Modeling Human Reasoning About Conditionals

Marcos Cramer[1] , **Steffen Hölldobler**[1,2] , **Marco Ragni**[3*]
[1]Technische Universität Dresden, Dresden, Germany
[2]North Caucasus Federal University, Stavropol, Russian Federation
[3]Danish Institute of Advanced Studies, South Denmark University
marcos.cramer@tu-dresden.de, sh43@posteo.de, ragni@sdu.dk

Abstract

Numerous results in psychology demonstrate that inferences humans draw from conditional statements (i.e., statements of the form if antecedent then consequent) differ systematically from classical two-valued logical inferences. Today, still no formal approach yet exists which captures the specifics of semantic differences between types of conditionals (obligation vs. factual) or types of antecedents (necessary vs. non-necessary). We claim that the three-valued, non-monotonic weak completion semantics can model human conditional reasoning adequately, especially with this distinction. We test the predictions of the weak completion semantics in a psychological experiment and demonstrate its cognitive adequacy. We situate the results within formal and cognitive theories and argue that we need logics that are descriptive for the human inference process.

1 Introduction

To demonstrate some specifics of human reasoning, we consider four examples: What follows in each of the following reasoning problems?

1. *If it rains, then the roofs must be wet* and *it rains* (AA).
2. *If Pauls rides a motorbike, then Paul must wear a helmet* and *Paul does not ride a motorbike* (DA).
3. *If the library is open, then Elisa is studying late in the library* and *Elisa is studying late in the library* (AC).
4. *If Nancy rides her motorbike, then Nancy goes to the mountains* and *Nancy does not go to the mountains* (DC).

In each example, a conditional is given together with a positive or negative fact, which is the affirmation of the antecedent (AA), the denial of the antecedent (DA), the affirmation of the consequent (AC), or the denial of the consequent (DC). The examples are adapted from the literature (Dietz Saldanha, Hölldobler, and Lourêdo Rocha 2017; Byrne 2005; Byrne 1989).

We claim that most humans answer *the roofs are wet*, *Paul does not wear a helmet*, *the library is open*, and *Nancy does not ride her motorbike*, respectively, if they have not been exposed to logic before. For the Examples 1 and 4 the answers can be obtained by applying modus ponens and modus tollens, respectively; two valid inference rules in classical two-valued logic. However, for Examples 2 and 3 the answers are invalid in classical two-valued logic.

Such a logic does not seem to be of great help when modeling human conditional reasoning as long as conditionals are represented by implications. Moreover, as Byrne has shown in (Byrne 1989) for each of the four types of inference, humans may suppress previously drawn conclusions when additional knowledge becomes available; this holds for valid as well as invalid inferences with respect to classical two-valued logic. Hence, this calls for a theory based on non-monotonic logic. The well established *mental model theory* (Johnson-Laird and Byrne 1991; Khemlani, Byrne, and Johnson-Laird 2018) claims that conditionals trigger the representation of sets of possibilities. The respective possibilities can be modulated by a reasoner's knowledge, or pragmatics, or semantics leading to different representations (Johnson-Laird and Byrne 2002). Barrouillet et al. demonstrated in (Barrouillet, Grosset, and Lecas 2000) that there is an implicit order on these possibilities in conditional reasoning. A default representation (not considering these modulations above) correctly predicts the answers in the cases AA and DC, but in the cases DA and AC it predicts that humans will answer *nothing follows*. It is well-known that humans sometimes consider conditionals as bi-conditionals (see e.g. (Johnson-Laird and Byrne 1991)), but it is surprising that this seems to hold for all four examples if our earlier claim is correct.

Returning to human conditional reasoning, the main question tackled in this paper is *how can human conditional reasoning be adequately modeled*? Following Bibel (Bibel 1991) we believe that *there is an adequate general proof method that can automatically discover any proof done by humans provided the problem (including all required knowledge) is stated in appropriately formalized terms* where adequateness, roughly speaking, is understood as the property of a theorem proving method that *for any given knowledge base, the method solves simpler problems faster than more difficult ones*.

In this paper we will show that the *weak completion semantics (WCS)*, a three-valued, and non-monotonic cognitive theory, can adequately model human conditional reasoning. In particular, it can adequately model the four examples discussed above. Moreover, it can also explain the differences humans seem to make in the cases AC and DC when

[*]Authors are given in alphabetical order.

dealing with conditionals classified as obligation or factual and antecedents classified as necessary or non-necessary. In the case of AC given a non-necessary antecedent, humans answer with *nothing follows* significantly more often. Whereas in the case of DC given a factual conditional, they answer with *nothing follows* much more often.

In order to validate the claims made above as well as the predictions made by the WCS we designed and performed an experiment involving 56 logically naive participants from Central Europe and Great Britain. The results confirm the claims made above as well as (most of) the predictions made by the WCS. But the results also point towards open research questions.

The paper is organized as follows. After presenting the WCS in Section 2, we introduce a classification of conditionals in Section 3. Taking this classification into account, we extend the WCS. As shown in Section 4, this will lead to a number of predictions made by the WCS. These predictions including the claims made at the beginning of this paper are tested in an experiment specified in Section 5. The experiment will be evaluated in Section 6. A discussion and an outlook to future work concludes the paper in Section 7.

2 The Weak Completion Semantics

We assume the reader to be familiar with logic and logic programming as presented in e.g. (Fitting 1996) and (Lloyd 1984). Let $\top$, $\bot$, and $\cup$ be truth constants denoting *true*, *false*, and *unknown*, respectively. A *(logic) program* is a finite set of clauses of the form $B \leftarrow body$, where B is an atom and *body* is either $\top$, or $\bot$, or a finite, non-empty set of literals. Clauses of the form $B \leftarrow \top$, $B \leftarrow \bot$, and $B \leftarrow L_1, \ldots, L_n$ are called *facts*, *assumptions*, and *rules*, respectively, where L_i, $1 \leq i \leq n$, are literals.

Throughout this paper, $\mathcal{P}$ will denote a program. An atom B is *defined* in $\mathcal{P}$ iff $\mathcal{P}$ contains a clause of the form $B \leftarrow body$. We restrict our attention to propositional programs although the WCS extends to first-order programs as well (Hölldobler 2015). As an example consider the program

$$\mathcal{P}_c = \{C \leftarrow A \wedge \neg ab,\ ab \leftarrow \bot\},$$

where A, C, and ab are atoms. C and ab are defined, whereas A is undefined. ab is an abnormality predicate which is assumed to be false. In the WCS, this program represents the conditional *if A then C*.

Consider the following transformation: (1) For all defined atoms B occurring in $\mathcal{P}$, replace all clauses of the form $B \leftarrow body_1$, $B \leftarrow body_2$, ... by $B \leftarrow body_1 \vee body_2 \vee \ldots$. (2) Replace all occurrences of $\leftarrow$ by $\leftrightarrow$. The resulting set of equivalences is called the *weak completion* of $\mathcal{P}$. It differs from the completion defined in (Clark 1978) in that undefined atoms are not mapped to false, but to unknown instead.

As shown in (Hölldobler and Kencana Ramli 2009a), each weakly completed program admits a least model under the three-valued Łukasiewicz logic (Łukasiewicz 1920) (see Table 1). This model will be denoted by $\mathcal{M}_\mathcal{P}$. It can be computed as the least fixed point of a semantic operator introduced in (Stenning and van Lambalgen 2008). Let $\mathcal{P}$ be a program and I a three-valued interpretation represented by the pair $\langle I^\top, I^\bot \rangle$, where $I^\top$ and $I^\bot$ are the sets of atoms mapped to true and false by I, respectively, and atoms which are not listed are mapped to unknown by I. We define $\Phi_\mathcal{P}\, I = \langle J^\top, J^\bot \rangle$,[1] where

$$\begin{aligned} J^\top &= \{B \mid \text{there is } B \leftarrow body \in \mathcal{P} \text{ and } I\ body = \top\}, \\ J^\bot &= \{B \mid \text{there is } B \leftarrow body \in \mathcal{P} \text{ and} \\ &\qquad \text{for all } B \leftarrow body \in \mathcal{P} \text{ we find } I\ body = \bot\}. \end{aligned}$$

Following (Kakas, Kowalski, and Toni 1992) we consider an *abductive framework* $\langle \mathcal{P}, \mathcal{A}_\mathcal{P}, \mathcal{IC}, \models_{wcs} \rangle$, where $\mathcal{P}$ is a logic program, $\mathcal{A}_\mathcal{P} = \{B \leftarrow \top \mid B \text{ is undefined in } \mathcal{P}\} \cup \{B \leftarrow \bot \mid B \text{ is undefined in } \mathcal{P}\}$ is the *set of abducibles*, $\mathcal{IC}$ is a finite set of *integrity constraints*,[2] and $\mathcal{M}_\mathcal{P} \models_{wcs} F$ iff $\mathcal{M}_\mathcal{P}$ maps the formula F to true. Let $\mathcal{O}$ be an *observation*, i.e., a finite set of literals. $\mathcal{O}$ is *explainable* in the abductive framework $\langle \mathcal{P}, \mathcal{A}_\mathcal{P}, \mathcal{IC}, \models_{wcs} \rangle$ iff there exists a non-empty $\mathcal{X} \subseteq \mathcal{A}_\mathcal{P}$ called an *explanation* such that (1) $\mathcal{M}_{\mathcal{P}\cup\mathcal{X}} \models_{wcs} L$ for all $L \in \mathcal{O}$ and (2) $\mathcal{M}_{\mathcal{P}\cup\mathcal{X}}$ satisfies $\mathcal{IC}$. Formula F *follows credulously* from $\mathcal{P}$ and $\mathcal{O}$ iff there exists an explanation $\mathcal{X}$ for $\mathcal{O}$ such that $\mathcal{M}_{\mathcal{P}\cup\mathcal{X}} \models_{wcs} F$. F *follows skeptically* from $\mathcal{P}$ and $\mathcal{O}$ iff $\mathcal{O}$ can be explained and for all explanations $\mathcal{X}$ for $\mathcal{O}$ we find $\mathcal{M}_{\mathcal{P}\cup\mathcal{X}} \models_{wcs} F$. One should observe that if an observation $\mathcal{O}$ cannot be explained, then *nothing follows* credulously as well as skeptically. In case of skeptical consequences this is an application of the so-called *Gricean implicature* (Grice 1975): humans normally do not quantify over things which do not exist.

Given premises, general knowledge, and observations, *reasoning in the WCS* is hence modeled in five steps:

1. Reasoning towards a program $\mathcal{P}$ following (Stenning and van Lambalgen 2008).
2. Weakly completing the program.
3. Computing the least model $\mathcal{M}_\mathcal{P}$ of the weak completion of $\mathcal{P}$ under the three-valued Łukasiewicz logic.
4. Reasoning with respect to $\mathcal{M}_\mathcal{P}$.
5. If observations cannot be explained, then applying skeptical abduction.

In Section 4 we will explain how these five steps work in the case of the conditional reasoning tasks considered in this paper. More examples can be found, for example, in (Dietz, Hölldobler, and Ragni 2012) or (Oliviera da Costa et al. 2017).

3 A Classification of Conditionals

Obligation versus Factual Conditionals A conditional whose consequent appears to be obligatory given the antecedent is called an *obligation conditional*. As pointed out by Byrne (Byrne 2005), for each obligation conditional there are two initial possibilities people think about. The first possibility is the conjunction of the antecedent and the consequent; it is permitted. The second possibility is the conjunction of the antecedent and the negation of the consequent; it

[1] Whenever we apply a unary operator like $\Phi_\mathcal{P}$ to an argument like I, then we omit parenthesis and write $\Phi_\mathcal{P}\, I$ instead.

[2] In all examples discussed in this paper $\mathcal{IC} = \emptyset$.

F	$\neg F$
$\top$	$\bot$
$\bot$	$\top$
U	U

$\wedge$	$\top$	U	$\bot$
$\top$	$\top$	U	$\bot$
U	U	U	$\bot$
$\bot$	$\bot$	$\bot$	$\bot$

$\vee$	$\top$	U	$\bot$
$\top$	$\top$	$\top$	$\top$
U	$\top$	U	U
$\bot$	$\top$	U	$\bot$

$\leftarrow$	$\top$	U	$\bot$
$\top$	$\top$	$\top$	$\top$
U	U	$\top$	$\top$
$\bot$	$\bot$	U	$\top$

$\leftrightarrow$	$\top$	U	$\bot$
$\top$	$\top$	U	$\bot$
U	U	$\top$	U
$\bot$	$\bot$	U	$\top$

Table 1: The truth tables for the Łukasiewicz logic. One should observe that $\mathsf{U} \leftarrow \mathsf{U} = \mathsf{U} \leftrightarrow \mathsf{U} = \top$ as shown in the grey cells.

is forbidden. Reconsidering Example 1, the permitted possibility is *it rains* and *the roofs are wet*, whereas the forbidden possibility is *it rains* and *the roofs are not wet*. Obligations are deontic obligations, i.e. legal, moral, or societal obligations of a person to perform certain actions, or naive physical obligations that cannot be avoided under normal circumstances. The fact that the consequence is obligatory may be explicitly marked with a word like *must*, but this is unnecessary. The exemplary conditionals 1 and 2 presented in the Introduction appear to be obligations.

If the consequent of a conditional is not obligatory, then it is called a *factual conditional*. In particulr, there is no forbidden possibility in such a case. This appears to hold for Examples 3 and 4 given in the Introduction.

Necessary versus Non-Necessary Antecedents The antecedent A of a conditional $if\,A\;then\,C$ is said to be *necessary* if and only if its consequent C cannot be true unless A is true. More precisely, A may be true while C is not, but C cannot be true while A is not. For example, the *library being open* is a necessary antecedent for *studying late in the library*, but visitors of a library can have varying reasons like *reading textbooks* or *having an essay to write* for *studying late in the library*. In the examples presented in the Introduction, it appears that the antecedents of Examples 1 and 3 are necessary, whereas the antecedents of Examples 2 and 4 appear to be non-necessary.

Pragmatics Humans may classify conditionals as obligation or factual and antecedents as necessary or non-necessary. This is an informal and pragmatic classification. It depends on the background knowledge and experience of a person as well as on the context. For example, the conditional *if it is cloudy, then it is raining* discussed in (Khemlani, Byrne, and Johnson-Laird 2018) may be classified as an obligation conditional with necessary antecedent by people living in Java, whereas it may be classified as a factual conditional by people living in Central Europe.

WCS The classification of conditionals can be taken into account by extending the definition of the set of abducibles:

$$\mathcal{A}^e_{\mathcal{P}} = \mathcal{A}_{\mathcal{P}} \cup \mathcal{A}^{nn}_{\mathcal{P}} \cup \mathcal{A}^f_{\mathcal{P}},$$

where $\mathcal{A}_{\mathcal{P}}$ is as defined above,

$$\begin{array}{rcl} \mathcal{A}^{nn}_{\mathcal{P}} & = & \{C \leftarrow \top \mid C \text{ is the head of a rule in } \mathcal{P} \\ & & \quad \text{representing a conditional with} \\ & & \quad \text{non-necessary antecedent}\}, \\ \mathcal{A}^f_{\mathcal{P}} & = & \{ab \leftarrow \top \mid ab \text{ occurs in the body of a rule in } \mathcal{P} \\ & & \quad \text{representing a factual conditional}\}. \end{array}$$

$C \leftarrow A \wedge \neg ab$	A non-necessary	A necessary
Factual	$ab \leftarrow \top,\ C \leftarrow \top$	$ab \leftarrow \top$
Obligation	$C \leftarrow \top$	

Table 2: The additional facts in the set of abducibles for a rule of the form $C \leftarrow A \wedge \neg ab$ representing a conditional $if\,A\;then\,C$.

The set $\mathcal{A}^{nn}_{\mathcal{P}}$ contains facts for the consequents of conditionals with non-necessary antecedent. If an antecedent of a conditional is non-necessary then there may be other unknown reasons for establishing the consequent of the conditional. The set $\mathcal{A}^f_{\mathcal{P}}$ contains facts for the abnormalities occurring in the bodies of the the representation of factual conditionals. The antecedent of a factual conditional may be true, yet the consequent of the conditional may still not hold. Adding a fact for the abnormality predicate occurring in the bodyl will force this abnormality to become true and its negation to become false. Hence, the body of the clause containing the abnormality predicate will be false.[3] Table 2 illustrates the new facts in the set of abducibles.

4 Predictions of WCS for Human Responses

If a conditional premise $if\,A\;then\,C$ is given as the first premise, then according to (Stenning and van Lambalgen 2008) this shall be represented as a *license for inference* by the program $\mathcal{P}_c$ presented in Section 2. It is called a licence as in human reasoning it is usually not the case that all antecedents which are necessary to enforce a conclusion are mentioned. Weakly completing the program we obtain

$$\{C \leftrightarrow A \wedge \neg ab,\ ab \leftrightarrow \bot\}.$$

Computing its least model we obtain $\langle \emptyset, \{ab\} \rangle$. In this model A and C are mapped to unknown, whereas ab is mapped to false. Please note that this model is the least fixed point of the $\Phi_{\mathcal{P}_c}$ operator which can be computed by iterating the operator starting with the empty interpretation $\langle \emptyset, \emptyset \rangle$.

In the following subsections we assume that a conditional $if\,A\;then\,C$ is given as the first premise and consider the four different cases which occur if a second premise is added.

[3]This technique is used in (Dietz, Hölldobler, and Ragni 2012) to represent an enabling relation and model the suppression effect. In particular, *a library not being open* prevents a person from *studying in it*.

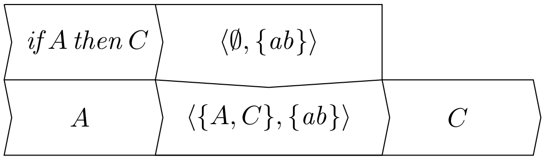

Figure 1: AA reasoning. The left column shows the premises, the middle column the constructed least models, and the right column the generated responses.

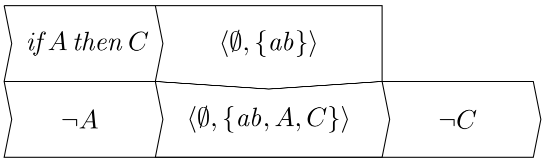

Figure 2: DA reasoning.

4.1 Affirmation of the Antecedent

If the antecedent A of the conditional *if* A *then* C is affirmed as a second premise, then this is represented by the program

$$\mathcal{P}_{aa} = \mathcal{P}_c \cup \{A \leftarrow \top\}.$$

Weakly completing the program and computing its least model we obtain $\langle\{A, C\}, \{ab\}\rangle$. Reasoning with respect to this model we conclude C (see Figure 1). This is independent of the classification of conditionals as obligation or factual or that of antecedent as necessary or non-necessary. Example 1 presented in the Introduction belongs to this category with A and C denoting *it rains* and *the roofs must be wet*, respectively.

Predictions of WCS for Human Responses on AA In AA inferences with the premises *if* A *then* C and A, most humans will answer C, and this is independent of the classification of the conditional and the antecedent.

4.2 Denial of the Antecedent

If the antecedent A of the conditional *if* A *then* C is denied as a second premise, then this is represented by the program

$$\mathcal{P}_{da} = \mathcal{P}_c \cup \{A \leftarrow \bot\}.$$

Weakly completing the program and computing its least model we obtain $\langle\emptyset, \{ab, A, C\}\rangle$. Reasoning with respect to this model we conclude $\neg C$ (see Figure 2). This is independent of the classification of the conditional as well as the antecedent. Example 2 presented in the Introduction belongs to this category with A and C denoting *Paul rides a motorbike* and *Paul is wearing a helmet*, respectively.

Prediction of WCS for Human Responses on DA In DA inferences with the premises *if* A *then* C and $\neg A$, most humans will answer $\neg C$, and this is independent of the classification of the conditional and the antecedent.

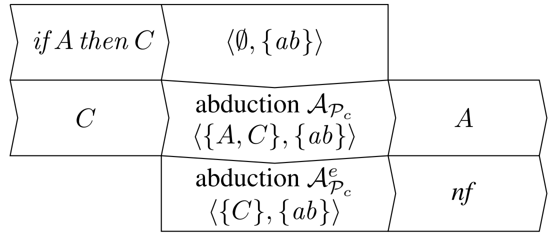

Figure 3: AC reasoning. The answer *nothing follows (nf)* is given if the antecedent of the conditional is non-necessary, the reasoner considers $\mathcal{A}^e_{\mathcal{P}_c}$ and is reasoning skeptically.

4.3 Affirmation of the Consequent

If the consequent C of the conditional *if* A *then* C is affirmed as a second premise, then this is considered to be an observation to be explained because C is already defined in the program $\mathcal{P}_c$. But A is undefined. Hence, we obtain $\mathcal{A}_{\mathcal{P}_c} = \{A \leftarrow \top, A \leftarrow \bot\}$. $\{A \leftarrow \top\}$ is the only minimal explanation for $\{C\}$. Let

$$\mathcal{P}_{ac} = \mathcal{P}_c \cup \{A \leftarrow \top\}.$$

Weakly completing the program and computing its least model we obtain $\langle\{A, C\}, \{ab\}\rangle$. Reasoning with respect to this least model we conclude A.

However, if the classification of antecedents is taken into account and if the conditional has a non-necessary antecedent, then the set of abducibles will be extended by the fact $C \leftarrow \top$. In this case, there is a second minimal explanation for $\{C\}$, viz. $\{C \leftarrow \top\}$. Let

$$\mathcal{P}'_{ac} = \mathcal{P}_c \cup \{C \leftarrow \top\}.$$

Weakly completing the program and computing its least model we obtain $\langle\{C\}, \{ab\}\rangle$. Taking both explanations into account and reasoning skeptically, we conclude *nothing follows (nf)* (see Figure 3).[4] The case of a conditional with necessary antecedent is exemplified by Example 3 from the Introduction, with A and C denoting *the library is open* and *Elisa is studying late in the library*, respectively. Here we conclude that *the library is open*.

Let us now consider an everyday conditional with non-necessary antecedent. What follows from

5. *if Paul rides a motorbike, then Paul must wear a helmet* and *Paul wears a helmet*?

We expect that a significant number of humans will answer *nothing follows*.

Prediction of WCS for Human Responses on AC In AC inferences with the premises *if* A *then* C and C, most humans will answer A. If A is a non-necessary antecedent, then the number of *nf* answers will increase. Moreover, the

[4]Formally, C and $\neg ab$ follow skeptically, but this is nothing new as C is the observation and ab is assumed to be false in the weak completion of $\mathcal{P}_c$. We would like to draw conclusions which preserve semantic information, are parsimonious, and state something new (Johnson-Laird and Byrne 1991).

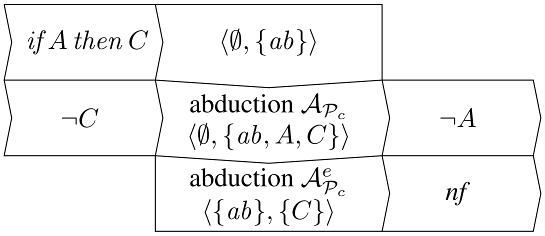

Figure 4: DC reasoning. The answer *nf* is given if the conditional is a factual one, the reasoner considers $\mathcal{A}^e_{\mathcal{P}_c}$ and is reasoning skeptically.

time to generate an *nf* answer will be longer than the time to generate the answer A.

4.4 Denial of the Consequent

If the consequent C of the conditional $\textit{if } A \textit{ then } C$ is denied as a second premise, then this is again considered to be an observation because C is already defined in $\mathcal{P}_c$. But A is undefined. Hence, we obtain $\mathcal{A}_{\mathcal{P}_c} = \{A \leftarrow \top,\ A \leftarrow \bot\}$. $\{A \leftarrow \bot\}$ is the only minimal explanation for $\{\neg C\}$. Let

$$\mathcal{P}_{dc} = \mathcal{P}_c \cup \{A \leftarrow \bot\}.$$

Weakly completing the program and computing its least model we obtain $\langle \emptyset, \{ab, A, C\} \rangle$. Reasoning with respect to this least model we conclude $\neg A$.

However, if the classification of conditionals is taken into account and if the conditional is a factual one, then the set of abducibles will be extended by the fact $ab \leftarrow \top$. In this case, there is a second minimal explanation for $\{\neg C\}$, viz. $\{ab \leftarrow \top\}$. Let

$$\mathcal{P}'_{dc} = \mathcal{P}_c \cup \{ab \leftarrow \top\}.$$

Weakly completing the program and computing its least model we obtain $\langle \{ab\}, \{C\} \rangle$. Taking both explanations into account and reasoning skeptically, we conclude *nf* (see Figure 4). Example 4 presented in the Introduction belongs to this category with A and C denoting *Nancy rides her motorbike* and *Nancy goes to the mountains*, respectively. As this was classified as a factual conditional we expect that a significant number of humans will answer *nothing follows*.

Let us now consider an everyday obligation conditional. What follows from

6. *if Paul rides a motorbike, then Paul must wear a helmet* and *Paul does not wear a helmet*?

We expect that most participants will conclude that *Paul does not ride a motorbike*.

Prediction of WCS for Human Responses on DC In DC inferences with the premises $\textit{if } A \textit{ then } C$ and $\neg C$, most humans will answer $\neg A$. If the conditional is a factual one, then the number of *nf* answers will increase. Moreover, the time to generate an *nf* answer will be longer than the time to generate the answer $\neg A$.

5 Putting it to the Test

The goal of our investigation is to test the predictions made in the previous section in an everyday context, i.e., in a context familiar to the participants.

5.1 Participants, materials and methods

We tested 56 logically naive participants on an online website (Prolific, prolific.co). We restricted the participants to Central Europe and Great Britain to have a similar background knowledge about weather etc. We assume that the participants had not received any education in logic beyond high school training. We took the usual precautions for such a procedure; for example, the website checked that participants were proficient speakers of English. The participants were first presented with a story followed by a first assertion ("a conditional premise"), and a second assertion ("a (possibly negated) atomic premise"), and then for each problem they had to answer the question "What follows?". Both parts were presented simultaneously. The participants responded by clicking one of the answer options. They could take as much time as they needed. Participants acted as their own controls.

The participants carried out 48 problems consisting of the 12 conditionals listed in the Appendix and solved all four inference types (AA, DA, AC, DC). They could select one of three responses: *nothing follows*, the fact that had not been presented in the second premise, and the negation of this fact. We chose the content based on (i) previously tested conditionals in the literature and (ii) on everyday context. The classification of the conditionals was done by the authors.

As an example consider the following story: *Peter has a lawn in front of his house. He is keen to make sure that the grass on lawn does not dry out, so whenever it has been dry for multiple days, he turns on the sprinkler to water the lawn.* Then, the conditional *if it rains, then the lawn is wet* and the negated atomic proposition *the lawn is not wet* are given. In this case, the three answers from which participants could select were *it rains*, *it does not rain*, and *nothing follows*.

6 Evaluation

6.1 Affirmation of the Antecedent

The total number of selected responses as well as the median response time (in milliseconds) for C ($\textit{Mdn } C$) and *nf* ($\textit{Mdn nf}$) responses can be found in Table 3 for AA inferences that is for a given conditional $\textit{if } A \textit{ then } C$ and fact A.

The everyday context elicited a high response rate of AA inferences of about 95% (640 out of 672) for C-answers. The number of participants answering $\neg C$ or *nf* as well as the classification of conditionals appears to be irrelevant. The WCS models human AA inferences adequately.

6.2 Denial of the Antecedent

The total number of selected responses as well as the median response time (in milliseconds) for $\neg C$ ($\textit{Mdn } \neg C$) and *nf* ($\textit{Mdn nf}$) responses can be found in Table 4 for DA inferences that is for a given conditional $\textit{if } A \textit{ then } C$ and fact $\neg A$.

Class	C	$\neg C$	nf	Sum	$Mdn\, C$	$Mdn\, nf$
(1)	55	1	0	56	3343	na
(2)	55	1	0	56	3487	na
(3)	53	3	0	56	3516	na
O+nec	163	5	0	168	3408	na
(4)	53	1	2	56	3403	3472
(5)	53	2	1	56	3903	3572
(6)	54	1	1	56	3088	6959
O+non-nec	160	4	4	168	3543	4183
(7)	49	1	6	56	3885	7051
(8)	54	1	1	56	3559	7349
(9)	54	1	1	56	3710	3826
F+nec	157	3	8	168	3615	6926
(10)	51	2	3	56	3929	6647
(11)	54	1	1	56	3777	5073
(12)	55	1	0	56	2977	na
F+non-nec	160	4	4	168	3644	5860
Obligation	323	9	4	336	3516	4183
Factual	317	7	12	336	3640	6575
Necessary	320	8	8	336	3546	6926
Non-nec	320	8	8	336	3588	4934
Total	640	16	16	672	3570	5925

Table 3: The results for AA inferences. The grey line shows the numbers for Example 1. '*na*' is an acronym for *not applicable*. 'O+nec' refers to obligation conditionals with necessary antecedent, which are the conditionals (1) - (3) in the experiment. 'O+non-nec' refers to obligation conditionals with non-necessary antecendent, which are the conditionals (4) - (6) in the experiment. 'F+nec' refers to factual conditionals with necessary antecendent, which are the conditionals (7) - (9) in the experiment. 'F+non-nec' refers to factual conditionals with non-necessary antecedent, which are the conditionals (10) - (12) in the experiment. In the lines labeled 'Obligation' and 'Factual' the results for obligation and factual conditionals are shown, respectively. In the lines labeled 'Necessary' and 'Non-nec' the results for conditionals with necessary and non-necessary antecedents are shown, respectively. The line labeled 'Total' shows the results for all experiments.

Class	C	$\neg C$	nf	Sum	$Mdn\, \neg C$	$Mdn\, nf$
(1)	0	45	11	56	2863	4901
(2)	2	54	0	56	3367	na
(3)	2	51	3	56	3647	10477
O+nec	4	150	14	168	3356	5115
(4)	1	40	15	56	3722	7189
(5)	3	28	25	56	5735	7814
(6)	4	36	16	56	3602	6240
O+non-nec	8	104	56	168	4064	7471
(7)	2	51	3	56	3928	7273
(8)	1	47	8	56	3296	5728
(9)	1	52	3	56	3549	8735
F+nec	4	150	14	168	3605	6582
(10)	1	39	16	56	3725	6874
(11)	0	41	15	56	3374	5887
(12)	1	41	14	56	3205	7002
F+non-nec	2	121	45	168	3374	6221
Obligation	12	254	70	336	3583	6613
Factual	6	271	59	336	3518	6221
Necessary	8	300	28	336	3474	5808
Non-nec	10	225	101	336	3646	6700
Total	18	525	129	672	3558	6450

Table 4: The results for DA inferences. The grey line shows the numbers for Example 2. If the antecedent is non-necessary, then *nf* is answered significantly more often (grey cells).

The everyday context elicited a high response rate of DA inferences of about 78% (525 out of 672) for $\neg C$-answers, but the case of *nf*-answers varied from 8% (14 out of 168) up to 33% (56 out of 168). The number of participants answering C is irrelevant.

The answer *nf* was more often given in case of conditionals with non-necessary antecedents than in the case of conditionals with necessary antecedents (30% vs. 8%, Wilcoxon signed rank, $W = 0$, $p < .001$). The WCS predicts the answer $\neg C$ given by the majority of the participants, but it cannot model the difference of the *nf*-answers. We speculate that in case of an *nf*-answer the clauses representing conditionals with non-necessary antecedents should not be weakly completed. This would require a modification to the semantic definitions of WCS, whose theoretical and algorithmic properties have not yet been investigated.

6.3 Affirmation of the Consequent

The total number of selected responses as well as the median response time (in milliseconds) for A ($Mdn\, A$) and *nf* ($Mdn\, nf$) responses can be found in Table 5 for DA inferences that is for a given conditional $if\, A\, then\, C$ and fact C.

The everyday context elicited a high response rate of AC

Class	A	$\neg A$	*nf*	Sum	*Mdn A*	*Mdn nf*
(1)	37	1	18	56	3952	7995
(2)	48	1	7	56	4003	4170
(3)	43	1	12	56	3458	9001
O+nec	128	3	37	168	3797	8175
(4)	42	1	13	56	3659	8828
(5)	32	1	23	56	4704	6044
(6)	29	1	26	56	3593	4396
O+non-nec	103	3	62	168	3968	5939
(7)	51	1	4	56	3767	4397
(8)	42	1	13	56	3798	4565
(9)	45	1	10	56	3492	4598
F+nec	138	3	27	168	3699	4565
(10)	34	2	20	56	5224	6289
(11)	29	2	25	56	3218	6205
(12)	33	1	22	56	3483	4992
F+non-nec	96	5	67	168	3885	6116
Obligation	231	6	99	336	3888	6044
Factual	234	8	94	336	3769	5650
Necessary	266	6	64	336	3735	5450
Non-nec	199	8	129	336	3906	6039
Total	465	14	193	672	3826	5802

Table 5: The results for AC inferences. The grey lines show the results for Examples 3 (line marked (7)) and 5 (line marked (4)). If the antecedent is non-necessary, then *nf* is answered significantly more often (grey cells).

Class	A	$\neg A$	*nf*	Sum	*Mdn* $\neg A$	*Mdn nf*
(1)	1	45	10	56	3449	4758
(2)	0	50	6	56	4058	7922
(3)	2	46	8	56	3796	4517
O+nec	3	141	24	168	3767	5732
(4)	3	46	7	56	3872	4154
(5)	1	54	1	56	4946	8020
(6)	0	36	20	56	4062	5235
O+non-nec	4	136	28	168	4293	5803
(7)	1	37	18	56	5974	4744
(8)	3	42	11	56	4367	5013
(9)	0	47	9	56	4208	3966
F+nec	4	126	38	168	4849	4574
(10)	2	35	19	56	4879	4167
(11)	0	39	17	56	4411	5647
(12)	0	34	22	56	3726	3813
F+non-nec	2	108	58	168	4338	4542
Obligation	7	277	52	336	4053	4790
Factual	6	234	96	336	4459	4345
Necessary	7	267	62	336	4096	4758
Non-nec	6	244	86	336	4325	4555
Total	13	511	148	672	4311	5162

Table 6: The results for DC inferences. The grey lines show the results for Examples 4 (line marked (10)) and 6 (line marked (4)). In case of factual conditionals, *nf* is answered significantly more often (grey cells).

inferences of about 69% (465 out of 672) for A-answers, but the case of *nf*-answers varied from 16% (27 out of 168) up to 40% (67 out of 168). The number of participants answering $\neg A$ is irrelevant.

As predicted by the WCS, the answer *nf* was more often given in case of conditionals with non-necessary antecedents than in the case of conditionals with necessary antecedents (38% vs. 19%, Wilcoxon signed rank, $W = 82$, $p < .001$).

6.4 Denial of the Consequent

The total number of selected responses as well as the median response time (in milliseconds) for $\neg A$ (*Mdn* $\neg A$) and *nf* (*Mdn nf*) responses can be found in Table 6 for DC inferences that is for a given conditional *if A then C* and fact $\neg C$.

The everyday context elicited a high response rate of DC inferences of about 76% (511 out of 672) for $\neg A$-answers, but the case of *nf*-answers varied from 14% (24 out of 168) up to 35% (58 out of 168). The number of participants answering A is irrelevant.

As predicted by the WCS, the answer *nf* was more often given in case of a factual conditional than in case of an obligation conditional (35% vs. 14%, Wilcoxon signed rank, $W = 133$, $p < .001$). So the predicted increase in the selection of *nf* can be confirmed.

6.5 Interpreting the Results

For each conditional used in the experiments and for each type of inference, the WCS correctly predicted the answer given by a majority of the participants. This can be explained in classical, two-valued logic if one assumes that each conditional used in the experiments was erroneously considered to be a bi-conditional by the majority. This is quite surprising given that six of the twelve antecedents of the conditionals used in the experiment were classified as non-necessary (see Appendix). Moreover, the WCS correctly predicted the rising number of *nf*-answers in AC inferences if the antecedent was non-necessary and in DC inferences if the conditional was a factual one.

Given an AA inference task, reasoners just conclude the consequent of the conditional. This corresponds to modus ponens. Reasoners are familiar with this kind of inference and make almost no mistakes.

Given a DA inference task, most reasoners conclude the negation of the consequent of the conditional as predicted. One should note that the median response time of the answer C in AA inferences and the median response time of the answer $\neg C$ in DA inferences are almost identical (3570 vs. 3558). This can also be explained by the WCS in that the steps taken to construct the least models in AA and DA inference tasks are very similar. In each case, the semantic operator $\Phi_{\mathcal{P}}$ needs to be applied twice to reach a fixed point. In the connectionist network implementing the semantic operator (Hölldobler and Kencana Ramli 2009b) the stable states corresponding to the least fixed points are computed in six steps in both cases.

Given an AC or DC inference task, reasoners may search for a minimal explanation of $\{C\}$ or $\{\neg C\}$ using the set $\mathcal{A}_{\mathcal{P}}$ of abducibles. Such a minimal explanation always exists and gives rise to a model that maps the antecedent A of the given conditional *if A then C* to either true or false, respectively. This model may be called the *preferred* model in the sense of (Ragni and Knauff 2013). Once the preferred model has been constructed, a reasoner may upon further thought search for models using the extended set $\mathcal{A}^e_{\mathcal{P}} \supseteq \mathcal{A}_{\mathcal{P}}$ of abducibles and find a second minimal explanation, giving rise to a second model. In this second model A is unknown. Reasoning skeptically, the reasoner will answer *nf*. This not only explains the difference between necessary and non-necessary antecedents or obligation and factual conditionals but also why a significantly larger number of participants answered *nf* in the case of non-necessary antecedents and factual conditionals, respectively. In order to fully support his interpretation of the results, further experiments recording the time of deliberation are required.

WCS correctly predicts that the answer *nf* appears significantly less frequently for AC inferences with a necessary antecedent as well as for DC inferences with an obligation conditional. However, even though the answer *nf* does appear significantly less frequently in these cases, the number of *nf* answers for these inference tasks is not as insignificantly small as in the case of AA inferences. This is not predicted by WCS, but may have multiple various reasons: Some reasoners might not consider C or $\neg C$ as an observation that needs to be explained. Rather, if they might just add $C \leftarrow \top$ or $C \leftarrow \bot$ to the program, in which case no model assigning A to true or false can be constructed. Some reasoners might consider C or $\neg C$ as an observation that needs to be explained but not necessarily by A or $\neg A$, respectively; moreover, the classification of the given conditional may depend on the cultural background of the reasoner. Or the reasoners may make a mistake in constructing the preferred model, which is – as mentioned before – the least fixed point of the semantic operator introduced in (Stenning and van Lambalgen 2008).

One should observe that it took participants less time to answer C and $\neg C$ in AA and DA inferences compared to the time to answer A and $\neg A$ in AC and DC inferences. This is a well-known phenomena, as AC and DC inferences are considered to be more difficult than AA and DA ones (see e.g. (Barrouillet, Grosset, and Lecas 2000)). This can also be explained by the WCS: In AA and DA inferences it suffices to compute the least fixed point of the semantic operator, whereas in AC and DC inferences abduction needs to be added considering the consequent or its negation as an observation to be explained. Apart from that, at the moment we can only speculate why the median response time of $\neg A$-answers in DC inferences is larger than the median response time of A-answers in AC inferences: This may depend on the sequences in which possible explanations for the observations are considered; in particular, the possible explanation $\{A \leftarrow \top\}$ may have been considered before the possible explanation $\{A \leftarrow \bot\}$.

The second hypothesis is, however, a core question: is answering *nf* an indication that a participant did not know the answer, or is it at the end of a deliberation process that might follow the predicted process in the WCS? While this answer cannot be given in general, the median response times for *nf* are higher than for the respective responses A, $\neg A$, C, or $\neg C$. This often indicates more thinking and less guessing, because participants do not quickly and easily respond *nf* to avoid thinking. These findings indicate towards the processes predicted for each inference type in the WCS, but for further support more studies are necessary.

7 Summary and Outlook

As shown in this paper, the WCS adequately models human conditional reasoning in that it generates the answers given by a majority of the participants. It is based on several principles: (1) Conditionals are represented as licenses for inference in a (logic) program. (2) Abnormality predicates are used to represent unknown additional conditions; they are initially assumed to be false. (3) The definitions given in a program are weakly completed. (4) Programs are interpreted under the three-valued Łukasiewicz logic. (5) A positive or negative fact given as a premise is considered to be an observation which needs to be explained if the program already contains a definition for the fact. (6) Skeptical abduction is applied. (7) The Gricean implicature is applied.

These principles are well justified. In most human reasoning scenarios not all necessary antecedents of a conditional will or can be given. The abnormality predicate takes care of this. If a detail becomes important later on, then this can be added. Reconsider Example 4, the *absence of gas will prevent Nancy from riding the motorbike to the mountains*. *Sufficient amount of gas* is an enabling relation for riding a motorbike. This can be modeled by adding the rule $ab \leftarrow \neg gas$ to the program representing Example 4. The new rule will override $ab \leftarrow \bot$ when the weak completion is computed as $ab \leftrightarrow \bot \vee \neg gas$ is semantically equivalent to $ab \leftrightarrow \neg gas$. In general, positive information will override a negative one. During the weak completion process, the only-if halves of definitions[5] are added, which is based on ideas underlying *conditional perfection* in linguistics (Van der Auwera 1997).

It has been shown that two-valued logics cannot model human reasoning (Ragni et al. 2016). Furthermore, under the three-valued Łukasiewics logic, programs and their weak completions have least models. This does not hold

[5]The weak completion of the definition $A \leftarrow C \wedge \neg ab$ is $A \leftrightarrow C \wedge \neg ab$, whereas the bi-conditional corresponding to the conditional *if A then C* is *A iff C*.

for the three-valued Kleene logic (Kleene 1952). There, $\cup \leftarrow \cup = \cup$ and, consequently, programs like $\{a \leftarrow b\}$ have two minimal models $\langle\{a, b\}, \emptyset\rangle$ and $\langle\emptyset, \{a, b\}\rangle$, but no least one. Reasoning credulously does not adequately model human reasoning as shown in this paper, because credulous reasoning does not account for the growing number of *nf*-responses in AC and DC inferences if conditionals are classified. The Gricean implicature has been applied at various occasions: an atom can only be false if there is evidence for this falsehood; otherwise it is unknown. This can be seen in the definition of weak completion as well as in the definition of the Φ-operator. Likewise, skeptical consequences are only defined if the observations are explainable.

The WCS constructs models, which are considered to be mental models in the sense of (Craik 1945) and (Johnson-Laird 1983). Reasoning is performed with respect to the constructed mental models. The WCS is non-monotonic, multi-valued, and the background knowledge need not be consistent which might be closer to humans than to formal databases. Furthermore, the semantic operator, which is used to construct the models can be represented as a feed-forward network (Hölldobler and Kencana Ramli 2009b; Dietz Saldanha et al. 2018a), can be learned (d'Avila Garcez and Zaverucha 1999; Besold et al. 2017), and can be applied to model the average human reasoner. Thus, it suggests a solution for the five fundamental problems for logical models of human reasoning discussed in (Oaksford and Chater 2020). Moreover, the WCS is computational, meaning answers to queries are computed. It is also comprehensive in that different human reasoning tasks can be modeled without changing the theory. For example, in (Dietz, Hölldobler, and Ragni 2012) it is shown that the suppression task (Byrne 1989) is modeled adequately by the WCS. In (Oliviera da Costa et al. 2017) it is shown that the WCS models human syllogistic reasoning better than the twelve cognitive theories investigated in (Khemlani and Johnson-Laird 2012) and in (Dietz Saldanha et al. 2018b) it was shown how ethical decision problems can be modeled by the WCS. The WCS as shown here, is a formally founded approach that can explain human reasoning and is a bridging system at the intersection between human and formal reasoning.

However, much remains to be done and we have already raised various open questions in the paper. How can the increased number of *nf*-answers in DA inferences be explained? How can we distinguish between the third truth value 'unknown' and 'I don't have a clue' in the answers of participants? How is WCS related to MMT?

Acknowledgements

We like to thank Meghna Bhadra, Islam Hamada, and Claudia Schon for many suggestions to improve the readability of the paper and Sangeet S. Khemlani for many fruitful discussions on cognitive and formal theories.

Appendix: Conditionals of the Experiment

Obligation Conditionals with Necessary Antecedent (O+nec) (1) *If it rains, then the roofs must be wet.* (2) *If water in the cooking pot is heated over* $99°C$*, then the water starts boiling.* (3) *If the wind is strong enough, then the sand is blowing over the dunes.*

Obligation Conditionals with Non-Necessary Antecedent (O+non-nec) (4) *If Paul rides a motorbike, then Paul must wear a helmet.* (5) *If Maria is drinking alcoholic beverages in a pub, then Maria must be over 19 years of age.* (6) *If it rains, then the lawn must be wet.*

Factual Conditionals with Necessary Antecedent (F+nec) (7) *If the library is open, then Sabrina is studying late in the library.* (8) *If the plants get water, then they will grow.* (9) *If my car's start button is pushed, then the engine will start running.*

Factual Conditionals with Non-Necessary Antecedent (F+non-nec) (10) *If Nancy rides her motorbike, then Nancy goes to the mountains.* (11) *If Lisa plays on the beach, then Lisa will get sunburned.* (12) *If Ron scores a goal, then Ron is happy.*

References

Barrouillet, P.; Grosset, N.; and Lecas, J.-F. 2000. Conditional reasoning by mental models: Chronometric and developmental evidence. *Cognition* 75:237–266.

Besold, T. R.; d'Avila Garcez, A. S.; Bader, S.; Bowman, H.; Domingos, P. M.; Hitzler, P.; Kühnberger, K.-U.; Lamb, L. C.; Lowd, D.; Lima, P. M. V.; de Penning, L.; Pinkas, G.; Poon, H.; and Zaverucha, G. 2017. Neural-symbolic learning and reasoning: A survey and interpretation. *CoRR* abs/1711.03902.

Bibel, W. 1991. Perspectives on automated deduction. In Boyer, R. S., ed., *Automated Reasoning: Essays in Honor of Woody Bledsoe*. Dordrecht: Kluwer Academic. 77–104.

Byrne, R. M. J. 1989. Suppressing valid inferences with conditionals. *Cognition* 31(1):61–83.

Byrne, R. M. J. 2005. *The Rational Imagination: How People Create Alternatives to Reality*. Cambridge, MA, USA: MIT Press.

Clark, K. L. 1978. Negation as failure. In Gallaire, H., and Minker, J., eds., *Logic and Databases*. New York: Plenum. 293–322.

Craik, K. J. W. 1945. *The Nature of Explanation*. Cambridge: Cambridge University Press.

d'Avila Garcez, A. S., and Zaverucha, G. 1999. The connectionist inductive learning and logic programming system. *Applied Intelligence* 11(1):69–77.

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. Cognitive Science Society.

Dietz Saldanha, E.-A.; Hölldobler, S.; Kencana Ramli, C. D. P.; and Palacios Medinacelli, L. 2018a. A core method for the weak completion semantics with skeptical abduction. *Journal of Artificial Intelligence Research* 63:51–86.

Dietz Saldanha, E.-A.; Hölldobler, S.; Schwarz, S.; and Stefanus, L. Y. 2018b. The weak completion semantics and equality. In G. Barthe, G. Sutcliffe, M. V., ed., *LPAR-22. Proceedings of the 22nd International Conference on Logic for Programming, Artificial Intelligence, and Reasoning*, volume 57, 326–242. EPiC series in Computing.

Dietz Saldanha, E.-A.; Hölldobler, S.; and Lourêdo Rocha, I. 2017. Obligation versus factual conditionals under the weak completion semantics. In Hölldobler, S.; Malikov, A.; and Wernhard, C., eds., *Proceedings of the Second Young Scientists' International Workshop on Trends in Information Processing*, volume 1837, 55–64. CEUR-WS.org. http://ceur-ws.org/Vol-1837/.

Fitting, M. 1996. *First–Order Logic and Automated Theorem Proving*. Berlin: Springer-Verlag, 2nd edition.

Grice, H. P. 1975. Logic and conversation. In Cole, P., and Morgan, J. L., eds., *Syntax and Semantics*, volume 3. Academic Press, New York. 41–58.

Hölldobler, S., and Kencana Ramli, C. D. P. 2009a. Logic programs under three-valued Łukasiewicz's semantics. In Hill, P. M., and Warren, D. S., eds., *Logic Programming*, volume 5649 of *Lecture Notes in Computer Science*, 464–478. Springer-Verlag Berlin Heidelberg.

Hölldobler, S., and Kencana Ramli, C. D. P. 2009b. Logics and networks for human reasoning. In Alippi, C.; Polycarpou, M. M.; Panayiotou, C. G.; and Ellinasetal, G., eds., *Artificial Neural Networks – ICANN*, volume 5769 of *Lecture Notes in Computer Science*, 85–94. Springer-Verlag Berlin Heidelberg.

Hölldobler, S. 2015. Weak completion semantics and its applications in human reasoning. In Furbach, U., and Schon, C., eds., *Bridging 2015 – Bridging the Gap between Human and Automated Reasoning*, volume 1412 of *CEUR Workshop Proceedings*, 2–16. CEUR-WS.org. http://ceur-ws.org/Vol-1412/.

Johnson-Laird, P. N., and Byrne, R. M. J. 1991. *Deduction*. Hove and London (UK): Lawrence Erlbaum Associates.

Johnson-Laird, P. N., and Byrne, R. M. J. 2002. Conditionals: A theory of meaning, pragmatics, and inference. *Psychological Review* 109:646–678.

Johnson-Laird, P. N. 1983. *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*. Cambridge: Cambridge University Press.

Kakas, A. C.; Kowalski, R. A.; and Toni, F. 1992. Abductive Logic Programming. *Journal of Logic and Computation* 2(6):719–770.

Khemlani, S., and Johnson-Laird, P. N. 2012. Theories of the syllogism: A meta-analysis. *Psychological Bulletin* 138(3):427–457.

Khemlani, S. S.; Byrne, R. M. J.; and Johnson-Laird, P. N. 2018. Facts and possibilities: A model-based theory of sentenial reaoning. *Cognitive Science* 1–38.

Kleene, S. C. 1952. *Introduction to Metamathematics*. North-Holland.

Lloyd, J. W. 1984. *Foundations of Logic Programming*. Springer-Verlag.

Łukasiewicz, J. 1920. O logice trójwartościowej. *Ruch Filozoficzny* 5:169–171.

Oaksford, M., and Chater, N. 2020. New paradigms in the psychology of reasoning. *Annual Review of Psychology* 71:12.1–12.26.

Oliviera da Costa, A.; Dietz Saldanha, E.-A.; Hölldobler, S.; and Ragni, M. 2017. A computational logic approach to human syllogistic reasoning. In Gunzelmann, G.; Howes, A.; Tenbrink, T.; and Davelaar, E. J., eds., *Proceedings of the 39th Annual Conference of the Cognitive Science Society*, 883–888. Austin, TX: Cognitive Science Society.

Ragni, M., and Knauff, M. 2013. A theory and a computational model of spatial reasoning. *Psychological Review* 120:561–588.

Ragni, M.; Dietz, E.-A.; Kola, I.; and Hölldobler, S. 2016. Two-valued logic is not sufficient to model human reasoning, but three-valued logic is: A formal analysis. In Furbach, U., and Schon, C., eds., *Bridging 2016 – Bridging the Gap between Human and Automated Reasoning*, volume 1651 of *CEUR Workshop Proceedings*, 61–73. CEUR-WS.org. http://ceur-ws.org/Vol-1651/.

Stenning, K., and van Lambalgen, M. 2008. *Human Reasoning and Cognitive Science*. MIT Press.

Van der Auwera, J. 1997. Pragmatics in the last quarter century: The case of conditional perfection. *Journal of Pragmatics* 27(3):261–274.