

An Abductive Reasoning Approach to the Belief-Bias Effect

Luís Moniz Pereira^{*1}
Emmanuelle-Anna Dietz^{†1,2}
Steffen Hölldobler^{‡2}

¹Centro de Inteligência Artificial (CENTRIA), Departamento de Informática
Faculdade de Ciências e Tecnologia, Universidade Nova de Lisboa, 2829-516 Caparica, Portugal
²International Center for Computational Logic (ICCL), TU Dresden, D-01062 Dresden, Germany

Abstract

The tendency to accept or reject arguments based on own beliefs or prior knowledge rather than on the reasoning process is called the belief-bias effect. A psychological syllogistic reasoning task shows this phenomenon, wherein participants were asked whether they accept or reject a given syllogism. We discuss one case which is commonly assumed to be believable but not logically valid. By introducing abnormalities, abduction and background knowledge, we model this case under the weak completion semantics. Our formalization reveals new questions about observations and their explanations which might include some relevant prior abductive contextual information concerning some side-effect. Inspection points, introduced by Pereira and Pinto, allow us to express these definitions syntactically and intertwine them into an operational semantics.

Keywords: Abductive Reasoning; Side-effects; Inspection Points; Human Reasoning; Belief-Bias Effect

1 Introduction

Whenever discovering abductive explanations for some given primary observation, one may wish to check too whether some other given additional secondary observations are logical consequences of the abductive explanations found for the primary observation. The primary explanations may be used to explain the secondary observations, whereas the secondary explanations cannot be used to explain the primary observations. We show this type of reasoning requires the characterization of a new abduction concept and mechanism, that of *contextual abduction*. We formalize contextual abduction and employ it to understand and justify the belief-bias effect in human reasoning by addressing syllogisms.

Evans, Barston, and Pollard (1983) made a psychological study showing possibly conflicting processes in human reasoning. Participants were presented syllogisms and had to decide whether they were logically valid. Consider S_{add} :

PREMISE1 *No addictive things are inexpensive (not costly).*
PREMISE2 *Some cigarettes are inexpensive (not costly).*
CONCLUSION *Therefore, some addictive things are not cigarettes.*

^{*}Imp@fct.unl.pt

[†]{dietz,sh}@iccl.tu-dresden.de

Copyright © 2014, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Although the conclusion does not necessarily follow from the premises, participants assumed the syllogism to be logically valid. Evans, Barston, and Pollard (1983) concluded that they have been unduly influenced by their own beliefs. In the following, we explain the belief-bias effect by abductive reasoning and its side-effects.

2 Preliminaries

We restrict ourselves to datalog (first-order) programs, ie. the set of terms consists only of constants and variables. A logic program \mathcal{P} is a finite set of clauses of the form

$$A \leftarrow A_1 \wedge \dots \wedge A_n \wedge \overline{B_1} \wedge \dots \wedge \overline{B_m},$$

where A and A_i , $1 \leq i \leq n$, are atoms and $\overline{B_j}$, $1 \leq j \leq m$, are negated atoms. A is the head and $A_1 \wedge \dots \wedge A_n \wedge \overline{B_1} \wedge \dots \wedge \overline{B_m}$ is the body of the clause. We abbreviate the body of a clause by *body*. $A \leftarrow \top$ and $A \leftarrow \perp$ are special cases of clauses denoting positive and negative facts, respectively. If an argument is written with an upper case letter, it is a variable; otherwise it is a constant. In the sequel, we assume \mathcal{P} to be ground, containing all the ground instances of its clauses. The set of all atoms occurring in \mathcal{P} is $\text{atoms}(\mathcal{P})$. An atom is undefined in \mathcal{P} if it is not the head of some clause in \mathcal{P} and the corresponding set of these atoms is $\text{undef}(\mathcal{P})$.

2.1 Three-Valued Łukasiewicz Semantics

Table 1 shows the three-valued Łukasiewicz (1920) semantics. Interpretations are represented by pairs $\langle I^\top, I^\perp \rangle$, where

$$\begin{aligned} I^\top &= \{A \in \mathcal{B}_{\mathcal{P}} \mid A \text{ is mapped to } \top\}, \\ I^\perp &= \{A \in \mathcal{B}_{\mathcal{P}} \mid A \text{ is mapped to } \perp\}. \end{aligned}$$

$\mathcal{B}_{\mathcal{P}}$ is the Herbrand base wrt \mathcal{P} . A model of \mathcal{P} is an interpretation which maps each clause occurring in \mathcal{P} to \top .

2.2 Weak Completion Semantics

Weak completion semantics (Hölldobler and Kencana Ramli 2009a) seems to adequately model Byrne's (1989) suppression task (Dietz, Hölldobler, and Ragni 2012) and Wason's (1968) selection task (Dietz, Hölldobler, and Ragni 2013). Consider the following transformation for \mathcal{P} :

1. Replace all clauses with the same head $A \leftarrow \text{body}_1, \dots, A \leftarrow \text{body}_n$ by $A \leftarrow \text{body}_1 \vee \dots \vee \text{body}_n$.
2. For all atoms A , if $A \in \text{undef}(\mathcal{P})$ then add $A \leftarrow \perp$.
3. Replace all occurrences of \leftarrow by \leftrightarrow .

$F \mid \bar{F}$	$\wedge \mid \top \quad \text{U} \quad \perp$	$\vee \mid \top \quad \text{U} \quad \perp$	$\leftarrow_L \mid \top \quad \text{U} \quad \perp$	$\leftrightarrow_L \mid \top \quad \text{U} \quad \perp$
$\top \mid \perp$	$\top \mid \top \quad \text{U} \quad \perp$	$\top \mid \top \quad \top \quad \top$	$\top \mid \top \quad \top \quad \top$	$\top \mid \top \quad \text{U} \quad \perp$
$\perp \mid \top$	$\text{U} \mid \text{U} \quad \text{U} \quad \perp$	$\text{U} \mid \top \quad \text{U} \quad \text{U}$	$\text{U} \mid \text{U} \quad \top \quad \top$	$\text{U} \mid \text{U} \quad \top \quad \text{U}$
$\text{U} \mid \text{U}$	$\perp \mid \perp \quad \perp \quad \perp$	$\perp \mid \top \quad \text{U} \quad \perp$	$\perp \mid \perp \quad \text{U} \quad \top$	$\perp \mid \perp \quad \text{U} \quad \top$

Table 1: \top , \perp , and U denote *true*, *false*, and *unknown*, respectively.

The resulting set of equivalences is called the *completion* of \mathcal{P} (Clark 1978). If Step 2 is omitted, then the resulting set is called the *weak completion* of \mathcal{P} ($\text{wc } \mathcal{P}$).

Hölldobler and Kencana Ramli (2009b) showed that each weakly completed program admits a least model. Stenning and van Lambalgen (2008) devised an operator which has been generalized for first-order programs by Hölldobler and Kencana Ramli (2009a): Let I be an interpretation in $\Phi_{\mathcal{P}}(I) = \langle J^{\top}, J^{\perp} \rangle$, where

$$J^{\top} = \{A \mid \text{there exists } A \leftarrow \text{body} \in \mathcal{P} \text{ with } I(\text{body}) = \top\},$$

$$J^{\perp} = \{A \mid \text{there exists } A \leftarrow \text{body} \in \mathcal{P} \text{ and for all } A \leftarrow \text{body} \in \mathcal{P} \text{ we find } I(\text{body}) = \perp\}.$$

As shown by Hölldobler and Kencana Ramli, the least fixed point of $\Phi_{\mathcal{P}}$ is identical to the least model of the weak completion of \mathcal{P} ($\text{lm}_L \text{wc } \mathcal{P}$). Starting with the empty interpretation $I = \langle \emptyset, \emptyset \rangle$, $\text{lm}_L \text{wc } \mathcal{P}$ can be computed by iterating $\Phi_{\mathcal{P}}$.¹

3 Reasoning in an Appropriate Logical Form

We follow the approach of Stenning and van Lambalgen (2005; 2008), who proposed to model human reasoning by a two step process: Firstly, human reasoning is to be modeled by setting up an *appropriate representation* and, secondly, *reasoning* wrt this representation.

3.1 Integrity Constraints

One should observe that the first premise *no addictive things are inexpensive* is equivalent to *if something is inexpensive, then it is not addictive*. As weak completion semantics does not allow negative heads in clauses, for every negative conclusion $\overline{p(X)}$ we need to introduce an auxiliary formula $p'(X)$ together with the clause $p(X) \leftarrow p'(X)$. We obtain a (preliminary) representation of the first premise of S_{add} :

$$\begin{aligned} add'(X) &\leftarrow \overline{inex(X)}, \\ add(X) &\leftarrow add'(X), \end{aligned}$$

where $add(X)$, $add'(X)$, and $inex(X)$ denote that X is addictive, not addictive, and inexpensive, respectively.

With the introduction of these auxiliary atoms, the need for integrity constraints arises. An interpretation which maps both, $add(a)$ and $add'(a)$, to true for some constant a should not be a model. This can be represented by a specific type of *integrity constraint* (IC), a logical formula of the form $\perp \leftarrow IC_body$, where IC_body is a conjunction of atoms. For our running example we obtain the integrity constraint:

$$\perp \leftarrow add(X) \wedge add'(X)'.$$

However, the Φ operator does not consider clauses of this form. We need to apply a two step approach: firstly, compute the least model and secondly, verify whether it respects

¹Dietz, Hölldobler, and Wernhard (2014) showed that the weak completion semantics corresponds to the well-founded semantics (Van Gelder, Ross, and Schlipf 1991) for tight logic programs.

the IC , where an interpretation I respects IC iff for each $\perp \leftarrow IC_body$ in IC , we find that $I(IC_body) \in \{\perp, \text{U}\}$.²

For the following examples, whenever there is a $p(X)$ and its $p(X)'$ counterpart, we implicitly assume its corresponding integrity constraint, $IC_p: \perp \leftarrow p(X) \wedge p(X)'$.

3.2 Abnormalities

Stenning and van Lambalgen proposed the idea to implement conditionals by default licenses for implications. This can be achieved by adding an *abnormality predicate* to the antecedent of the implication and initially assuming that the abnormality predicate is false. Following this idea, the initial program representing PREMISE1 in S_{add} is extended to:

*If something is inexpensive and not abnormal,
then it is not addictive.*

Nothing (by default) is abnormal (regarding PREMISE1).

3.3 Background Knowledge and Belief-Bias

A direct first-order representation of PREMISE2 is

There exists a cigarette which is inexpensive. (1)

which includes two pieces of information. Firstly, a more specific one, namely that there exists something, say a , which is a cigarette. Secondly, as Evans et al. have discussed, it contains background knowledge or belief that humans seem to have, which, in this context, we assume to be

Cigarettes are inexpensive (2)

compared to other addictive things, which implies (1) and biases the reasoning towards a representation. This belief bias together with the idea to represent conditionals by default licenses for implication leads to

*If something is a cigarette and not abnormal,
then it is inexpensive.*

Nothing (by default) is abnormal (regarding PREMISE2).

Additionally, we have a second piece of background knowledge, viz. it is commonly known that

Cigarettes are addictive. (3)

This belief together with (2), leads to

*If something is a cigarette,
then it is abnormal (regarding PREMISE1).*

The information contained in the two premises of S_{add} is encoded as the program \mathcal{P}_{add} as follows:

$$\begin{aligned} cig(a) &\leftarrow \top, \\ add'(X) &\leftarrow \overline{inex(X)} \wedge \overline{ab_1(X)}, \\ add(X) &\leftarrow add'(X), \\ ab_1(X) &\leftarrow \perp, \\ ab_1(X) &\leftarrow cig(X), \\ inex(X) &\leftarrow cig(X) \wedge \overline{ab_2(X)}, \\ ab_2(X) &\leftarrow \perp. \end{aligned}$$

²This corresponds to the definition applied for the well-founded semantics in (Pereira, Aparício, and Alferes 1991).

Finally, we impose $IC_{add} : \perp \leftarrow add(X) \wedge add'(X)$.

4 Abductive Framework

Following Kakas, Kowalski, and Toni (1993) we consider an *abductive framework* consisting of a program \mathcal{P} as knowledge base, a collection of atoms \mathcal{A} of abducibles syntactically represented by the set of the (positive and negative) facts for each undefined ground atom in \mathcal{P} , a set of integrity constraints IC , and the logical consequence relation $\models_{\mathcal{L}}^{lmwc}$, where $\mathcal{P} \models_{\mathcal{L}}^{lmwc} F$ if and only if $\text{lm}_{\mathcal{L}}\text{wc}(\mathcal{P}(F)) = \top$ for a formula F . An *observation* is a non-empty set of literals. An explanation \mathcal{E} is a consistent subset of \mathcal{A} .

Definition 1 Let $\langle \mathcal{P}, \mathcal{A}, IC, \models_{\mathcal{L}}^{lmwc} \rangle$ be an abductive framework, \mathcal{O} an observation, \mathcal{E} an explanation, IC a set of integrity constraints, and $\models_{\mathcal{L}}^{lmwc}$ the consequence relation defined for all formulas F .

\mathcal{O} is **explained by \mathcal{E} given \mathcal{P} and IC** iff $\mathcal{P} \cup \mathcal{E} \models_{\mathcal{L}}^{lmwc} \mathcal{O}$, where $\mathcal{P} \not\models_{\mathcal{L}}^{lmwc} \mathcal{O}$ and $\text{lm}_{\mathcal{L}}\text{wc}(\mathcal{P} \cup \mathcal{E})$ respects IC .

\mathcal{O} is **explained given \mathcal{P} and IC** iff there exists an \mathcal{E} such that \mathcal{O} is explained by \mathcal{E} given \mathcal{P} and IC .

We distinguish between two forms of reasoning:

F **follows skeptically from \mathcal{P} , IC and \mathcal{O}** iff \mathcal{O} can be explained given \mathcal{P} and IC , and for all minimal (or otherwise preferred) explanations \mathcal{E} we find that $\mathcal{P} \cup \mathcal{E} \models_{\mathcal{L}}^{lmwc} \mathcal{O}$.

F **follows credulously from \mathcal{P} , IC and \mathcal{O}** iff there exists a minimal (or otherwise preferred) explanation \mathcal{E} such that $\mathcal{P} \cup \mathcal{E} \models_{\mathcal{L}}^{lmwc} \mathcal{O}$.

An important extension of abduction pertains to the issue that, whenever discovering explanations for some given primary observation, one may wish to check too whether some other given additional secondary observations are true, being a logical consequence of the abductive explanations found for the primary observation.³

5 Contextual Abductive Explanations

Let's consider again \mathcal{P}_{add} , where its weak completion consists of the equivalences:

$$\begin{aligned} cig(a) &\leftrightarrow \top, \\ add'(a) &\leftrightarrow \overline{inex(a) \wedge ab_1(a)}, \\ add(a) &\leftrightarrow \overline{add'(a)}, \\ ab_1(a) &\leftrightarrow cig(a), \\ inex(a) &\leftrightarrow \overline{cig(a) \wedge ab_2(a)}, \\ ab_2(a) &\leftrightarrow \perp. \end{aligned}$$

Its least model, $\text{lm}_{\mathcal{L}}\text{wc} \mathcal{P}_{add}$, is

$\{\{cig(a), inex(a), add(a), ab_1(a)\}, \{add'(a), ab_2(a)\}\}$, from which we cannot derive the CONCLUSION in S_{add} . Obviously, the CONCLUSION is something about an object which is not a . The first part of this conclusion is an observation, let's say about b : $\mathcal{O}_{add(b)} = \{add(b)\}$, which we need to explain as described in Section 4. The set of abducibles is: $\mathcal{A}_{(\mathcal{P}_{add} \cup \mathcal{O}_{add(b)})} = \{cig(b) \leftarrow \top, cig(b) \leftarrow \perp\}$.

We have the following minimal explanations for $\mathcal{O}_{add(b)}$:

$$\mathcal{E}_{\overline{cig(b)}} = \{cig(b) \leftarrow \perp\} \text{ and } \mathcal{E}_{cig(b)} = \{cig(b) \leftarrow \top\}.$$

³Reuse of contextual abductions, by resorting to an implementation of tabled abduction for well-founded semantics, is reported in (Saptawijaya and Pereira 2013).

The corresponding least models of the weak completion are: $\text{lm}_{\mathcal{L}}\text{wc}(\mathcal{P}_{add} \cup \mathcal{E}_{\overline{cig(b)}}) = \{\{add(b)\}, \{cig(b), inex(b), \dots\}\}$, $\text{lm}_{\mathcal{L}}\text{wc}(\mathcal{P}_{add} \cup \mathcal{E}_{cig(b)}) = \{\{add(b), cig(b), inex(b), \dots\}, \dots\}$. Under credulous reasoning we conclude, given explanation $\mathcal{E}_{\overline{cig(b)}}$ that the CONCLUSION of S_{add} is true, as there exists something additive which is not a cigarette. However, $\mathcal{O}_{add(b)}$ can also be explained by $\mathcal{E}_{cig(b)}$. We would prefer to reflect that the first explanation describes the usual case and the second describes the exceptional one.

6 Inspection Points

In the example we can easily identify from the context what is the exceptional case. But this is not stated in our program yet. For this purpose we apply inspection points, originally presented in (Pereira and Pinto 2011). We now discuss this approach and show how inspection points can be modeled in our example, concerning its abducibles. Given an atom A , we introduce the following two reserved (meta-)predicates:

$$\text{inspect}(A) \quad \text{and} \quad \text{inspect}_{\neg}(A),$$

that are special cases of abducibles. They differ from usual abducibles in the way that they can only be abduced whenever A or $\neg A$ have been abduced somewhere else already: an abductive solution or explanation \mathcal{E} is only valid when for each $\text{inspect}(A)$ it contains, respectively $\text{inspect}_{\neg}(A)$, it also contains a corresponding A , respectively $\neg A$.

One should observe that for a treatment of inspection points for all literals in a program and not just the abducible ones, we would simply need to adopt the program transformation technique in (Pereira and Pinto 2011), which recursively relays inspection of non-terminal literals to the base inspection of terminals.

6.1 Usual Contextual Abduction

Let us modify our previous example with inspection points.

The new program, \mathcal{P}_{add}^{insp} , has a new ab_1 -clause:

$$(\mathcal{P}_{add} \setminus \{ab_1(X) \leftarrow cig(X)\}) \cup \{ab_1(X) \leftarrow \text{inspect}(cig(X))\}.$$

Suppose again that b is additive, i.e. $\mathcal{O}_{add(b)} = \{add(b)\}$. As $cig(b)$ is unknown, $\text{inspect}(cig(b))$ is false, and $ab_1(b)$ is false rather than unknown that is, its falsity is obtained because nothing is known about $cig(b)$, and there exists no observation to be explained and which abduces $cig(b)$. The only minimal explanation for $\mathcal{O}_{add(b)}$ is generated by $inex(b)$ being false, achieved by $\text{lm}_{\mathcal{L}}\text{wc}(\mathcal{P}_{add}^{insp} \cup \mathcal{E}_{\overline{cig(b)}})$: $\{\{add(b)\}, \{add'(b), cig(b), ab_1(b), inex(b), ab_2(b)\}\}$. Now, even under skeptical reasoning, we conclude that there exists an additive thing which is not a cigarette.

6.2 Contextual Side-effects

Let us extend S_{add} with the conditional:

If something is sold in the streets and there is no reason to believe it is cigarettes, then it is illegal.

Accordingly, our new program \mathcal{P}_{ill}^{insp} is defined as:

$$\mathcal{P}_{add}^{insp} \cup \{\text{ill}(X) \leftarrow \text{streets}(X) \wedge \text{inspect}_{\neg}(cig(X))\}.$$

We only conclude that something is illegal when we have abduced somewhere else that it is not a cigarette. Note, that we can abduce $\mathcal{E}_{\text{streets}(b)} = \{\text{streets}(b) \leftarrow \top\}$ to

partly explain $\mathcal{O}_{ill(b)} = \{illegal(b)\}$. Assume additionally, we observe $\mathcal{O}_{add(b)}$ which is explained by $\mathcal{E}_{cig(b)}$. The $lm_{wc}(\mathcal{P}_{ill}^{insp} \cup \mathcal{E}_{cig(b)} \cup \mathcal{E}_{streets(b)})$ is:

$$\{\{add(b), streets(b), illegal(b), inspect_{-}(cig(b))\}, \\ \{add'(b), cig(b), ab_1(b), inex(b), ab_2(b)\}\}$$

thus, $illegal(b)$ is true. We define side-effects as follows:

Definition 2 Given a background knowledge \mathcal{P} and two observations \mathcal{O}_1 and \mathcal{O}_2 , both explained given \mathcal{P} and IC .

\mathcal{O}_2 is a **necessary contextual side-effect** of \mathcal{O}_1 given \mathcal{P} and IC iff for all explanations \mathcal{E}_1 for \mathcal{O}_1 , there exists some explanation \mathcal{E}_2 for \mathcal{O}_2 containing abducibles of the form $inspect(A)$, or respectively $inspect_{-}(A)$, and for all such abducibles for which abducible A , respectively $\neg A$, is not in \mathcal{E}_2 , and at least one exists, then it is in \mathcal{E}_1 .

\mathcal{O}_2 is a **strict necessary contextual side-effect** of \mathcal{O}_1 given \mathcal{P} and IC iff \mathcal{O}_2 is a necessary contextual side-effect of \mathcal{O}_1 given \mathcal{P} and IC and all explanations \mathcal{E}_1 for \mathcal{O}_1 are also all the explanations for \mathcal{O}_2 .

\mathcal{O}_2 is a **possible contextual side-effect of explained** \mathcal{O}_1 given \mathcal{P} and IC iff there exists an explanation \mathcal{E}_1 for \mathcal{O}_1 and an explanation \mathcal{E}_2 for \mathcal{O}_2 containing abducibles of the form $inspect(A)$, or respectively $inspect_{-}(A)$, and for all such abducibles for which abducible A , respectively $\neg A$, is not in \mathcal{E}_2 , at least one exists, then it is in \mathcal{E}_1 .

\mathcal{O}_2 is a **strict possible contextual side-effect** of \mathcal{O}_1 given \mathcal{P} and IC iff \mathcal{O}_2 is a possible contextual side-effect of \mathcal{O}_1 given \mathcal{P} and IC and there exists an explanation \mathcal{E}_1 for \mathcal{O}_1 such that \mathcal{E}_1 is also an explanation for \mathcal{O}_2 .

The idea behind necessary contextual side-effects, is, that every explanation \mathcal{E}_1 for \mathcal{O}_1 affords us with one *complete* explanation under which some *incomplete* explanation \mathcal{E}_2 for \mathcal{O}_2 is necessarily completed. The idea behind possible contextual side-effects, is, that there is at least one explanation \mathcal{E}_1 for \mathcal{O}_1 which affords us with one *complete* explanation under which some *incomplete* explanation \mathcal{E}_2 for \mathcal{O}_2 can be completed. Summing up, $\mathcal{O}_{add(b)}$ is a strict necessary contextual side-effect of explained $\mathcal{O}_{inex(b)}$ given \mathcal{P}_{add}^{insp} and IC_{add} and $\mathcal{O}_{ill(b)}$ is a possible contextual side-effect of $\mathcal{O}_{add(b)}$ given \mathcal{P}_{ill}^{insp} and IC_{add} but not vice versa.

7 Conclusion

Taking weak completion semantics as starting point, we showed with a running example the need for possible extensions in abductive reasoning. Introducing the concept of inspection points in our framework by applying reserved (meta-)predication for all abductive ground atoms, gives us the possibility to distinguish between different kinds of abducibles in a logic program. This distinction allows to straightforwardly implement the concepts of contextual side-effects. Torasso et al. (1991) do abduction with the completion semantics and do not address side-effects.

8 Acknowledgments

We thank Pierangelo DellAcqua for comments and are grateful to the 3 reviewers. The stay of the 2nd author at UNL was supported by DAAD within the FITweltweit program.

References

- Byrne, R. M. J. 1989. Suppressing valid inferences with conditionals. *Cognition* 31:61–83.
- Clark, K. L. 1978. Negation as failure. In Gallaire, H., and Minker, J., eds., *Logic and Data Bases*, volume 1. New York, NY: Plenum Press. 293–322.
- Dietz, E.-A.; Hölldobler, S.; and Ragni, M. 2012. A computational logic approach to the suppression task. In Miyake, N.; Peebles, D.; and Cooper, R. P., eds., *Proceedings of the 34th Annual Conference of the Cognitive Science Society*, 1500–1505. Austin, TX: Cognitive Science Society.
- Dietz, E.-A.; Hölldobler, S.; and Ragni, M. 2013. A computational logic approach to the abstract and the social case of the selection task. In *Proceedings of the 11th International Symposium on Logical Formalizations of Commonsense Reasoning*.
- Dietz, E.-A.; Hölldobler, S.; and Wernhard, C. 2014. Modeling the suppression task under weak completion and well-founded semantics. *J. of Applied Non-Classical Logics*. accepted.
- Evans, J.; Barston, J. L.; and Pollard, P. 1983. On the conflict between logic and belief in syllogistic reasoning. *Memory & Cognition* 11(3):295–306.
- Hölldobler, S., and Kencana Ramli, C. D. 2009a. Logic programs under three-valued Łukasiewicz semantics. In Hill, P. M., and Warren, D. S., eds., *Logic Programming, 25th International Conference, ICLP 2009*, volume 5649 of *LNCS*, 464–478. Heidelberg: Springer.
- Hölldobler, S., and Kencana Ramli, C. D. 2009b. Logics and networks for human reasoning. In Aliippi, C.; Polycarpou, M. M.; Panayiotou, C. G.; and Ellinas, G., eds., *International Conference on Artificial Neural Networks, ICANN 2009, Part II*, volume 5769 of *LNCS*, 85–94. Heidelberg: Springer.
- Kakas, A. C.; Kowalski, R. A.; and Toni, F. 1993. Abductive logic programming. *Journal of Logic and Computation* 2(6):719–770.
- Łukasiewicz, J. 1920. O logice trójwartościowej. *Ruch Filozoficzny* 5:169–171. English translation: On three-valued logic. In: Łukasiewicz J. and Borkowski L. (ed.). (1990). *Selected Works*, Amsterdam: North Holland, pp. 87–88.
- Pereira, L. M., and Pinto, A. M. 2011. Inspecting side-effects of abduction in logic programs. In Balduccini, M., and Son, T. C., eds., *Logic Programming, Knowledge Representation, and Non-monotonic Reasoning: Essays in honour of Michael Gelfond*, volume 6565 of *LNAI*. Springer. 148–163.
- Pereira, L. M.; Aparício, J. N.; and Alferes, J. J. 1991. Hypothetical reasoning with well founded semantics. In Mayoh, B., ed., *Scandinavian Conference on Artificial Intelligence: Proc. of the SCAI'91*. Amsterdam: IOS Press. 289–300.
- Saptawijaya, A., and Pereira, L. M. 2013. Tabled abduction in logic programs. *Theory and Practice of Logic Programming* 13(4-5-Online-Supplement).
- Stenning, K., and van Lambalgen, M. 2005. Semantic interpretation as computation in nonmonotonic logic: The real meaning of the suppression task. *Cognitive Science* 6(29):916–960.
- Stenning, K., and van Lambalgen, M. 2008. *Human Reasoning and Cognitive Science*. A Bradford Book. Cambridge, MA: MIT Press.
- Torasso, P.; Console, L.; Portinale, L.; and Dupré, D. T. 1991. On the relationship between abduction and deduction. *Journal of Logic and Computation* 1(5):661–690.
- Van Gelder, A.; Ross, K. A.; and Schlipf, J. S. 1991. The well-founded semantics for general logic programs. *Journal of the ACM* 38(3):619–649.
- Wason, P. 1968. Reasoning about a rule. *Quarterly Journal of Experimental Psychology* 20(3):273–281.