

Abduction through Semantic Tableaux versus Abduction through Goal-Directed Proofs¹

JOKE MEHEUS and DAGMAR PROVIJN

Centre for Logic and Philosophy of Science, Ghent University
Blandijnberg 2, 9000 Ghent, Belgium
Joke.Meheus@UGent.be
Dagmar.Provijn@UGent.be

ABSTRACT

In this paper, we present a goal-directed proof procedure for abductive reasoning. This procedure will be compared with Aliseda's approach based on semantic tableaux. We begin with some comments on Aliseda's algorithms for computing conjunctive abductions and show that they do not entirely live up to their aims. Next we give a concise account of goal-directed proofs and we show that abductive explanations are a natural spin-off of these proofs. Finally, we show that the goal-directed procedure solves the problems we encountered in Aliseda's algorithms.

KEYWORDS: Abduction, Semantic Tableaux, Goal-directed proof procedures

1. Introduction

Very few philosophers of science will object to the idea that abduction plays a key role in scientific reasoning. Still, the logical foundations of abduction have been seriously neglected, both by logicians and philosophers of science. For centuries, the former have discarded abduction as a fallacy ("affirming the consequent"), about which nothing 'decent' can be said. The latter are usually more sympathetic to the concept, despite (or, as is sometimes the case, thanks to) the fact that also they believe there is nothing logical about it.

There is, however, a gradual change in attitude. Two developments seem important for this. On the one hand, the advent of non-standard logics (especially non-monotonic ones) made it possible to broaden the domain of logic to include all kinds of reasoning forms that are traditionally viewed as non-logical. On the other hand, abduction has been intensively studied within computer science and Artificial Intelligence, which led to a large variety of logic-based approaches to abduction.

Important as these developments are, the results are fragmented and their application to problems of philosophy of science are often far from evident. Aliseda (2006) is a welcome exception to this. In this book, Aliseda presents a general framework for the logical study of abduction, that has roots in formal logic as well as computer science, and shows how it can be applied to problems in the philosophy of science (such as explanation, empirical progress and epistemic change). We mention only a few examples of what we consider to be central contributions: the taxonomy for abduction (which is the most elaborated one available in the literature) and the analysis of the different "abductive styles", the attention for "abductive anomaly" (which is too often ignored), and the structural characterization of abductive inference. In view of this, we are convinced that the book will soon count as a standard reference for the study of abduction.

¹ Research for this paper was supported by subventions from Ghent University and from the Research Foundation – Flanders (FWO - Vlaanderen). The authors are indebted to Christian Straßer for checking some prolog code.

The reader will forgive us for adding a personal note. For about ten years now, the first author has been convinced that abduction should be modelled as a kind of “forward reasoning”, in which one tries to derive the explanations *from* the theory together with the *explanandum*.² This resulted in a number of papers on formal logics for abduction that lead from a theory and one or more *explananda* to a set of possible explanations (see, for instance, Meheus and Batens (2006)). It is under the influence of Aliseda’s book that Meheus began to see a number of limitations to this particular approach and that she started to look elsewhere.³ This eventually led to a cooperation with the second author, who, for applications totally unrelated to abduction, had been working on a specific type of goal-directed proofs (see Batens and Provijn (2001)).

The goal-directed proofs from Batens and Provijn (2001) bear similarities with semantic tableaux (which are at the heart of Aliseda’s computational approach to abduction). For instance, unlike ordinary proofs, these goal-directed proofs form a decision method for $A_1, \dots, A_n \vdash B$. Moreover, they are primarily based on the analysis of formulas. One of the differences is that, for most sets of premises, goal-directed proofs are more efficient (require less computation) than semantic tableaux.

As it turns out, abductive explanations are a natural spin-off of goal-directed proofs. Moreover, they seem a nice compromise between forward and backward approaches: although the *explanandum* is not included among the premises, it is the starting point of the derivation, and moreover directs the inferences that are made.

In this paper, we shall present the basic ideas behind goal-directed proofs for abduction and make some comparisons with Aliseda’s approach based on semantic tableaux. We will necessarily have to be brief, but we refer the reader to Meheus and Provijn (To Appear) for a more detailed discussion.

We begin with some comments on Aliseda’s approach, acknowledging at once that, in this very short contribution, we shall not be able to do justice to it.

2. Abduction Through Semantic Tableaux

In this section, we shall argue that Aliseda’s algorithms for computing *conjunctive* abductions by means of tableaux (Aliseda, 2006, pp. 113-116) do not entirely live up to their aims: the procedures do not warrant that the generated explanations are non-redundant (in the sense of the definition on p. 111) and that the one for *consistent* conjunctive abductions moreover does not warrant consistency.

We begin with the algorithm for conjunctive plain abductions (Aliseda, 2006, p. 113). It is claimed in footnote 3 that the algorithm does not lead to redundant solutions, because each conjunctive explanation is a conjunction of *partial* explanations. However, in instruction 3 of the algorithm, conjunctions are constructed that contain one literal from *each* $BPC(\Gamma_i)$;⁴ in the next and final instruction only *repeated* conjuncts are eliminated. Hence, there is no

² Abduction is usually viewed as a kind of “backward reasoning”: “given an *explanandum* B and a theory T, find an explanation A such that B can be derived from T and A”.

³ One limitation is that the logics seem suitable only *one* of Aliseda’s abductive styles, namely the one where one is interested in the *minimal* abductions.

⁴ $BPC(\Gamma_i)$ – the Branch Partial Closure of Γ_i – is the set of literals that close the open branch Γ_i but do not close all the other open branches of the tableau.

warrant that the conjunctions do not contain conjuncts that, in view of the other conjuncts, are superfluous to close the tableau.⁵ Here is an example.⁶

Example 1. $\Theta = \{((p \ \& \ q) \ \& \ r) \supset s, ((p \ \& \ q) \ \vee \ t) \supset s\}$ $\varphi = s$.

The only non-redundant conjunctive explanation is $p \ \& \ q$, but also $(p \ \& \ q) \ \& \ r$ is generated.

One way to remedy this first problem is to add an instruction that eliminates all non-minimal conjunctions – we call an abductive solution $\bigwedge(\Sigma)$ minimal iff there is no $\Sigma' \subset \Sigma$ for which $\bigwedge(\Sigma')$ is also an abductive solution.

The procedure for conjunctive consistent explanations (p. 115) is plagued by the same problem. Actually, the problem is even worse, because now *all* possible conjunctions of partial explanations are generated as abductive solutions. Moreover, although it is claimed that the “production of any inconsistency whatsoever” is avoided (Aliseda, 2006, pp. 114), the procedure does not warrant this. Here are but two examples – note that in each case the theory is consistent and the *explanandum* is compatible with it.

Example 2. $\Theta = \{p \supset q, r \vee s\}$ $\varphi = q$.

The procedure leads to $\sim r \ \& \ \sim s$ as a consistent explanation for q , which is incompatible with $r \vee s$.

Example 3. $\Theta = \{p \supset q, (r \ \& \ \sim r) \supset q\}$ $\varphi = q$.

In view of the *tautology* $(r \ \& \ \sim r) \supset q$, the *contradiction* $r \ \& \ \sim r$ is obtained as a consistent explanation for q .

The cause for this second problem is simple: even if the literals are compatible with the theory (which is checked in instruction 3), their conjunctions need not be. Hence, what is needed is an additional instruction in which those ρ_i from instruction 6 are selected for which $\mathcal{T}(\Theta + \rho_i)$ is a semi-closed extension.

These additional instructions are needed also for another reason. On p. 116 it is claimed that, in order to modify the algorithms to handle the production of *explanatory* abductions, one only needs to avoid self-explanations (the case where the abductive solution α is identical to the *explanandum* φ). Aliseda justifies this claim by her Fact 4 – the idea that, given *any* Θ and φ , the algorithms never produce abductive solutions α such that $\alpha \models \varphi$. As in some cases abductive solutions will be generated that are internally inconsistent (see example 3), this Fact does not hold in general.⁷ It also does not hold true in general (as Atocha claims on p. 114) that the trivial solution (where $\alpha = \varphi$) can be avoided by first running the algorithm for atomic explanations. The reason is that the algorithms for atomic explanations do not always rule out self-explanations. For instance, where $\Theta = \{s \supset p, p \supset s\}$, p is generated as an explanation for p . The following is a bit less obvious:⁸

⁵ According to Aliseda (personal communication), we are mistaken in not taking into account that what can be found on p. 113 of Aliseda (2006) is only a *sketch* of the algorithm as it is presented in Aliseda (1997) (see also the claims on p. 110 of Aliseda (2006)). On her view, the redundancy problem does not occur because after the last instruction of the algorithm, the generated solutions have to be checked against the conditions listed in the output of the algorithm.

⁶ We follow Aliseda's convention to use Θ as metavariable for background theories and φ as a metavariable for *explananda*.

⁷ Given the generality of Aliseda's Fact 4, it cannot be held against us that we chose an example in which one of the premises is a tautology.

⁸ On Aliseda's procedures, any theory from which the explanandum is not derivable can be transformed into one that leads to atomic self-explanations: where φ is the *explanandum*, it suffices to add $\varphi \vee \sim\varphi$ to the theory.

Example 4. $\Theta = \{s \supset p, (p \ \& \ (q \vee r)) \supset s\}$ $\varphi = p$.
 Not only s is produced as an explanation for p , but also p .

We end this section with a remark on the notion of “partial explanation”. On Aliseda’s account, given a theory Θ and *explanandum* φ , a formula β counts as a partial explanation for φ iff $\mathcal{T}((\Theta + \sim\varphi) + \beta)$ is semi-closed. This definition is one of the main reasons why, in the abductive solutions, irrelevant conjuncts may creep in: the definition does not warrant that there is any relation between the ‘partial explanation’ and the *explanandum*. One may respond to this that the problem disappears if only *relevant* premises are included in the theory. However, this is setting the cart before the horse. One of the main problems in searching for an abduction is precisely to identify the relevant premises. And, in all interesting cases, this problem is only solved when the abductive explanations have been generated.

In Section 4, we shall see that this problem is dealt with in a very natural way by goal-directed proofs. But first we present a very brief introduction to the format, which was elaborated in Batens and Provijn (2001).

3. Goal-Directed Proofs

The basic idea is that the formal elements of the search process, involved in the construction of a proof for $\Gamma \vdash G$, are pushed in the proof itself. In this section, Γ always denotes the premise set and G is used to refer to the main goal of the search process. In a goal-directed proof, formulas are derived that have the following form

$$[B_1, \dots, B_n] A$$

which indicates that A is derived on the condition $[B_1, \dots, B_n]$.⁹ The rules of the inference system are such that: ‘if $[\Delta]A$ is derivable from a set of premises Γ then $\Gamma \cup \Delta \vdash A$.’

A goal-directed proof starts with the introduction of the main goal by writing

$$[G] G$$

on the first line of the proof. This line is logically redundant but, as we shall see, guides the search procedure. The condition of $[\Delta]A$ reminds us that these formulas have not yet been derived and that they should be derived in order to derive A .

If G is a *positive part* of a premise A , it is allowed that A is entered in the proof and, if G is different from A , that A is analyzed. If G is not *positive part* of a premise, the condition of $[G]G$ is analyzed in such a way that a set of new goals is obtained that will guide the search process.

Inference systems satisfying this proof format allow for perspicuous and simple heuristics that warrant goal-directed and efficient proofs. However, in this introduction, we have to omit the discussion of the heuristics – we refer the reader to Batens and Provijn (2001) and to Meheus and Provijn (To Appear) for this.

The goal-directed proof procedure also needs *marking definitions*. Lines in a proof are *marked* if some goal in the condition of that line is useless in order to derive the main goal or a derived goal.

For a concise formulation of the *positive part* relation and the inference rules we distinguish between *a*- and *u*-formulas, based on a theme from Smullyan (1995). Let $*A$ denote the ‘complement’ of A , viz. B if A has the form $\sim B$ and $\sim A$ otherwise.

⁹ If the condition is empty (notation: $[\emptyset]A$ or simply A), A is said to be derived unconditionally.

\mathfrak{a}	\mathfrak{a}_1	\mathfrak{a}_2	\mathfrak{b}	\mathfrak{b}_1	\mathfrak{b}_2
$A \& B$	A	B	$\sim(A \& B)$	$*A$	$*B$
$A \equiv B$	$A \supset B$	$B \supset A$	$\sim(A \equiv B)$	$\sim(A \supset B)$	$\sim(B \supset A)$
$\sim(A \vee B)$	$*A$	$*B$	$A \vee B$	A	B
$\sim(A \supset B)$	A	$*B$	$A \supset B$	$*A$	B
$\sim\sim A$	A	A			

The following clauses constitute a recursive definition of the *positive part* relation for propositional classical logic (henceforth **CL**):

1. $pp(A, A)$.
2. $pp(A, \mathfrak{a})$ if $pp(A, \mathfrak{a}_1)$ or $pp(A, \mathfrak{a}_2)$.
3. $pp(A, \mathfrak{b})$ if $pp(A, \mathfrak{b}_1)$ or $pp(A, \mathfrak{b}_2)$.
4. $pp(A, B)$ and $pp(B, C)$, then $pp(A, C)$.

We now move to the instructions for **CL**.¹⁰ Two general restrictions have to be taken into account:

- R1** Formula analyzing rules are not applied on formulas introduced by the Goal rule.
R2 No rule is applied to repeat a marked or unmarked line.

The instruction *Goal* introduces the main goal in the proof:

Goal Start a goal-directed proof with:

1 [G] G Goal

Premises are only introduced in the proof if a goal of an unmarked line is a positive part of it:

Prem If A is a goal of an unmarked line, $B \in \Gamma$ and $pp(A, B)$, then one may add:

k B Prem

If a goal A is a *positive part* of a formula B that was introduced by Prem, the formula analyzing rules allow one to analyze B until $[\Delta]A$ is derived on a line in the proof. The formula analyzing rules can be summarized as follows:¹¹

$$\frac{[\Delta] \mathfrak{a}}{[\Delta] \mathfrak{a}_1 \quad [\Delta] \mathfrak{a}_2} \qquad \frac{[\Delta] \mathfrak{b}}{[\Delta \cup \{ * \mathfrak{b}_2 \}] \mathfrak{b}_1 \quad [\Delta \cup \{ * \mathfrak{b}_1 \}] \mathfrak{b}_2}$$

The general form of the rules is $[\Delta]A / [\Delta \cup \Delta']B$. Their application is governed by the following instruction (in which R refers to the name of the analyzing rule):

¹⁰ For reasons of space, we present a version in which the instructions are immediately linked to the inference rules – see Batens (2003). This version is less deterministic than the one presented in Batens and Provijn (2001).

¹¹ If two formulas occur at the bottom line of a rule, both variants may be derived (separately) in the proof.

FAR If C is a goal of an unmarked line, $[\Delta]A$ is the formula of an unmarked line i , $[\Delta]A / [\Delta \cup \Delta']B$ is a formula analyzing rule, and $pp(C, B)$, then one may add:

$$k \quad [\Delta \cup \Delta'] B \quad i \quad R$$

The names of the formula analyzing rules are $\dagger E$, with $\dagger \in \{\&, \vee, \supset, \sim, \equiv\}$ or $\sim\dagger E$ with $\dagger \in \{\&, \vee, \supset, \equiv\}$.

If no goal is a positive part of a premise or an analyzed formula, the condition analyzing rules lead to the analysis of the available goals. These rules can be summarized as:

$$\frac{[\Delta \cup \{a\}] A}{[\Delta \cup \{a_1, a_2\}] A} \qquad \frac{[\Delta \cup \{b\}] A}{[\Delta \cup \{b_1\}] A \quad [\Delta \cup \{b_2\}] A}$$

The general form of the rules is $[\Delta \cup \{B\}]A / [\Delta \cup \Delta']A$. Their application is governed by:

CAR If A is a goal of an unmarked line, $[\Delta \cup \{B\}]A$ is the formula of an unmarked line i , $[\Delta \cup \{B\}]A / [\Delta \cup \Delta']A$ is a condition analyzing rule, then one may add:

$$k \quad [\Delta \cup \Delta'] A \quad i \quad R$$

The names of the condition analyzing rules are equal to the names of the formula analyzing rules preceded by a C .

As $A \vee \sim A$ is valid in **CL**, *Excluded Middle* allows for the elimination of certain goals by the following instruction:

EM If A is a goal of an unmarked line, $[\Delta \cup \{B\}]A$ and $[\Delta' \cup \{\sim B\}]A$ are the respective formulas of the unmarked lines i and j , and $\Delta \subseteq \Delta'$ or $\Delta' \subseteq \Delta$, then one may add:

$$k \quad [\Delta \cup \Delta'] A \quad i, j \quad EM$$

The instruction *Transitivity* allows both for the elimination of goals that are derived unconditionally and for the generation of alternative conditions (if the goals of a certain condition are themselves conditionally derived in the proof):

Trans If A is a goal of an unmarked line, and $[\Delta \cup \{B\}]A$ and $[\Delta']B$ are the respective formulas of the unmarked lines i and j , then one may add:

$$k \quad [\Delta \cup \Delta'] A \quad i, j \quad TRANS$$

As we are concerned here with the goal-directed generation of abductive explanations from *consistent* and *finite* premise sets, both the instruction for the application of *Ex Falso Quodlibet* (EFQ)¹² and the restrictions for infinite premise sets can be skipped (see Batens and Provijn (2001)).

The last elements of the procedure are the *marking definitions*. For the current application, there are three reasons for *marking* a line: redundancy, inconsistency and loops.

¹² The instruction EFQ is completely isolated. Without it, a paraconsistent variant of the **CL** procedure is obtained – see Batens (To Appear).

A line on which $[\Delta \cup \Delta']A$ has been derived is redundant and hence *R-marked* if $[\Delta]A$ has been derived in the proof. Evidently, searching for the members of Δ' is useless to obtain A.

Definition 1 Line i on which $[\Delta]A$ is derived, is R-marked on a stage of a proof if on that stage $[\Delta']A$ is derived and $\Delta' \subset \Delta$.

A condition contains a flat inconsistency if both A and $\sim A$ occur in it. If the derivation of G relies on the derivation of an inconsistency, the procedure of Batens and Provijn (2001) takes care of this by means of EFQ. As abductive explanations should not be generated by means of inconsistencies, lines of which the condition contains a flat inconsistency are *I-marked*.

Definition 2 Line i on which $[\Delta]A$ is derived, is I-marked if Δ is flatly inconsistent.

Lines on which $[\Delta \cup \{A\}]A$ is derived can only lead to loops in the search process and are *L-marked*. This also indicates that the search process for A should be led by other conditions, if possible.

Definition 3 Line i on which $[\Delta]A$ is derived, is L-marked if $A \in \Delta$, unless line i was introduced by means of the Goal rule.

A proof is *finished* whenever G is derived. A proof is *stopped* if it is finished or if no further instructions can be applied.

4. Abduction through Goal-Directed Proofs

Recognizing abductive solutions on the basis of goal-directed proofs is absolutely straightforward. Here are the definitions:

Definition 4 Given a theory Θ and an explanandum φ , A is an *atomic explanation* for φ iff A is a literal and $[A]\varphi$ is derived on an unmarked line (different from the goal line) in a proof for $\Theta \vdash \varphi$ that is *stopped* but *not finished*.

Definition 5 Given a theory Θ and an explanandum φ , $A_1 \& \dots \& A_n$ is a *conjunctive explanation* for φ iff A_1, \dots, A_n are literals and $[A_1, \dots, A_n]\varphi$ ($n > 1$) is derived on an unmarked line in a proof for $\Theta \vdash \varphi$ that is *stopped* but *not finished*.

As is shown in Meheus and Provijn (To Appear), the requirement that the proof should be stopped but not finished warrants, together with the marking rules, that all abductive solutions that satisfy these definitions are *explanatory* (in the sense of (Aliseda, 2006, p. 74). That the proof should be stopped moreover warrants that all possible abductive solutions occur in the proof.

In the above definitions, we assume, like Aliseda, that φ is a literal. However, as is shown in Meheus and Provijn (To Appear), the approach can easily be generalized to handle other forms of *explananda*.

We shall now review the examples from Section 2 in terms of goal-directed proofs. All proofs are generated according to the heuristics discussed in Meheus and Provijn (To Appear).

The first example illustrates how goal-directed proofs warrant the non-redundancy of the abductive explanations.

Example 1. $\Theta = \{((p \ \& \ q) \ \& \ r) \supset s, ((p \ \& \ q) \ \vee \ t) \supset s\}$ $\varphi = s$.

1	[s] s		Goal
2	$((p \ \& \ q) \ \& \ r) \supset s$		Prem
3	$[(p \ \& \ q) \ \& \ r] s$	2	$\supset E$
4	$[p \ \& \ q, r] s$	3	C&E
5	$[p, q, r] s$	4	C&E R
6	$((p \ \& \ q) \ \vee \ t) \supset s$		Prem
7	$[(p \ \& \ q) \ \vee \ t] s$	6	$\supset E$
8	$[p \ \& \ q] s$	7	$C \vee E$
9	$[t] s$	7	$C \vee E$
10	$[p, q] s$	8	C&E

At stage 10 of the proof, line 5 is R-marked and the proof is stopped. Hence, $(p \ \& \ q) \ \& \ r$ is not retained as an explanation for s .

On Aliseda's method, the next two examples lead to inconsistent explanations. In the goal-directed proofs, the inconsistent explanations are either not generated or the lines on which they occur are I-marked.

Example 2. $\Theta = \{p \supset q, r \vee s\}$ $\varphi = q$.

As both q and p are not a positive part of $r \vee s$, the latter is never entered in the proof. Hence, $[\sim r, \sim s] q$ cannot be derived.

Example 3. $\Theta = \{p \supset q, (r \ \& \ \sim r) \supset q\}$ $\varphi = q$.

1	[q] q		Goal
2	$p \supset q$		Prem
3	[p] q	2	$\supset E$
4	$(r \ \& \ \sim r) \supset q$		Prem
5	$[r \ \& \ \sim r] q$	4	$\supset E$
6	$[r, \sim r] q$	5	C&E I

As soon as line 6 is entered, it is I-marked. Hence, in the stopped proof, only p is retained as an explanation for q .

The final example illustrates how self-explanations are ruled out.

Example 4. $\Theta = \{s \supset p, (p \ \& \ (q \ \vee \ r)) \supset s\}$ $\varphi = p$.

1	[p]		Goal
2	$s \supset p$		Prem
3	[s] p	2	$\supset E$
4	$(p \ \& \ (q \ \vee \ r)) \supset s$		Prem
5	$[p \ \& \ (q \ \vee \ r)] s$	4	$\supset E$
6	$[p, q \ \vee \ r] s$	5	C&E
7	$[p, q \ \vee \ r] p$	3,6	Trans L

Some more lines may be added to this proof (depending on the heuristics followed). For instance, $C\vee E$ can be applied to line 6, and $Trans$ to the resulting line and line 3. However, all these lines will be L-marked. Hence, the only explanation for p that is generated is s .

5. Comparing the Two Methods

Like tableaux, the goal-directed proofs presented here are essentially based on α - and β -type transformation rules. Still, as we discuss in more length in Meheus and Provijn (To Appear), generating abductions on the basis of goal-directed proofs seem to have a number of advantages. First, the procedure is (in general) more efficient. One of the reasons for this is that only *relevant* premises are introduced in the proof.¹³ Next, problems of redundancy, irrelevance, inconsistency and (partial or total) self-explanations can be dealt with in a very efficient and transparent way. Finally, the format can easily be generalized to handle other kinds of abductive inference, such as the case in which the *explanandum* is not a literal, the case in which the *explanandum* is inconsistent with the theory (“abductive anomalies”), and the case in which the theory itself is inconsistent.¹⁴

A possible disadvantage of the goal-directed proofs is that we see no elegant way to handle Aliseda’s disjunctive explanations (although we should add that we are not entirely convinced of the unrestricted way in which disjunctive abductions are generated in Aliseda’s tableaux either). Another advantage of tableaux is their graphical format (but see Batens (2006) for a diagrammatic representation of goal-directed proofs).

There is a more philosophical reason why we think the comparison between the two methods is useful. Most of the problems that we discussed in Section 2 could be considered as cases in which something is ‘wrong’ with the background theory (it contains tautologies, or irrelevant premises, or redundant premises, etc.). And, as Aliseda observes on p. 117, one’s algorithm is not necessarily to be blamed for this: “Bad theories produce bad explanations”. Unfortunately, in this messy and complex world, with our limited cognitive capacities and limited resources, we often have to work with ‘bad theories’. The more a method safeguards us from deriving bad explanations from such theories the better, so it seems. At least, if the cost for doing so is not too high.

This is the real discussion that we hope our comparison will contribute to.

References

Aliseda, A. (1997). Seeking Explanations: Abduction in Logic, Philosophy of Science and Artificial Intelligence. PhD thesis. University of Amsterdam, The Netherlands: Institute for Logic, Language and Computation (ILLC).

Aliseda, A. (2006). Abductive Reasoning. Logical Investigations into Discovery and Explanation. Dordrecht: Springer.

¹³ The conditions under which a premise is *relevant* in the search for an abductive explanation for an *explanandum* ϕ can be defined in a very precise way on the basis of goal-directed proofs – see Meheus and Provijn (To Appear).

¹⁴ The last case is adequately handled by simply omitting the rule EFQ from the procedure (which is a quite unnatural rule anyway).

Batens, D. (2003). "A formal approach to problem solving". In: Delrieux, C.; Legris, J. (Eds.), *Computer Modeling of Scientific Reasoning*, 15-26. Bahia Blanca, Argentina: Universidad Nacional del Sur.

Batens, D. (2006). "A diagrammatic proof search procedure as part of a formal approach to problem solving". In: Magnani, L. (Ed.), *Model Based Reasoning in Science and Engineering*, 265-284. London: King's College Publications.

Batens, D. (To Appear). "A paraconsistent proof procedure based on classical logic".

Batens, D.; Provijn, D. (2001). "Pushing the search paths in the proofs. A study in proof heuristics". *Logique et Analyse* 173-174-175, 113-134. Appeared 2003.

Meheus, J.; Batens, D. (2006). "A formal logic for abductive reasoning". *Logic Journal of the IGPL* 14, 221-236.

Meheus, J.; Provijn, D. (To Appear). "A goal-directed proof procedure for abductive reasoning".

Smullyan, R. M. (1995). *First Order Logic*. New York: Dover. Original edition: Springer, 1968.