

Exploiting Default Logic to Detect Outliers

Fabrizio Angiulli¹ and Rachel Ben-Eliyahu - Zohary²

¹ ICAR-CNR c/o DEIS

Univ. della Calabria

87030 Rende (CS), Italy

angiulli@icar.cnr.it

² Comm. Systems Engineering Dept.

Ben-Gurion Univ. of the Negev

Beer-Sheva 84105, Israel

rachel@bgumail.bgu.ac.il

Abstract. Default logic is used to describe regular behavior and normal properties. We suggest to exploit the framework of default logic for detecting *outliers* – individuals who behave in an unexpected way or feature abnormal properties. We first formally define the notion of an *outlier* and an *outlier witness*. We then show that finding outliers is quite complex. Indeed, we show that several versions of the outlier detection problem lie over the second level of the polynomial hierarchy. For example, the question of establishing if at least one outlier can be detected in a given propositional default theory is Σ_2^P -complete. In addition, we show that outlier detection can be done in polynomial time for both the class of acyclic normal unary defaults and the class of acyclic dual normal unary defaults.

1 Introduction

Default logics were developed as a tool for reasoning with incomplete knowledge. By using default rules, we can describe how things work in general and then make some assumptions about individuals and draw conclusions about their properties and behavior.

In this paper, we suggest a somewhat different usage of default logics. The basic idea is as follows. Since default rules are used for describing regular behavior, we can exploit them for detecting individuals or elements who *do not* behave normally according to the default theory at hand. We call such entities *outliers*. An outlier is an element that shows some properties that are contrary to those that can be logically justified.

Outlier detection can be useful in several application contexts, e.g., to single out exceptional behaving individuals or system components. Note that according to our approach, exceptions are not explicitly listed in the theory as “abnormals,” as is often done in logical-based abduction [11, 2, 3]. Rather, their “abnormality” is singled out exactly because some of the properties characterizing them do not have a justification within the theory at hand. For example, suppose that it usually takes about two seconds to download a one-megabyte file from some server. Then, one day, the system is slower – instead four seconds are needed to perform the same task. While four seconds may indicate a good performance it is helpful to find the source of the delay. Another example

might be that someone's car brakes are making a strange noise. Although they seem to be functioning properly, this is not normal behavior and the car should be serviced. In this case, the car brakes are outliers and the noise is their witness.

Outlier detection can also be used for examining database integrity. If an abnormal property is discovered in a database, the source who reported this observation would have to be double-checked.

Detecting abnormal properties, that is, detecting outliers, can also lead to an update of default rules. Suppose we have the rule that birds fly, and we observe a bird, say Tweety, that does not fly. We report this occurrence of an outlier in the theory to the knowledge engineer. The engineer investigates the case, finds out that Tweety is, for example, a penguin, and updates the knowledgebase with the default "penguins do not fly."

In this paper, we formally state the ideas briefly sketched above within the context of Reiter's default logic. For simplicity, we concentrate on the propositional fragment of default logic although the generalization of such ideas to the realm of first-order defaults is also worth exploring. So, whenever we use a default theory with variables, as in some of the following examples, we relate to it as an abbreviation of its grounded version.

The rest of the paper is organized as follows. In Section 2, we give preliminary definitions as well as a formal definition of the concept of an outlier. In Section 3, we describe the complexity of finding outliers in propositional default logic. Section 4 describes some tractable cases. Related work is discussed in Section 5. Conclusions are given in Section 6.

Because of space limitations, throughout the paper proofs of results are sketched or omitted. Full proofs can be found in [1].

2 Definitions

In this section we provide preliminary definitions for concepts we will be using throughout the paper.

2.1 Preliminaries

The following definitions will be assumed. Let T be a propositional theory. Then T^* denotes its logical closure. If S is a set of literals, then $\neg S$ denotes the set of all literals that are the negation of some literal in S .

Default logic was introduced by Reiter [12]. A *propositional default theory* Δ is a pair (D, W) consisting of a set W of propositional formulas and a set D of default rules. A *default rule* δ has the form $\frac{\phi:\psi}{\chi}$ (or, equivalently, $\phi : \psi/\chi$), where ϕ , ψ and χ are propositional formulas, called, respectively, *prerequisite*, *justification*, and *consequent* of δ . The prerequisite could be omitted, though justification and consequent are required. If $\psi = \chi$, the default rule is called *normal*. The informal meaning of a default rule δ is the following: if ϕ is known, and if it is consistent to assume ψ , then we conclude χ . An *extension* is a maximal set of conclusions that can be drawn from a theory. An extension E of a propositional default theory $\Delta = (D, W)$ can be finitely

characterized through the set D_E of *generating defaults* for E w.r.t. Δ , i.e., the set $D_E = \{\phi : \psi/\chi \in D \mid \phi \in E \wedge \neg\psi \notin E\}$. Indeed, $E = (W \cup \{\chi \mid \phi : \psi/\chi \in D_E\})^*$.

A propositional default theory $\Delta = (D, W)$ is *disjunction free* (DF for short), if W is a set of literals, and the prerequisite, justification and consequent of each default in D is a conjunction of literals. A DF default theory is *normal mixed unary* (NMU in short) iff its set of defaults contains only rules of the form $\frac{\alpha;\beta}{\beta}$, where α is either empty or a literal and β is a literal. An NMU default theory is *normal unary* (NU for short) iff the prerequisite of each default is either empty or positive. An NMU default theory is *dual normal unary* (DNU for short) iff the prerequisite of each default is either empty or negative.

Let Δ be a default theory and l a literal. Then $\Delta \models l$ means that l belongs to every extension of Δ . Similarly, for a set of literals S , $\Delta \models S$ means that every literal $l \in S$ belongs to every extension of Δ . A default theory is *coherent* if it has at least one extension.

We review some basic definitions about complexity theory, particularly, the polynomial hierarchy. The reader is referred to [6] for more on complexity theory. The classes Σ_k^P and Π_k^P are defined as follows: $\Sigma_0^P = \Pi_0^P = P$ and for all $k \geq 1$, $\Sigma_k^P = \text{NP}^{\Sigma_{k-1}^P}$, and $\Pi_k^P = \text{co-}\Sigma_k^P$. Σ_k^P models computability by a nondeterministic polynomial-time algorithm which may use an oracle, loosely speaking a subprogram that can be run with no computational cost, for solving a problem in Σ_{k-1}^P . The class D_k^P , $k \geq 1$, is defined as the class of problems that consists of the conjunction of two independent problems from Σ_k^P and Π_k^P , respectively. Note that for all $k \geq 1$, $\Sigma_k^P \subseteq D_k^P \subseteq \Sigma_{k+1}^P$.

A problem A is *complete* for the class \mathcal{C} iff A belongs to \mathcal{C} and every problem in \mathcal{C} is reducible to A by polynomial-time transformations. A well known Σ_k^P -complete problem is to decide the validity of a formula $QBE_{k,\exists}$, that is, a formula of the form $\exists X_1 \forall X_2 \dots Q X_k f(X_1, \dots, X_k)$, where Q is \exists if k is odd and is \forall if k is even, X_1, \dots, X_k are disjoint set of variables, and $f(X_1, \dots, X_k)$ is a propositional formula in X_1, \dots, X_k . Analogously, the validity of a formula $QBE_{k,\forall}$, that is a formula of the form $\forall X_1 \exists X_2 \dots Q X_k f(X_1, \dots, X_k)$, where Q is \forall if k is odd and is \exists if k is even, is complete for Π_k^P . Deciding the conjunction $\Phi \wedge \Psi$, where Φ is a $QBE_{k,\exists}$ formula and Ψ is a $QBE_{k,\forall}$ formula, is complete for D_k^P .

2.2 Defining outliers

Next we formalize the notion of an outlier in default logic. In order to motivate the definition and make it easy to understand, we first look at an example.

Example 1. Consider the following default theory which represents the knowledge that birds fly and penguins are birds that do not fly, and the observations that Tweety and Pini are birds and Tweety does not fly.

$$D = \left\{ \frac{Bird(x) : Fly(x)}{Fly(x)}, \frac{Penguin(x) : Bird(x)}{Bird(x)}, \frac{Penguin(x) : \neg Fly(x)}{\neg Fly(x)} \right\}$$

$$W = \{Bird(Tweety), Penguin(Pini), \neg Fly(Tweety)\}$$

This theory has two extensions. One extension is the logical closure of $W \cup \{Bird(Pini), \neg Fly(Pini)\}$ and the other is the logical closure of $W \cup \{Bird(Pini), Fly(Pini)\}$.

If we look carefully at the extensions, we note that Tweety not flying is quite strange, since we know that birds fly and Tweety is a bird. Therefore, there is no apparent justification for the fact that Tweety does not fly (other than the fact $\neg Fly(Tweety)$ belonging to W). Had we been told that Tweety is a penguin, we could have explained why Tweety does not fly. But, as the theory stands now, we are not able to explain why Tweety does not fly, and, thus, Tweety is an exception. Moreover, if we are trying to nail down what makes Tweety an exception, we notice that if we would have dropped the observation $\neg Fly(Tweety)$ from W , we would have concluded the exact opposite, namely, that Tweety does fly. Thus, $\neg Fly(Tweety)$ “induces” such an exceptionality (we will call *witness* a literal like $\neg Fly(Tweety)$). Furthermore, if we drop from W both $\neg Fly(Tweety)$ and $Bird(Tweety)$, we are no longer able to conclude that Tweety flies. This implies that $Fly(Tweety)$ is a consequence of the fact that Tweety is a bird, and thus $Bird(Tweety)$ is the property of Tweety that behaves exceptionally (or the *outlier*).

From the above example, one could be induced to define an outlier as an individual, i.e., a constant, in our case *Tweety*, that possesses an exceptional property, denoted by a literal having the individual as one of its arguments, in our case $Bird(Tweety)$. However, for a conceptual viewpoint, it is much more general and flexible to single out a property of an individual which is exceptional, rather than simply the individual. That assumed, we also note that within the propositional context we deal with here, we do not explicitly have individuals distinct from their properties and, therefore, the choice is immaterial.

Based on the example and considerations mentioned above, we can define the concept of an outlier as follows.

Definition 1. Let $\Delta = (D, W)$ be a propositional default theory such that W is consistent and $l \in W$ is a literal. If there exists a set of literals $S \subseteq W$ such that:

1. $(D, W_S) \models \neg S$, and
2. $(D, W_{S,l}) \not\models \neg S$.

where $W_S = W \setminus S$ and $W_{S,l} = W_S \setminus \{l\}$, then we say that l is an *outlier* in Δ and S is an *outlier witness set* for l in Δ .

According to this definition, a literal l is an outlier if and only if there is an exceptional property, denoted by a set of literals S , holding in *every* extension of the theory.

The exceptional property is the outlier witness for l . Thus, according to this definition, in the default theory of Example 1 above we should conclude that $Bird(Tweety)$ denotes an outlier and $\{\neg Fly(Tweety)\}$ is its witness. Note that we have defined an outlier witness to be a set, not necessarily a single literal, since in some theories taking a single literal does not suffice to form a witness for a given outlier, being that all witnesses of such an outlier have a cardinality strictly larger than one.

Example 2. Consider the default theory $\Delta = (D, W)$, where the set of default rules D conveys the following information about weather and traffic in a small town in southern California:

1. $\frac{July \wedge Weekend: \neg Traffic_Jam \wedge \neg Rain}{\neg Traffic_Jam \wedge \neg Rain}$ - that is, normally during a July weekend there are no traffic jams nor any rain.
2. $\frac{January: Rain}{Rain}, \frac{January: \neg Rain}{\neg Rain}$ - in January it sometimes rains and sometimes it doesn't rain.
3. $\frac{Weekend \wedge Traffic_Jam: Accident \vee Rain}{Accident \vee Rain}$ - If there is a traffic jam in the weekend then normally it must be raining or there would have been an accident.

Suppose also that $W = \{July, Weekend, Traffic_Jam, Rain\}$. Then, the set $S = \{Traffic_Jam, Rain\}$ is an outlier witness for both *Weekend* and *July*. Moreover, S is a *minimal* outlier witness set for either *Weekend* or *July*, since deleting one of the members from S will render S not being a witness set.

Here is another example.

Example 3. Consider the following default theory Δ :

$$D = \left\{ \frac{Income(x) \wedge Adult(x): Works(x)}{Works(x)}, \frac{FlyingS(x): InterestTakeOff(x)}{InterestTakeOff(x)}, \frac{FlyingS(x): InterestNavigate(x)}{InterestNavigate(x)} \right\}$$

$$W = \{Income(Johnny), Adult(Johnny), \neg Works(Johnny), FlyingS(Johnny), \neg InterestTakeOff(Johnny)\}$$

This theory claims that normally adults who have a monthly income work, and students who take flying lessons are interested in learning how to take off and navigate. The observations are that Johnny is an adult who has a monthly income, but he does not work. He is also a student in a flying school but he is not interested in learning how to take-off. Based on the events of September 11, 2001, we'd like our system to conclude that Johnny is the argument of two outliers. Indeed, the reader can verify that the following facts are true:

1. $(D, W_{\neg Works(Johnny)}) \models Works(Johnny)$,
2. $(D, W_{\neg InterestTakeOff(Johnny)}) \models InterestTakeOff(Johnny)$,
3. $(D, W_{\neg Works(Johnny), Adult(Johnny)}) \not\models Works(Johnny)$, and
4. $(D, W_{\neg InterestTakeOff(Johnny), FlyingS(Johnny)}) \not\models InterestTakeOff(Johnny)$.

Hence, both $\neg Works(Johnny)$ and $\neg InterestTakeOff(Johnny)$ are outlier witnesses, while $Adult(Johnny)$ and $FlyingS(Johnny)$ are outliers. Note that $Income(Johnny)$ is also an outlier, with the witness $\neg Works(Johnny)$.

2.3 Defining outlier detection problems

In order to state the computational complexity of detecting outliers, in the rest of the work we refer to the following problems (also referred to as *queries*) defined for an input default theory $\Delta = (D, W)$:

- $Q0$: Given Δ , does there exist an outlier in Δ ?
 $Q1$: Given Δ and a literal $l \in W$, is there any outlier witness for l in Δ ?
 $Q2$: Given Δ and a set of literals $S \subseteq W$, is S a witness for any outlier l in Δ ?
 $Q3$: Given Δ , a set of literals $S \subseteq W$, and a literal $l \in W$, is S a witness for l in Δ ?

3 General complexity results

In this and in the following section we will analyze the complexity associated with detecting outliers in general, DF, NMU, NU and DNU propositional default theories (this section) and illustrate some tractable classes of theories (Section 4). Here, we limit ourselves to provide, for each of the outlier problems defined above (i.e. $Q0$, $Q1$, $Q2$, and $Q3$), a general discussion of the proof (for the full proofs, see [1]).

Theory \ Query	$Q0$	$Q1$	$Q2$	$Q3$
Propositional	$\Sigma_3^P\text{-c}$	$\Sigma_3^P\text{-c}$	$D_2^P\text{-c}$	$D_2^P\text{-c}$
DF, NMU	$\Sigma_2^P\text{-c}$	$\Sigma_2^P\text{-c}$	$D^P\text{-c}$	$D^P\text{-c}$
NU, DNU	NP-c	NP-c	P	P
Acy. NU, Acy. DNU	P	P	P	P

Table 1. Complexity results for outlier detection

The complexity results are summarized in Table 1, where \mathcal{C} -c stands for \mathcal{C} -complete.

We recall that the entailment problem is Π_2^P -complete for general propositional default theories [13, 7], co-NP-complete for DF [8] and NMU [1] propositional default theories, and in P for NU [8] and DNU [14] propositional default theories.

3.1 Queries $Q0$ and $Q1$

We start commenting about query $Q0$, the most general form of query that we have defined above. Given a default theory, this query asks for the existence of an outlier in the theory. When general propositional default theories are considered, this query is rather complex as it lies at the third level of the polynomial hierarchy.

Theorem 1. *$Q0$ on general propositional default theories is Σ_3^P -complete under polynomial time transformations.*

We note that a problem lying at the k -th level of the polynomial hierarchy is characterized by exactly k independent “sources of complexity”. Each source of complexity consists of a search space composed by an exponential number of candidate solutions. In the case of general propositional default theories, two of the three sources underly to the associated entailment problem, that are (i) the exponential number of generating defaults $D_{\mathcal{E}} \subseteq D$ and, thus, of possible extensions \mathcal{E} of the default theory $\Delta = (D, W_S)$ ($\Delta = (D, W_{S,l})$ resp.), and (ii) the propositional deductive inference needed to check

that $D_{\mathcal{E}}$ generates an extension \mathcal{E} of Δ and that $\neg S \in \mathcal{E}$ ($\neg S \notin \mathcal{E}$ resp.). The third one is determined by the exponential number of subsets of literals $S \cup \{l\}$ of W candidate to play the role of an outlier witness set (the set S) and an outlier (the literal l) in Δ .

If we restrict query $Q0$ to DF theories then the following holds.

Theorem 2. *$Q0$ restricted to DF propositional default theories is Σ_2^P -complete under polynomial time transformations.*

The complexity of $Q0$ for DF theories goes down one level in the polynomial hierarchy w.r.t. general theories as, in this case, the deductive inference check reduces to simple set operations, and, therefore, we are left with only two sources of complexity.

The complexity associated with $Q0$ does not decrease even if we consider such a simplified form of DF theories as NMU theories, as stated below.

Theorem 3. *$Q0$ restricted to NMU propositional default theories is Σ_2^P -complete under polynomial time transformations.*

This result is explained since the complexity of the entailment problem for NMU theories is the same for DF theories.

To obtain a further reduction in complexity, we have to consider simpler theories than the NMU ones.

Theorem 4. *$Q0$ restricted to propositional NU default theories is NP-complete under polynomial time transformations.*

Theorem 5. *$Q0$ restricted to propositional DNU default theories is NP-complete under polynomial time transformations.*

Query $Q0$ on these theories lies at the first level of the polynomial hierarchy since the entailment problem for NU and DNU default theories is polynomial time decidable. Note that, however, this NP-completeness result tells that searching for outliers even in these simple form of theories is a very complex task.

Summarizing, the query $Q0$ is complete for the class Σ_3^P for general default theories, is complete for the class Σ_2^P for DF and NMU propositional default theories, and is complete for the class NP and NU and DNU default theories, hence its complexity lies exactly one level of the polynomial hierarchy over the level of the complexity of the entailment problem for these theories.

As for query $Q1$, hold considerations analogous to that of query $Q0$ and, thus, the complexity results for these two queries are identical.

Intuitively, this can be justified noting that none of the sources of complexity involved in query $Q0$ is relieved by the knowledge of the outlier l . In particular, the number of possible outlier witness set $S \subseteq W \setminus \{l\}$ for l is still exponential.

3.2 Queries Q2 and Q3

Given a default theory and a set of literals S , query $Q2$ asks whether S is a witness for any outlier in the theory. The complexity of $Q2$ lies in the polynomial hierarchy one level below the complexity of $Q0$. Indeed, one of the sources of complexity involved

with query $Q0$, that is the exponential number of outlier witnesses, falls off when query $Q2$ is considered.

In particular, $Q2$ on general theories is the conjunction of two independent problems, one from Π_2^P and one from Σ_2^P .

Theorem 6. *$Q2$ on general propositional default theories is D_2^P -complete under polynomial time transformations.*

For DF and NMU default theories, $Q2$ is the conjunction of two independent problems from co-NP and NP respectively.

Theorem 7. *$Q2$ restricted to DF propositional default theories is D^P -complete under polynomial time transformations.*

Theorem 8. *$Q2$ restricted to NMU propositional default theories is D^P -complete under polynomial time transformations.*

Finally, $Q2$ on NU and DNU theories is the conjunction of two polynomial time solvable problems, and, hence, in these cases, this query is in P.

Theorem 9. *$Q2$ restricted to NU and DNU propositional default theories is in P.*

Finally, consider query $Q3$. This query is important as it constitutes the basic operator to be implemented in a system of outlier detection on propositional default theories. Indeed, given a default theory, set of literals S and a literal l , this query asks whether S is an outlier witness set for l in the input theory.

Analogously to what holding between query $Q0$ and $Q1$, both query $Q2$ and $Q3$ have the same complexity.

4 Some Tractable Cases

In this section, we look for some classes of default theories for which outlier detection is computationally tractable.

Restricting our attention to NMU, NU, or DNU theories does not suffice to attain tractability of the most general queries $Q0$ and $Q1$. Some further restriction is needed, which is considered next.

Theorem 10 ([8] [14]). *Suppose Δ is a normal (dual normal) unary default theory. We can decide whether a literal belongs to every extension of Δ in time $O(n^2)$, where n is the length of the theory.*

Definition 2. The *atomic dependency graph* of an NMU default theory Δ is a directed graph whose nodes are all atoms in the language of Δ , and such that there is an arc directed from p to q iff there is a default in Δ in which p or $\neg p$ is a prerequisite and q or $\neg q$ is a consequence.

Definition 3. A normal (dual normal) unary default theory is *acyclic* iff its atomic dependency graph is acyclic.

Theorem 11. *Queries $Q0$, $Q1$, $Q2$ and $Q3$, restricted to the class of acyclic NU or acyclic DNU default theories can be solved in polynomial time in the size of the input theory.*

Proof. It can be shown that for any acyclic NMU default theory $\Delta = (D, W)$ such that W is consistent and for any literal l in W , any minimal outlier witness set for l in Δ is at most 1 in size. Theorem's statement then follows from Theorem 10.

5 Related Work

The research on logical-based abduction [11, 2, 3] is related to outlier detection. In the framework of logic-based abduction, the domain knowledge is described using a logical theory T . A subset X of hypotheses is an abduction explanation to a set of manifestations M if $T \cup X$ is a consistent theory that entails M . Abduction resembles outlier detection in that it deals with exceptional situations.

The work most relevant to our study is perhaps the paper by Eiter, Gottlob, and Leone on abduction from default theories [4]. There, the authors have presented a basic model of abduction from default logic and analyzed the complexity of the main abductive reasoning tasks. They presented two modes of abductions: one based on brave reasoning and the other on cautious reasoning. According to those authors, a default abduction problem (DAP) is a tuple $\langle H, M, W, D \rangle$ where H is a set of ground literals called *hypotheses*, M is a set of ground literals called *observations*, and (D, W) is a default theory. Their goal, in general, was to explain some observations from M by using various hypotheses in the context of the default theory (D, W) . They suggest the following definition for an explanation:

Definition 4 ([4]). *Let $P = \langle H, M, D, W \rangle$ be a DAP and let $E \subseteq H$. Then, E is a skeptical explanation for P iff*

1. $(D, W \cup E) \models M$, and
2. $(D, W \cup E)$ has a consistent extension.

There is a close relationship between outliers and skeptical explanations, as the following theorem states. The theorem also holds for ordered semi-normal default theories [5].

Theorem 12. *Let $\Delta = (D, W)$ be a normal default theory, where W is consistent. Let $l \in W$ and $S \subseteq W$. S is an outlier witness set for l iff $\{l\}$ is a minimal skeptical explanation for $\neg S$ in the DAP $P = \langle \{l\}, \neg S, D, W_{S,l} \rangle$*

Hence, we can say that S is an outlier witness for l if $l \in W$, l is a skeptical explanation for S , but still $\neg S$ holds in every extension of the theory.

Despite the close relationship between outlier detection and abduction demonstrated by the above theorem (especially for normal defaults) we believe that there is a significant difference between the two concepts. In abduction, we have to single out a set of manifestations and a set of potential explanations. Outlier detection, on the other hand, has much more to do with knowledge discovery. The task in outlier detection is to *learn*

who the exceptionals (the outliers), or the suspects, if you wish, are, and to justify the suspicion (that is, list the outlier witnesses).

It also turns out that reducing outlier detection queries to abduction and vice versa is not straightforward, and therefore, when analyzing the computational complexities involved in answering outlier detection queries we have preferred to use the classical Boolean formula satisfiability problems (see [1]).

6 Conclusion

Suppose you are walking down the street and you see a blind person walking in the opposite direction. You believe he is blind because he is feeling his way with a walking stick. Suddenly, something falls out of his bag, and to your surprise, he finds it immediately without probing about with his fingers, as you would expect for a blind person. This kind of behavior renders the “blind” person suspicious.

The purpose of this paper has been to formally mimic this type of reasoning using default logic. We have formally defined the notion of an outlier and an outlier witness, and analyzed the complexities involved, pointing out some non-trivial tractable subsets. As explained in the introduction, outlier detection can also be used for maintaining database integrity and completeness.

This work can be extended in several ways. First, we can develop the concept of outliers in other frameworks of default databases, like System Z [10] and Circumscription [9]. Second, we can look for intelligent heuristics that will enable us to perform the involved heavy computational task more efficiently. Third, we can study the problem from the perspective of default theories as a “semantic check toolkit” for relational databases.

References

1. Fabrizio Angiulli, Rachel Ben-Eliyahu-Zohary, and Luigi Palopoli. Outlier detection using default logic, 2003. Full manuscript.
2. Luca Console, Daniele Theseider Dupré, and Pietro Torasso. On the relationship between abduction and deduction. *Journal of Logic and Computation*, 1(5):661–690, 1991.
3. Thomas Eiter and Georg Gottlob. The complexity of logic-based abduction. *JACM*, 42(1):3–42, 1995.
4. Thomas Eiter, Georg Gottlob, and Nicola Leone. Semantics and complexity of abduction from default theories. *Artificial Intelligence*, 90(1-2):177–223, 1997.
5. David W. Etherington. Formalizing nonmonotonic reasoning systems. *Artificial Intelligence*, 31:41–85, 1987.
6. M. R. Garey and D. S. Johnson. *Computers and intractability: a guide to the theory of NP-completeness*. W.H. Freeman, 1979.
7. George Gottlob. Complexity results for nonmonotonic logics. *J. Logic Comput.*, 2(3):397–425, 1992.
8. Henry A. Kautz and Bart Selman. Hard problems for simple default logics. *Artificial Intelligence*, 49:243–279, 1991.
9. John McCarthy. Circumscription - a form of non-monotonic reasoning. *Artificial Intelligence*, 13:27–39, 1980.

10. Judea Pearl. System z: a natural ordering of defaults with tractable applications to nonmonotonic reasoning. In *TARK-90*, pages 121–135, Monterey, CA, 1990.
11. David Poole. Normality and faults in logic-based diagnosis. In *IJCAI-89*, pages 1304–1310, 1989.
12. Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.
13. Jonathan Stillman. The complexity of propositional default logics. In *AAAI-92: Proceedings of the 10th national conference on artificial intelligence*, pages 794–799, 1992.
14. Rachel Ben-Eliyahu Zohary. Yet some more complexity results for default logic. *Artificial Intelligence*, 139(1):1–20, 2002.