first-order logic in order to effectively use tableaux for computing solutions to abductive problems formulated in first-order logic, a task that has traditionally been considered impenetrable simply for the reason that satisfiability in first-order logic is undecidable.

5 Outlook

To summarise, abductive reasoning in the context of ontologies is highly relevant from an application point of view, but represents an almost entirely open research area, including

- The identification of interesting syntactic restrictions on solutions.
- The definition of appropriate notions of minimality.
- The identification of restrictions on solutions, specifically motivated by reasoning with ontologies.
- The identification of DLs that admit tableaux-based reasoning techniques for abduction, and the study of their computational complexity.

Acknowledgements We would like to thank Thomas Meyer for valuable discussions.

Bibliography

- [1] ALCHOURRÓN, C. A., GÄRDENFORS, P., AND MAKINSON, D. The Logic of Theory Change: Partial Meet Contraction and Revision Functions. *The Journal of Symbolic Logic* 50, 2 (1985).
- [2] ALISEDA-LLERA, A. *Seeking Explanations: Abduction in Logic, Philosophy of Science and Artificial Intelligence*. PhD thesis, University of Amsterdam, 1997.
- [3] BAADER, F., AND KÜSTERS, R. Non-standard inferences in description logics: The story so far. In *Mathematical Problems from Applied Logic I. Logics for the XXIst Century*, D. M. Gabbay, S. S. Goncharov, and M. Zakharyashev, Eds., vol. 4 of *International Mathematical Series*. Springer-Verlag, 2006, pp. 1–75.
- [4] BONATTI, P., LUTZ, C., AND WOLTER, F. Expressive non-monotonic description logics based on circumscription. In *Proc. of KR-06* (2006), pp. 400–410.
- [5] COLUCCI, S., NOIA, T. D., SCIASCIO, E. D., DONINI, F. M., AND MONGIELLO, M. Concept abduction and contraction in description logics. In *Proc. of DL-03* (2003), vol. 81, CEUR Workshop Proceedings.
- [6] COLUCCI, S., NOIA, T. D., SCIASCIO, E. D., DONINI, F. M., AND MONGIELLO, M. A Uniform Tableaux-Based Approach to Concept Abduction and Contraction in \mathcal{ALN} . In *Proc. of DL-04* (2004), CEUR Workshop Proceedings, p. 104.
- [7] CUENCA-GRAU, B., HORROCKS, I., KUTZ, O., AND SATTLER, U. Will my ontologies fit together? In *Proc. of DL-06* (2006), CEUR.
- [8] ELSENBROICH, C. *Instinct for Detection*. PhD thesis, King’s College London, 2006.
- [9] ELSENBROICH, C., GABBAY, D. M., AND RODRIGUES, O. Getting Possibilities from the Impossible. In *Proc. of NMR-06* (2006), J. Dix and A. Hunter, Eds., Clausthal University of Technology, pp. 505–513.
- [10] GODERIS, A., SATTLER, U., LORD, P., AND GOBLE, C. Seven bottlenecks to workflow reuse and repurposing. In *ISWC 2005* (2005), vol. 3792, pp. 323–337.
- [11] HORROCKS, I., O.KUTZ, AND SATTLER, U. The Even More Irresistible *SROIQ*. In *Proc. of KR-06* (2006).
- [12] KALYANPUR, A. *Debugging and Repair of OWL Ontologies*. PhD thesis, The University of Maryland, 2006.
- [13] LACHICHE, N. Abduction and induction from a non-monotonic reasoning perspective. In *Abduction and Induction: Essays on their Relation and Integration*, P. Flach and A. Kakas, Eds., vol. 18. Kluwer, 2000, pp. 107–116.
- [14] LOBO, J., AND UZCÁTEGUI, C. Abductive Change Operators. *Fundamenta Informaticae* 27, 4 (1996), 385–411.
- [15] MARQUIS, P. Extending Abduction from Propositional to First-Order Logic. In *Fundamentals of AI Research: International Workshop FAIR’91 Proc.*, P. Jorrand and J. Kelemen, Eds. Springer, 1991, pp. 141–155.
- [16] MAYER, M. C., AND PIRRI, F. First-order abduction via tableau and sequent calculi. *Logic Journal of the IGPL* 1, 1 (1993), 99–117.
- [17] MAYER, M. C., AND PIRRI, F. Propositional Abduction in Modal Logic. *Logic Journal of the IGPL* 3, 6 (1995), 907–919.
- [18] MEYER, T., LEE, K., AND BOOTH, R. Knowledge integration for description logics. In *Proc. of AAAI-05* (2005), pp. 645–650.
- [19] MOTIK, B., HORROCKS, I., ROSATI, R., AND SATTLER, U. Can OWL and Logic Programming Live Together Happily Ever After? In *Proc. of ISWC-06* (2006), LNCS, SV. To appear.
- [20] PAUL, G. Approaches to abductive reasoning: an overview. *AI Review* 7 (1993), 109–152.
- [21] PEIRCE, C. S. Deduction, Induction, and Hypothesis. *Popular Science Monthly* 13 (1878), 470–482.
- [22] REYES-CABELLO, A. L., ALISEDA-LLERA, A., AND NEPOMUCENO-FERNÁNDEZ, A. Towards Abductive Reasoning in First-order Logic. *Logic Journal of the IGPL* 14, 2 (2006), 287–304.
- [23] SCHLOBACH, S. Diagnosing terminologies. In *AAAI-05* (2005), pp. 670–675.
- [24] SHANK, G., AND CUNNINGHAM, D. J. Six Modes of Peircean Abduction. In *Online Proc. of the 1996 Midwest AI and Cognitive Science Conference* (1996), M. Gasser, Ed., Indiana University.