Human Conditional Reasoning – The Case of Denying the Antecedent

Anonymous CogSci submission

Abstract

The weak completion semantics is a formal cognitive theory. It is multi-valued, non-monotonic, knowledge-rich, allows learning, can handle inconsistent background knowledge, and can be applied to model the average reasoner. In this paper we show that the case, where the antecedent of a conditional sentence is denied, can be fully modeled within the weak completion semantics, and we demonstrate that abduction can be restricted to find reasons for observations.

Keywords: conditionals; reasoning; countermodels; abduction; weak completion semantics

Introduction

The weak completion semantics (WCS) is a formal theory whose origin is a bug in the approach by Stenning and van Lambalgen (2008). The bug was corrected by switching from Kleene (1952) to Łukasiewicz (1920) three-valued logic (Hölldobler & Kencana Ramli, 2009). The main difference between the two logics is that conditionals with unknown antecedents and unknown consequents are true under Łukasiewicz, but unknown under Kleene logic; likewise for equivalences. This switch made it possible to formally prove that the so-called weak completion of a set of facts, assumptions, and rules has a least model. Moreover, this model can be computed as the least fixed point of a semantic operator specified by Stenning and van Lambalgen (2008). It is assumed to be a mental model in the sense of Craik (1945) and the initial reasoning in the WCS is defined with respect to this least model.

Since then, the WCS has been applied to various human reasoning tasks. For example, Dietz, Hölldobler, and Ragni (2012) have shown that the average reasoner in the suppression task (Byrne, 1989) can be adequately modeled. In this application, abduction (Hartshorne & Weiss, 1932) is applied to explain given observations. Several different explanations may exist, in which case skeptical reasoning is adequate, whereas credulous reasoning is not. In another application, Oliviera da Costa, Dietz Saldanha, Hölldobler, and Ragni (2017) have demonstrated that the WCS performs better than twelve cognitive theories in human syllogistic reasoning (Khemlani & Johnson-Laird, 2012).

This paper focuses on modeling human conditional reasoning within the WCS. In an experiment, Cramer, Hölldobler, and Ragni (2021b, 2021a) have asked participants about the consequences of a conditional sentence together with an atomic, possibly negated sentence. The conditional sentences were classified as obligational or factual conditionals with necessary or non-necessary antecedent by the authors of the experiment. The atomic sentences were either 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) of the conditional sentence. It turned out that the WCS adequately models the AC and the DC case, but failed to model the DA case. This was corrected by Bhadra (2021), who introduced counterexamples to the WCS framework and compared WCS to the mental model theory (Johnson-Laird, 1983, 2004); see also (Bhadra & Hölldobler, 2021). But these counterexamples were computed outside of the formal system of the WCS. In particular, the counterexamples were constructed outside of the abductive framework used by the WCS.

In this paper we will modify the WCS such that DA inferences can be computed fully within the formal system of the WCS. The modified approach will consider the given atomic, possibly negated sentences as observations which need to be explained. This seemingly simple modification leads to the problem that an observation like, for example, *Maria is drinking alcoholic beverages* is explained by itself, which is against the basic ideas underlying abuction. We will show in this paper that in a rich everyday setting this problem can be overcome and abduction can be applied to justify observations.

The Weak Completion Semantics

It is assumed that the reader is 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 $\top$, 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

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

where A, C, and ab are atoms. C and ab are defined, whereas A is undefined. The atom ab is an abnormality predicate which is assumed to be false. In the

WCS, this program represents the conditional sentence *if* A *then* C. In their everyday lives humans are often required to reason in situations where the information of all factors affecting the situation might not be complete. They still reason, unless new information which needs consideration comes to light. As suggested by Stenning and van Lambalgen (2005), the abnormality predicate occurring in the program serves the purpose of this (default) assumption.

Consider the following transformation: (1) For all defined atoms B occurring in $\mathcal{P}$, replace all clauses of the form $B \leftarrow body_1,\ B \leftarrow body_2,\ \ldots$ 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 $\mathcal{P}$ and will be denoted by $wc\mathcal{P}$. It differs from the program completion defined by Clark (1978) in that undefined atoms in the weakly completed program are not mapped to false, but to unknown instead.

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

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

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 program,

$$\begin{array}{rcl} \mathcal{A}_\mathcal{P} &=& \{B \leftarrow \top \mid B \text{ is undefined in } \mathcal{P}\}\\ &\cup& \{B \leftarrow \perp \mid B \text{ is undefined in } \mathcal{P}\} \end{array}$$

is the *set of abducibles*, $\mathcal{IC}$ is a finite set of *integrity constraints*, and $\mathcal{M}_{wc\mathcal{P}} \models_{wcs} F$ iff $\mathcal{M}_{wc\mathcal{P}}$ maps the formula F to true. Let O be an *observation*, i.e., a finite set of literals. 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 $\mathcal{M}_{wc(\mathcal{P}\cup\mathcal{X})} \models_{wcs} L$ for all $L \in O$ and $\mathcal{M}_{wc(\mathcal{P}\cup\mathcal{X})}$ satisfies $\mathcal{IC}$. Formula F *follows credulously* from $\mathcal{P}$ and O iff there exists an explanation $\mathcal{X}$ for O such that $\mathcal{M}_{wc(\mathcal{P}\cup\mathcal{X})} \models_{wcs} F$. F *follows skeptically* from $\mathcal{P}$ and O iff O can be explained and for all explanations $\mathcal{X}$ for O we find $\mathcal{M}_{wc(\mathcal{P}\cup\mathcal{X})} \models_{wcs} F$. One should observe that if an observation 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. Meaning, (unlike classical logic) *all* explanations for an observation O may only be taken into account to skeptically decide on a formula F, when O is explainable and these so-called explanations exist in the first place. If a formula F does not follow skeptically from $\mathcal{P}$ and O, we conclude *nothing follows*. Furthermore, one should also observe that if an observation O cannot be explained, then *nothing follows* credulously as well as skeptically. In some examples discussed in this paper, the set of integrity constraints is empty. However, they are needed in other applications of the WCS like human disjunctive reasoning (Hamada & Hölldobler, 2021).

[1] If 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.

A Classification of Conditionals

Obligational versus Factual Conditionals

Obligations are a familiar part of human lives. They legally, morally or even physically bind us to a certain course of action, following a premise. A conditional sentence whose consequent appears to be obligatory in a certain context, given the antecedent, is therefore called an *obligational conditional*. Byrne (2005) provides a good insight into the pragmatics of obligational conditionals. For each obligational conditional there are two initial possibilities humans think about. The first possibility is the conjunction of the antecedent and the consequent which is permitted. The second possibility is the conjunction of the antecedent and the negation of the consequent which is forbidden. As an example for an obligational conditional consider *if Maria is drinking alcoholic beverages in a pub, then she must be over 18 years of age*.

If the consequent of a conditional sentence is not obligatory given the antecedent, then it is called a *factual conditional*. In particular, the truth of the antecedent is *inconsequential* to that of the consequent; (even) if the antecedent is true, the consequent may or may not be true. This can be exemplified by the conditional sentence *if the plants get water, then they will grow*. While we would like to hope that simply watering plants would make them grow, it need not necessarily be the case. There are many other factors (for example lack of light, pest infestation etc.) which might yet hinder their growth. This implies that *plants get water* and *they will grow* is a permitted possibility, but *plants get water* and *they will not grow* is *not* a forbidden one; it is also permitted. Unlike the case of obligational conditionals, there is no forbidden possibility in case of factual conditionals.

Necessary versus Non-Necessary Antecedents

The antecedent A of a conditional sentence *if* A *then* C is said to be *necessary* (for the consequent C) if and only if C cannot be true unless A is true. This implies that if A does not hold, C cannot either. Following Byrne (2005), the possibility that while A is false, C is true, can be discounted in the case of necessary antecedents. For example, *plants get water* is a necessary antecedent for *plants will grow*. If a plant is not watered at all, it will very likely die.

The above does not imply however, that the antecedent need always be a precondition for the consequent, per se. The antecedent A of a conditional sentence *if* A *then* C is said to be *non-necessary* with re-

if A then C			
ON	ONN	FN	FNN
A C	A C	A C	A C
$\neg A$ $\neg C$	$\neg A$ $\neg C$	A $\neg C$	A $\neg C$
	$\neg A$ C	$\neg A$ $\neg C$	$\neg A$ $\neg C$
			$\neg A$ C

Table 1: All possibilities of antecedent A and consequent C for each classification of everyday conditional sentences.

$C \leftarrow A \wedge \neg ab$	A non-nec.	A nec.
Factual	$ab \leftarrow \top$, $C \leftarrow \top$	$ab \leftarrow \top$
Obligational	$C \leftarrow \top$	

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*, where A and C are ground atoms, ab is an abnormality predicate, and nec. is an abbreviation for necessary.

spect to the consequent C, if C can be true irrespective of the truth or falsity of A. In particular this implies, if A does not hold, C may or may not hold. *drinking alcoholic beverages in a pub* is inconsequential to the truth of the statement *older than 18 years*. There are plenty of adults (over 18 years) who do not drink alcohol. The antecedent of the conditional sentence *if Maria is drinking alcoholic beverages in a pub, then Maria must be over 18 years of age* is therefore called non-necessary.

Classification of Conditional Sentences

The previous subsections introduced the notions of obligation and factual conditionals, and necessary and non-necessary antecedents. This in turn gives rise to an informal and pragmatic classification of four kinds: obligation conditional with necessary antecedent (ON) or non-necessary antecedent (ONN), and factual conditional with necessary antecedent (FN) or non-necessary antecedent (FNN). That is, given a conditional sentence *if A then C*, the sentence may be classified as ON, ONN, FN, or FNN depending on how it is comprehended. Combining these aforementioned distinguishing set of possibilities for each feature of a conditional, then allows us to characterize the set of possibilities for each classification. This is illustrated in the Table 1.

The classification of everyday conditionals often depends on pragmatics: the context, the background knowledge and experience of a person. For example, the conditional sentence *if it is cloudy, then it is raining* discussed by Khemlani, Byrne, and Johnson-Laird (2018) may be classified as ON by people living in Java, whereas it may be classified as FN by people living in Central Europe.

Representing the Semantics of Conditionals

Obligational and factual conditionals are represented using programs as before. A conditional *if A then C*, where A and C are atoms, is represented by

$$\mathcal{P}_0 = \{C \leftarrow A \wedge \neg ab,\ ab \leftarrow \bot\}.$$

However, the semantics of conditionals is taken into consideration by extending the set of abducibles for a given program $\mathcal{P}$ to $\mathcal{A}^e_{\mathcal{P}} = \mathcal{A}_{\mathcal{P}} \cup \mathcal{A}^{nn}_{\mathcal{P}} \cup \mathcal{A}^f_{\mathcal{P}}$, where $\mathcal{A}^{nn}_{\mathcal{P}}$ is the set of all facts of the form $C \leftarrow \top$, where C is head of a rule occurring in $\mathcal{P}$ representing a conditional with non-necessary antecedent, and $\mathcal{A}^f_{\mathcal{P}}$ is the set of all facts of the form $ab \leftarrow \top$ where ab occurs in the body of a rule occurring in $\mathcal{P}$ representing a factual conditional. In other words, 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 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 occurring in the body of the representation of a factual conditional 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. Table 2 illustrates the new facts in the set of abducibles. These facts were first specified in (Dietz Saldanha, Hölldobler, & Lourêdo Rocha, 2017).

An Experiment

Cramer et al. (2021b, 2021a) have performed an experiment concerning conditional reasoning, where 56 logically naive participants were tested on an online website (Prolific, `prolific.co`). The participants were restricted to Central Europe and Great Britain to have a similar background knowledge about weather etc. It was also assumed that the participants had not received any education in logic beyond high school training. The participants were given a conditional premise and an atomic, possibly negated premise. 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. Each participant solved all four inference types (AA, DA, AC, DC) for each of the 12 conditionals, thereby carrying out 48 inference tasks in total. For each task they could select one of three responses: *nothing follows (nf)*, the fact that had not been presented in the second premise, and the negation of this fact. For example in the case of a DA inference task, the first assertion was of the form *if A then C*, the second assertion was $\neg A$, and they could choose among answers C, $\neg C$ or *nf*. The classification of the 12 conditionals into the four aforementioned kinds was done by the authors beforehand, and was not mentioned to the participants.

As shown by Cramer et al. (2021b, 2021a), the WCS can adequately handle the DC, the AC, as well as the AA inference tasks, but fails to handle the DA

Class	C	$\neg C$	nf	Sum	$Mdn\,\neg C$	$Mdn\,nf$
O	12	254	70	336	3583	6613
F	6	271	59	336	3518	6221
N	8	300	28	336	3474	5808
NN	10	225	101	336	3646	6700
Total	18	525	129	672	3558	6450

Table 3: The results for DA inferences given a conditional sentence *if A then C* and a negated atomic fact $\neg A$. O and F denote obligational and factual conditionals, respectively. N and NN denote necessary and non-necessary antecedents, respectively. $Mdn\,\neg C$ and $Mdn\,nf$ show the median response time in milliseconds for $\neg C$ and *nf*, respectively. It may be observed that when the antecedent is non-necessary, *nf* is answered significantly more often, whereas the number for $\neg C$ answers decreases (see grey cells).

case. Hence, in this paper we focus on the DA inference task, whose result is summarized in Table 3. The reader might first observe that as per the experiment data *nothing follows* was answered much more often in case of conditional sentences with non-necessary antecedents than in the case of conditional sentences with necessary ones (30% vs. 8%, Wilcoxon signed rank, $W = 0$, $p < .001$). More importantly, as soon as the classification of the antecedent changed from necessary to non-necessary, the number of $\neg C$ responses decreased to 225 and *nf* simultaneously increased to (a significant) 101.

As an example consider the conditional sentence *if Lisa plays on the beach, then Lisa will get sunburned* or, for short, *if b then s*, where b and s denote *Lisa plays on the beach* and *Lisa will get sunburned*, respectively. The conditional is classified as FNN. In a DA reasoning task the second premise is *Lisa does not play on the beach*. In the original version of the WCS (see e.g. (Cramer et al., 2021a; Dietz et al., 2012)) this reasoning problem was represented by the program

$$\mathcal{P}_1 = \{s \leftarrow b \wedge \neg ab_b,\ ab_b \leftarrow \bot,\ b \leftarrow \bot\}.$$

Weakly completing the program we obtain

$$wc\mathcal{P}_1 = \{s \leftrightarrow b \wedge \neg ab_b,\ ab_b \leftrightarrow \bot,\ b \leftrightarrow \bot\}$$

whose least model $\mathcal{M}_{wc\mathcal{P}_1}$ is

$$\langle \emptyset, \{s, b, ab_b\}\rangle. \tag{1}$$

Hence, we conclude that *Lisa will not get sunburned.* It adequately models the majority of the human reasoners in the experiment. In fact, 41 of the 56 participants answered $\neg b$, whereas 15 answered *nf* (see (Cramer et al., 2021a)). On the other hand, it cannot account for the number of *nf* answers in this case. One should observe that in this line of modeling, abduction was not applied at all and, hence, the classification of conditionals was not taken into account.

Bhadra and Hölldobler (2021) had noted that there is in fact the possibility that *Lisa will get sunburned* although *Lisa does not play on the beach.* As shown in Table 1, a conditional sentence classified as FNN like *if b then c* admits the possibility that the antecedent is false and the consequent is true. This possibility is expressed in the interpretation

$$\langle \{s\}, \{b, ab_b\}\rangle. \tag{2}$$

Comparing (1) and (2) while reasoning skeptically we can no longer conclude that *Lisa will not get sunburned*. Rather, *nothing follows*.

But this line of reasoning is outside the formal system defining the WCS so far. Thus, the question remains whether this line of reasoning can be incorporated or aligned with the WCS. One should note that the interpretation (2) is a model for $\mathcal{P}_1$ but not for $wc\mathcal{P}_1$.

Atomic Premises are Observations

In program $\mathcal{P}_1$ the atomic premise *Lisa does not play on the beach* was represented as the assumption $b \leftarrow \bot$ in program $\mathcal{P}_1$. What will happen if we only represent the conditional sentence as the program

$$\mathcal{P}_2 = \{s \leftarrow b \wedge \neg ab_b,\ ab_b \leftarrow \bot\}$$

and consider the atomic premise as observation $O = \{\neg b\}$? As b is undefined in $\mathcal{P}_2$ we obtain

$$\mathcal{A}_{\mathcal{P}_2} = \{b \leftarrow \top,\ b \leftarrow \bot\}.$$

As the conditional sentence is classified as FNN we obtain

$$\mathcal{A}^{nn}_{\mathcal{P}_2} = \{ab_b \leftarrow \top\} \quad \text{and} \quad \mathcal{A}^{f}_{\mathcal{P}_2} = \{s \leftarrow \top\}.$$

Hence,

$$\mathcal{A}^{e}_{\mathcal{P}_2} = \{b \leftarrow \top,\ b \leftarrow \bot,\ ab_b \leftarrow \top,\ s \leftarrow \top\}.$$

The observation $O = \{\neg b\}$ can be explained by the minimal explanation

$$\mathcal{X} = \{b \leftarrow \bot\}.$$

Adding $\mathcal{X}$ to the program $\mathcal{P}_2$ we obtain $\mathcal{P}_1$ and the model (1) of $wc\mathcal{P}_1$ as before.

But now there is a second, non-mininal explanation for $O = \{\neg b\}$, viz.

$$\mathcal{X}' = \{b \leftarrow \bot,\ s \leftarrow \top\}.$$

Adding $\mathcal{X}'$ to the program $\mathcal{P}_2$ we obtain

$$\mathcal{P}_2 \cup \mathcal{X}' = \{s \leftarrow b \wedge \neg ab_b,\ ab_b \leftarrow \bot,\ b \leftarrow \bot,\ s \leftarrow \top\}.$$

Weakly completing this extended program yields

$$wc(\mathcal{P}_2 \cup \mathcal{X}') = \{s \leftrightarrow (b \wedge \neg ab_b) \vee \top, ab_b \leftrightarrow \bot, b \leftrightarrow \bot\}$$

whose least model is (2). Taking the second non-minimal explanation into account and reasoning skeptically, we can no longer conclude that *Lisa will not get sunburned*, but must conclude *nothing follows*. Moreover, by considering atomic premises as observations, the reasoning process is completely within the formal framework of the WCS.

By considering the minimal explanation $\mathcal{X}$, a preferred model (1) is computed. Most of the participants in the experiment seem to reason with respect to this preferred model. As reported by (Cramer et al., 2021a), the participants needed on average 3374 ms to answer *Lisa will not get sunburned*. However, a significant number of the participants seem to deliberate and generate a second model, which does not map s to true. Hence, this is a countermodel to (1) with respect to s. As reported by (Cramer et al., 2021a), these participants needed 5887 ms on average for answering *nothing follows*. This can be explained by the fact that they need to generate a second model and need to compare the two models. The reasoning process seems to be in line with human spatial reasoning, in which case Ragni and Knauff (2013) have confirmed that most humans generate preferred models and reason with respect to them.

Justifying Observations

The solution presented in the previous section may be criticized by noting that the observation *Lisa is not playing on the beach* is explained by *Lisa is not playing on the beach*. This is not really at the heart of abduction where we want to justify observations and not simply repeat them.

As we will argue in this section, this critique can be overcome if we consider a richer background knowledge. Let us consider the conditional sentence *if Maria is drinking alcoholic beverages in a pub, then Maria must be over 18 years of age* or, for short, *if a then o*, where a and o denote that *Maria is drinking alcoholic beverages in a pub* and *Maria must be over 18 years of age*, respectively. This conditional sentence is classified as ONN and represented by the program

$$\mathcal{P}_4 = \{o \leftarrow a \wedge \neg ab_a,\ ab_a \leftarrow \bot\}$$

with

$$\mathcal{A}^e_{\mathcal{P}_4} = \{a \leftarrow \top,\ a \leftarrow \bot,\ o \leftarrow \top\}.$$

Now, consider the observation $\mathcal{O} = \{a\}$, ie., *Maria is drinking alcoholic beverages in a pub*. This observation is explained by the minimal explanation $\{a \leftarrow \top\}$.

Well, in most pubs there is a drinks menu and, very likely, this drinks menu has at least two columns: alcoholic and non-alcoholic beverages. Let us assume that among the alcoholic beverages there are beer (b), wine (w), and liquor (ℓ) and that among the non-alcoholic beverages there are juice (j) and soft-drinks (s). This background knowledge can be encoded by the clauses

$$\begin{array}{rclrcl} a & \leftarrow & b \wedge \neg ab_b, & ab_b & \leftarrow & \bot, \\ a & \leftarrow & w \wedge \neg ab_w, & ab_w & \leftarrow & \bot, \\ a & \leftarrow & \ell \wedge \neg ab_\ell, & ab_\ell & \leftarrow & \bot, \\ na & \leftarrow & j \wedge \neg ab_j, & ab_j & \leftarrow & \bot, \\ na & \leftarrow & s \wedge \neg ab_s, & ab_s & \leftarrow & \bot. \end{array}$$

Let $\mathcal{P}_5$ be the union of these clauses and $\mathcal{P}_4$. Moreover, we need to specify with the help of the integrity constraint

$$\bot \leftarrow a \wedge na \tag{3}$$

that a drink cannot be alcoholic or non-alcoholic at the same time.

In $\mathcal{P}_5$ the atom a is defined, whereas b, w, ℓ, j, and s are undefined. Hence,

$$\begin{array}{rcl} \mathcal{A}_{\mathcal{P}_5} & = & \{b \leftarrow \top,\ b \leftarrow \bot,\ w \leftarrow \top,\ w \leftarrow \bot\} \\ & \cup & \{\ell \leftarrow \top,\ \ell \leftarrow \bot,\ j \leftarrow \top,\ j \leftarrow \bot\} \\ & \cup & \{s \leftarrow \top,\ s \leftarrow \bot\}. \end{array}$$

Now, the observation $\mathcal{O} = \{a\}$ can be explained by the minimal explanations

$$\{b \leftarrow \top\}, \quad \{w \leftarrow \top\}, \quad \text{and} \quad \{\ell \leftarrow \top\}.$$

No matter, which explanation is added to the program $\mathcal{P}_5$, the weakly completed extended program admits a minimal model, where both, a and o, are mapped to true, whereas na is mapped to unknown. Hence, the integrity constraint (3) is satisfied and, reasoning skeptically, we will conclude that *Maria is drinking acoholic beverages in a pub* and *Maria must be over 18 years of age.*

If we observe $\mathcal{O} = \{\neg a\}$, then there is a slight complication because the head of a rule in a program must be an atom. But there is a standard method in logic programming to deal with such cases. This method has already been applied in human syllogistic reasoning by WCS in (Oliviera da Costa et al., 2017). It can be applied here as well: We may add to the program $\mathcal{P}_5$ the clauses

$$a \leftarrow \neg na \wedge ab_{na} \quad \text{and} \quad ab_{na} \leftarrow \bot$$

to obtain program $\mathcal{P}_7$. The weak completion of $\mathcal{P}_7$ contains

$$a \leftrightarrow \neg na \wedge ab_{na} \quad \text{and} \quad ab_{na} \leftrightarrow \bot.$$

Hence, we have specified that the atom na is the negation of a. Now, the observation $\mathcal{O} = \{\neg a\}$ can be explained by the minimal explanations

$$\{j \leftarrow \top\} \quad \text{and} \quad \{s \leftarrow \top\}.$$

In either case, we learn that na is true, a is false, and the integrity constraint (3) is satisfied.

But as in the previous section, there are also non-minimal explanations for the observation $\mathcal{O} = \{\neg a\}$ like

$$\{j \leftarrow \top,\ o \leftarrow \top\} \quad \text{or} \quad \{s \leftarrow \top,\ o \leftarrow \top\}.$$

In words, *Maria is drinking juice or a soft-drink* and is, nevertheless, *over 18 years of age*. Hence, a careful and skeptical reasoner taking into account minimal and non-minimal explanations for the observation $\mathcal{O} = \{\neg a\}$ will answer *nothing follows*. In fact, as reported by Cramer et al. (2021a), 25 out of 56 participants drew this conclusion and it took them on average 7814 ms to do so, whereas 28 out of 56 participants answered $\neg o$ in 5735 ms on average, and 3 out of 56 answered o.

Again, both non-minimal explanations explain *over 18 years of age* by being *over 18 years of age*. As before, this critique can be overcome if we consider a richer background knowledge, where – for example – it is stated that in Germany *if Maria is driving alone in a car, then Maria must be over 18 years* *of age* or *if Maria is signing legal documents, then Maria must over 18 years of age*. If these conditional statements are added to the program, then *being over 18 years of age* can be explained by *driving a car* or by *signing a legal document*.

Discussion

In this paper the WCS has been extended to handle everyday DA inference tasks adequately. The main idea was to consider atomic, possibly negated premises as observations that need to be explained. We showed that the WCS can distinguish two kinds of reasoners: A fast reasoner taking into account only minimal explanations leading to preferred models and a deliberative reasoner taking also non-minimal explanations into account which may lead to counter-models and, hence, to different and revised answers. The extended WCS consists of the following steps, where Step 6 is the novel contribution:

1. Reasoning towards a logic program $\mathcal{P}$ following Stenning and van Lambalgen (2008).
2. Weakly completing the program to obtain $wc\mathcal{P}$.
3. Computing its least model $\mathcal{M}_{wc\mathcal{P}}$ under the three-valued Łukasiewicz (1920) logic.
4. Reasoning with respect to $\mathcal{M}_{wc\mathcal{P}}$.
5. If known observations cannot be explained or given integrity constraints cannot be met, applying skeptical abduction.
6. Searching for counterexamples.

The extended WCS matches the experimental data not only with respect to the responses but also the time taken. Moreover, we also provided a critical discussion of how within the abduction step observations will not be explained by themselves if a richer everyday background knowledge is considered.

The extended WCS has been compared to the mental model theory (MMT) in (Bhadra, 2021; Bhadra & Hölldobler, 2021), however, it is beyond the scope of this paper to repeat this comparison in detail. Just some remarks are appropriate. WCS and MMT largely agree on most of the results, but whereas the WCS is a formal theory, MMT is an informal one. MMT does not match some of the times taken as reported by (Cramer et al., 2021b, 2021a).

In all examples we have added abnormality predicates although they are not used in the paper. Abnormality predicates are used to model enabling constraints or exceptions. For example, non-alcoholic beer has become fashionable these days in Germany. Taking brand names into account we can write

$$aX \leftarrow bX \wedge \neg ab_b X,\ ab_b X \leftarrow \bot$$

to encode that *beer of brand X is usually alcoholic*, where X is a (first-order) variable. Adding

$$ab_b\, e \leftarrow \top$$

now states that *beer of brand e is abnormal* in the sense that it is non-alcoholic, where e is a constant symbol. Considering the weak completion of the (grounded) program we can no longer conclude that *the beer of brand e is alcoholic*, whereas for all other brands this still holds.

The WCS can be implemented by a recursive connectionist network, where the semantic operator computing the least models is computed by a feed-forward subnetwork (see (Dietz Saldanha, Hölldobler, Kencana Ramli, & Palacios Medinacelli, 2018)). This subnetwork can be trained and – after training – a revised semantic operator for a revised program can be extracted. Thus, learning is possible.

Several open problems remain to be investigated in the future.

- It remains to be seen whether the distinction between a fast and a deliberative reasoner fits the specification of fast and slow thinking by Kahnemann (2011).

- The extended WCS allows to come up with non-minimal models in the case of AA inferences. For example, consider the conditional sentence *if the plants get water, they will grow* and the atomic sentence *the plants get water*. As discussed, this conditional sentence is classified as FN as *watering plants* may not necessarily lead to *their growth*. Using w and g to denote *plants get water* and *plants will grow*, respectively, the conditional sentence is represented by the program

  $$\mathcal{P}_7 = \{g \leftarrow w \wedge \neg ab_w,\ ab_w \leftarrow \bot\}$$

  with

  $$\mathcal{A}_{\mathcal{P}_7} = \{w \leftarrow \top,\ w \leftarrow \bot,\ ab_w \leftarrow \top\}.$$

  The observation $O = \{w\}$ can be explained by the minimal explanation $\{w \leftarrow \top\}$. Adding this explanation to the program and weakly completing the extended program yields the (preferred) model $\langle\{g, w\}, \{ab_w\}\rangle$. However, the observation can also be explained by the non-minimal explanation $\{w \leftarrow \top,\ ab_w \leftarrow \top\}$. Adding the non-minimal explanation to the program and weakly completing the extended program yields the model $\langle\{w, ab_w\}, \{g\}\rangle$ and one can no longer conclude that g is true. As reported by (Cramer et al., 2021a) 54 out of 56 participants answered with respect to the preferred model. This model can also be computed by modus ponens and there seems to be a strong reason to favor it.

- Finally, we need to better understand when to extend the set of abducibles. For example, when president Kennedy was killed by a bullet and either *Oswald shot* or *somebody else shot* (see (Adams, 1970)), then the disjunction of these statements completely describes the state of affairs. It does not make sense to consider yet another person. On the other hand, if beer, wine, and liquor are explicitely mentioned as alcoholic beverages, then this does not completely describe the state of affairs as other alcoholic beverages are known.

Acknowledgments

References

Adams, E. W. (1970). Subjunctive and indicative conditionals. Foundations of Language, 6(1), 89–94.

Bhadra, M. (2021). Comparative studies of the weak completion semantics and the mental model theory. Unpublished master's thesis, Technische Universität Dresden, Fakultät für Informatik.

Bhadra, M., & Hölldobler, S. (2021). The weak completion semantics and counter examples. In C. Beierle, M. Ragni, F. Stolzenburg, & M. Thimm (Eds.), Proceedings of the 7th workshop on formal and cognitive reasoning. (to appear)

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 H. Gallaire & J. Minker (Eds.), Logic and databases (p. 293-322). New York: Plenum.

Craik, K. J. W. (1945). The nature of explanation. Cambridge: Cambridge University Press.

Cramer, M., Hölldobler, S., & Ragni, M. (2021a). Modeling human reasoning about conditionals. In L. Amgoud & R. Booth (Eds.), Proceedings of the 19th international workshop on non-monotonic reasoning (pp. 223–232).

Cramer, M., Hölldobler, S., & Ragni, M. (2021b). When are humans reasoning with modus tollens? Proceedings of the Annual Conference of the Cognitive Science Society,, 43, 2337-2343. (Retrieved from `https://escholarship.org/uc/item/9x33q50g`)

Dietz, E.-A., Hölldobler, S., & Ragni, M. (2012). A computational logic approach to the suppression task. Proceedings of the Annual Conference of the Cognitive Science Society,, 34, 1500-1505. (Retrieved from https://escholarship.org/uc/item/2sd6d61q)

Dietz Saldanha, E.-A., Hölldobler, S., Kencana Ramli, C. D. P., & Palacios Medinacelli, L. (2018). 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., & Lourêdo Rocha, I. (2017). Obligation versus factual conditionals under the weak completion semantics. In S. Hölldobler, A. Malikov, & C. Wernhard (Eds.), Proceedings of the Second Young Scientists' International Workshop on Trends in Information Processing (Vol. 1837, p. 55-64). (`http://ceur-ws.org/Vol-1837/`)

Fitting, M. (1996). First–order logic and automated theorem proving (2nd ed.). Berlin: Springer-Verlag.

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), Syntax and semantics (Vol. 3, p. 41-58). Academic Press, New York.

Hamada, I., & Hölldobler, S. (2021). On disjunctions and the weak completion semantics. In Proceedings of the virtual mathpsych/iccm. via `mathpsych.org/presentation/571`.

Hartshorne, C., & Weiss, P. (Eds.). (1932). Collected papers of charles sanders peirce (Vol. 2). Harvard University Press.

Hölldobler, S. (2015). Weak completion semantics and its applications in human reasoning. In U. Furbach & C. Schon (Eds.), Bridging 2015 – bridging the gap between human and automated reasoning (Vol. 1412, p. 2-16). CEUR-WS.org. (`http://ceur-ws.org/Vol-1412/`)

Hölldobler, S., & Kencana Ramli, C. D. P. (2009). Logic programs under three-valued Łukasiewicz's semantics. In P. M. Hill & D. S. Warren (Eds.), Logic programming (Vol. 5649, p. 464-478). Springer-Verlag Berlin Heidelberg.

Johnson-Laird, P. N. (1983). Mental models: Towards a cognitive science of language, inference, and consciousness. Cambridge, MA, USA: Harvard University Press.

Johnson-Laird, P. N. (2004). The history of mental models. In K. Manktelow & M. C. Chung (Eds.), Psychology of reasoning (p. 179-212). Psychology Press.

Kahnemann, D. (2011). Thinking, fast and slow. Farrar, Straus and Giroux.

Kakas, A. C., Kowalski, R. A., & Toni, F. (1992). Abductive Logic Programming. Journal of Logic and Computation, 2(6), 719-770.

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

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

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. (English translation: On Three-Valued Logic. In: *Jan Łukasiewicz Selected Works*. (L. Borkowski, ed.), North Holland, 87-88, 1990.)

Oliviera da Costa, A., Dietz Saldanha, E.-A., Hölldobler, S., & Ragni, M. (2017). A computational logic approach to human syllogistic reasoning. Proceedings of the Annual Conference of the Cognitive Science Society,, 39, 883-888.

Ragni, M., & Knauff, M. (2013). A theory and a computational model of spatial reasoning. Psychological Review, 120, 561-588.

Stenning, K., & van Lambalgen, M. (2005). Semantic interpretation as computation in nonmonotonic logic: The real meaning of the suppression task. Cognitive Science, 29(919-960).

Stenning, K., & van Lambalgen, M. (2008). Human reasoning and cognitive science. MIT Press.