

Chapter 1

Our Themes on Abduction in Human Reasoning: A Synopsis

Emmanuelle-Anna Dietz Saldanha¹, Steffen Hölldobler^{1,2}, Luís Moniz Pereira^{3*}

Abstract Psychological experiments have shown that humans do not reason according to classical logic. Therefore, we might argue that logic-based approaches in general are not suitable for modeling human reasoning. Yet, we take a different view and are convinced that logic can help us as an underlying formalization of a cognitive theory, but claim rather that classical logic is not adequate for this purpose. In this chapter we investigate abduction and its link to human reasoning. In particular we discuss three different variations we have explored and show how they can be adequately modeled within a novel computational and integrated, cognitive theory, the Weak Completion Semantics.

1.1 Introduction

Originally, one of the objectives of using logic in artificial intelligence and knowledge representation and reasoning was the formalization of human and common-sense reasoning. During the past decades the original objective shifted out of focus. The problem description was not specified adequately, possibly because there was little or no communication with cognitive scientists. Non-classical logic approaches exist, however most of them are purely theoretical and not applied to actual human case studies. Instead, artificial examples have been constructed, which only show that a theory works within some very specific context. But what is the value of a

¹International Center for Computational Logic, TU Dresden, 01062 Dresden, Germany. {dietz,sh}@iccl.tu-dresden.de

²North-Caucasus Federal University, Stavropol, Russian Federation.

³NOVA Laboratory for Computer Science and Informatics, Departamento de Informática, Faculdade de Ciências e Tecnologia, Universidade Nova de Lisboa, 2829-516 Caparica, Portugal. lmp@fct.unl.pt

* The authors are mentioned in alphabetical order.

cognitive theory that has never been evaluated against the conclusions that humans actually draw?

Taking this observation as starting point, several human reasoning episodes, ranging from the suppression and the selection task [6, 9], to spatial, syllogistic [5] and counterfactual reasoning [7, 8] have been successfully modeled under a new, computational and integrated, cognitive theory, the *Weak Completion Semantics*. In a nutshell, under the *Weak Completion Semantics* the following steps are taken for given a human reasoning scenario:

1. Reasoning towards a (logic) program representation.
2. Weakly completing the program.
3. Computing its least model under Łukasiewicz logic.
4. Reasoning with respect to the least model.
5. If necessary, applying abduction.

The *Weak Completion Semantics* is based on ideas first presented in [22, 23]. Unfortunately, these first ideas contained a technical bug, which was corrected in [13] by switching from three-valued Kleene logic [16] to three-valued Łukasiewicz logic [18]. As shown in [13], each weakly completed program admits a least model which can be computed as the least fixed point of an appropriate semantic operator. Unsurprisingly, it turned out that some human reasoning tasks require abduction and, hence, the *Weak Completion Semantics* has been extended with abduction.

In this chapter, we will avoid introducing formal definitions, and assume the reader to be familiar with logic and logic programming. The interested reader is referred to [17, 12]. The goal of this chapter is rather to give the intuitions behind the different facets that abduction can take within human reasoning. By means of four episodes of human reasoning, *the suppression task*, *counterfactual reasoning*, *contextual reasoning* and *reasoning with obligation and factual conditionals*, we will show that abductive reasoning is required. Interestingly, a different variation of abduction applies for each of the four cases in order to be adequately modeled under the *Weak Completion Semantics*.

1.2 The Suppression Task

The *suppression task* is a famous psychological experiment originally carried out by Ruth Byrne [3] showing that humans suppress previously drawn conclusion if new information becomes available.

Suppose that participants are told that

if she has an essay to finish then she will study late in the library

and that

she will study late in the library.

Then, the subjects are asked whether they are willing to conclude that *she has an essay to write*. As Byrne reports, 71% of the subjects are willing to draw this conclusion.

Following [22, 23], the conditional is translated into the program

$$\{library \leftarrow essay \wedge \neg ab_1, ab_1 \leftarrow \perp\},$$

whose weak completion is

$$\{library \leftrightarrow essay \wedge \neg ab_1, ab_1 \leftrightarrow \perp\},^2$$

where *essay* denotes that *she has an essay to finish*, *library* that *she will study late in the library* and *ab₁* is an *abnormality* predicate that can deal with possible exceptions to the rule. The weakly completed program admits a least model under (three-valued) Łukasiewicz logic,³ where *library* is mapped to true, *ab₁* is mapped to false, and *essay* is mapped to unknown. Adding the fact *library* $\leftarrow \top$ to the program does not change this least model. Hence, we should not add it but rather consider *library* as an observation that needs to be supported by an explanation. In the given context, the minimal explanation being $\{essay \leftarrow \top\}$. If added to the program, we obtain the weakly completed program

$$\{library \leftrightarrow essay \wedge \neg ab_1, essay \leftrightarrow \top, ab_1 \leftrightarrow \perp\},$$

whose least model maps *library* and *essay* to true and *ab₁* to false. Reasoning with respect to this model allows to abductively conclude *essay*, which most subjects did.

Hypothesis 1: Humans reason abductively.

In another group of the same experiment participants were told:

If she has an essay to finish, then she will study late in the library.

If she has a textbook to read, then she will study late in the library.

All other information being equal to the previous experiment, the number of subjects willing to conclude that *she has an essay to write* drops to 13% [3]. The previously drawn conclusion is suppressed as soon as the second conditional is given. Modelling this reasoning episode similarly to before, we obtain the weakly completed program

$$\{library \leftrightarrow (essay \wedge \neg ab_1) \vee (textbook \wedge \neg ab_2), ab_1 \leftrightarrow \perp, ab_2 \leftrightarrow \perp\},$$

² One should observe that under completion as defined in [4], $essay \leftrightarrow \perp$ would be an element of this set. The *Weak Completion Semantics* however does not map undefined relations to false but rather considers them unknown.

³ In the sequel, all models are computed with respect to three-valued Łukasiewicz logic where not explicitly mentioned otherwise.

where *textbook* denotes that *she has a textbook to read*. Considering *library* again as an observation to be justified, there are two minimal explanations, viz. $\{essay \leftarrow \top\}$ and $\{textbook \leftarrow \top\}$. Reasoning credulously, i.e. when one explanation suffices to conclude *essay*, then one would conclude that *she has an essay to write*. However, humans do not do this. Their conclusions appear to be modelled adequately when skeptical abduction is applied.

Hypothesis 2: Humans prefer skeptical over credulous abduction.

As shown in [14, 6], all twelve cases of the suppression task can be adequately modelled under the *Weak Completion Semantics*. For six of these cases, skeptical abduction is required.

1.3 Counterfactual Reasoning

Consider the following scenario [1]: *President Kennedy was killed. There was a lengthy investigation about whether Oswald or somebody else shot the president. In the end, it was determined that Oswald did it.* Which of the following conditionals do humans accept easily?

If Kennedy is dead and Oswald did not shoot Kennedy then someone else did.
If Oswald had not shot Kennedy then someone else would have.

According to [1], humans accept the former, but reject the latter conditional.

In this case, the background knowledge is encoded in the program

$$\{ \text{kennedy} \leftarrow \text{oswald} \wedge \neg \text{ab}_o, \quad \text{ab}_o \leftarrow \perp, \\ \text{kennedy} \leftarrow \text{someone_else} \wedge \neg \text{ab}_s, \quad \text{ab}_s \leftarrow \perp, \quad \text{oswald} \leftarrow \top \},$$

where *kennedy* denotes that *Kennedy was killed*, *oswald* denotes *Oswald shot*, *someone_else* denotes *somebody else shot*, and *ab_o* as well as *ab_s* are introduced abnormality predicates. As the weak completion of this program we obtain

$$\{ \text{kennedy} \leftrightarrow (\text{oswald} \wedge \neg \text{ab}_o) \vee (\text{someone_else} \wedge \neg \text{ab}_s), \\ \text{ab}_o \leftrightarrow \perp, \quad \text{ab}_s \leftrightarrow \perp, \quad \text{oswald} \leftrightarrow \top \},$$

whose least model maps *oswald* and *kennedy* to true and *ab_o* and *ab_s* to false.

The antecedent $\neg \text{oswald}$ of the second conditional is false in this least model. Hence, the conditional is a *counterfactual*. In order to evaluate this counterfactual we must revise the background knowledge. The *Weak Completion Semantics* does this in a straightforward and minimal way by replacing the positive fact $\text{oswald} \leftarrow \top$ with the negative assumption $\text{oswald} \leftarrow \perp$ in the program. Consequently, the least model of the weakly completed modified program maps *oswald*, *ab_o* and *ab_s* to false and *kennedy* to unknown. If we evaluate the second conditional with respect

to this least model we find that the antecedent $\neg oswald$ of the conditional is true, whereas its consequence $kennedy$ is unknown. Thus, the conditional itself is evaluated as unknown.

If we evaluate the first conditional with respect to the least model of the weak completion of the original program, then its antecedent $kennedy \wedge \neg oswald$ is also false as the literal $\neg oswald$ is mapped to false. If the original program is revised as before, then $\neg oswald$ is mapped to true, but now $kennedy$ is unknown with respect to the least model of the weakly completed modified program. Hence, the antecedent $kennedy \wedge \neg oswald$ is unknown. However, $kennedy$ can be explained by abduction. The only minimal explanation is $\{someone_else \leftarrow \top\}$. Adding this explanation to the revised program and weakly completing it, we obtain

$$\{ kennedy \leftrightarrow (oswald \wedge \neg ab_o) \vee (someone_else \wedge \neg ab_s), \\ ab_o \leftrightarrow \perp, \quad ab_s \leftrightarrow \perp, \quad oswald \leftrightarrow \perp \quad someone_else \leftrightarrow \top \},$$

whose least model maps $someone_else$ and $kennedy$ to true and $oswald$, ab_o and ab_s to false. Under this model the antecedent $kennedy \wedge \neg oswald$ as well as the consequent $someone_else$ of the first conditional are mapped to true. Consequently, the conditional itself is evaluated to true.

Hypothesis 3: If needed, humans minimally revise their background knowledge before applying abduction.

Various counterfactual reasoning scenarios, such as Pearl's *firing squad* scenario [19] and Byrne's *forest fire* scenario [2] can also be adequately modeled by applying *minimal revision followed by abduction* under the *Weak Completion Semantics*, as detailed in [8].

1.4 Contextual Reasoning

How can we prefer explanations that explain the normal cases to those explanations that explain the exceptional cases? How can we express that some explanations have to be considered only if there is some evidence for considering the exception cases? We want to avoid having to consider all explanations if there is no evidence at all for considering exception cases. On the other hand, we don't want to state that all exception cases are false.

Let us consider the famous Tweety example from [20] is as follows:

Usually birds can fly. Tweety and Jerry are birds.

According to [20], if nothing else is known about Tweety and Jerry, then we can deduce that both can fly.

This scenario can be easily modeled in the *Weak Completion Semantics*. The background knowledge is encoded in the program

$$\{fly X \leftarrow bird X \wedge \neg ab X, ab X \leftarrow \perp, bird tweety \leftarrow \top, bird jerry \leftarrow \top\},$$

where X a variable, $fly X$ ⁴ denotes that X can fly, $bird X$ that X is a bird and ab is an abnormality predicate. The corresponding ground program is

$$\{ fly tweety \leftarrow bird tweety \wedge \neg ab tweety, \quad ab tweety \leftarrow \perp, \quad bird tweety \leftarrow \top, \\ fly jerry \leftarrow bird jerry \wedge \neg ab jerry, \quad ab jerry \leftarrow \perp, \quad bird jerry \leftarrow \top \}.$$

Weakly completing this program we obtain

$$\{ fly tweety \leftrightarrow bird tweety \wedge \neg ab tweety, \quad ab tweety \leftrightarrow \perp, \quad bird tweety \leftrightarrow \top, \\ fly jerry \leftrightarrow bird jerry \wedge \neg ab jerry, \quad ab jerry \leftrightarrow \perp, \quad bird jerry \leftrightarrow \top \}.$$

whose least model maps $bird tweety$, $bird jerry$, $fly tweety$ and $fly jerry$ to true and $ab tweety$ as well as $ab jerry$ to false. Reasoning with respect to this least model we conclude that *Tweety and Jerry can fly*.

Suppose we observe

Jerry does not fly.

This observation cannot be explained in the usual abductive framework as specified in [15]. In such a framework the set of abducibles is defined to be the set of undefined relations with respect to the given program, where a relation or ground atom A is said to be *undefined* given a program if (the ground instance of) the program does not contain a rule of the form $A \leftarrow Body$, where $Body$ is a conjunction of literals. In the program above all relations ($fly jerry$, $fly tweety$, $bird jerry$, $bird tweety$, $ab jerry$, $ab tweety$) are defined.

However, under the *Weak Completion Semantics*, negative assumptions like $ab jerry \leftarrow \perp$ can be *defeated* by positive facts like $ab jerry \leftarrow \top$. The weak completion of a negative assumption $A \leftarrow \perp$ and its *defeater* $A \leftarrow \top$ is $A \leftrightarrow \top \vee \perp$, which is semantically equivalent to $A \leftrightarrow \top$. We believe that humans may defeat negative assumptions. To this end we allow defeaters of negative assumption to be abducibles in an extended abductive framework. Now, the observation *Jerry does not fly* can be explained by the minimal explanation $\{ab jerry \leftarrow \top\}$. In other words, *Jerry does not fly because it is an abnormal bird*.

Hypothesis 4: Humans may defeat negative assumptions.

Consider this extended scenario:

*Usually birds can fly, but kiwis and penguins cannot.
Tweety and Jerry are birds.*

This background information can be encoded by the program

⁴ We omit parenthesis if unary predicates are applied to variables or constants.

$$\{ \text{fly } X \leftarrow \text{bird } X \wedge \neg \text{ab } X, \\ \text{ab } X \leftarrow \text{kiwi } X, \quad \text{ab } X \leftarrow \text{penguin } X, \\ \text{bird } \text{tweety} \leftarrow \top, \quad \text{bird } \text{jerry} \leftarrow \top \quad \}.$$

The corresponding ground program is

$$\{ \text{fly } \text{tweety} \leftarrow \text{bird } \text{tweety} \wedge \neg \text{ab } \text{tweety}, \\ \text{fly } \text{jerry} \leftarrow \text{bird } \text{jerry} \wedge \neg \text{ab } \text{jerry}, \\ \text{ab } \text{tweety} \leftarrow \text{kiwi } \text{tweety}, \quad \text{ab } \text{tweety} \leftarrow \text{penguin } \text{tweety}, \\ \text{ab } \text{jerry} \leftarrow \text{kiwi } \text{jerry}, \quad \text{ab } \text{jerry} \leftarrow \text{penguin } \text{jerry}, \\ \text{bird } \text{tweety} \leftarrow \top, \quad \text{bird } \text{jerry} \leftarrow \top \quad \}.$$

Weakly completing this program we obtain

$$\{ \text{fly } \text{tweety} \leftrightarrow \text{bird } \text{tweety} \wedge \neg \text{ab } \text{tweety}, \\ \text{fly } \text{jerry} \leftrightarrow \text{bird } \text{jerry} \wedge \neg \text{ab } \text{jerry}, \\ \text{ab } \text{tweety} \leftarrow \text{kiwi } \text{tweety} \vee \text{penguin } \text{tweety}, \\ \text{ab } \text{jerry} \leftarrow \text{kiwi } \text{jerry} \vee \text{penguin } \text{jerry}, \\ \text{bird } \text{tweety} \leftrightarrow \top, \quad \text{bird } \text{jerry} \leftrightarrow \top \quad \}.$$

whose least model maps *bird jerry* and *bird tweety* to true and all other ground atoms to unknown.

Suppose we observe

Jerry flies.

This observation can be justified by the minimal explanation

$$\{ \text{kiwi } \text{jerry} \leftarrow \perp, \text{penguin } \text{jerry} \leftarrow \perp \}.$$

One should observe that both atoms, *kiwi jerry* and *penguin jerry*, must be mapped to false in order to map *ab jerry* to false as well, which is a prerequisite for mapping *fly jerry* to true. In other words, we need to assume that *Jerry is not a kiwi and not a penguin* in order to conclude that *Jerry flies*.

There are several problems with this approach. Firstly, to the best of our knowledge there are currently 41 known classes of birds which do not fly. Secondly, only specialists in biology might know these 41 classes. Thirdly, there may be classes of non-flying birds which we are unaware of. Hence, it is unlikely that humans consider all known exceptions before concluding *Jerry flies*.

In order to overcome these problems, the *Weak Completion Semantics* was extended with a new truth-functional operator *ctxt* (called *context*) in [10], whose meaning is specified in Table 1.1. The meaning of *ctxt* can be understood as a mapping from three-valuedness to two-valuedness, which allows to locally capture *negation as failure* [4] under *Weak Completion Semantics*. Using the *ctxt* operator we may now specify the program for the last Tweety example as

L	$ctxt L$
\top	\top
\perp	\perp
U	\perp

Table 1.1 The truth table for $ctxt L$, where L denotes a literal.

$$\{ \text{fly } X \leftarrow \text{bird } X \wedge \neg ab X, \\ \text{ab } X \leftarrow \text{ctxt kiwi } X, \quad \text{ab } X \leftarrow \text{ctxt penguin } X, \\ \text{bird tweety} \leftarrow \top, \quad \text{bird jerry} \leftarrow \top \}.$$

The corresponding ground program is

$$\{ \text{fly tweety} \leftarrow \text{bird tweety} \wedge \neg ab tweety, \\ \text{fly jerry} \leftarrow \text{bird jerry} \wedge \neg ab jerry, \\ \text{ab tweety} \leftarrow \text{ctxt kiwi tweety}, \quad \text{ab tweety} \leftarrow \text{ctxt penguin tweety}, \\ \text{ab jerry} \leftarrow \text{ctxt kiwi jerry}, \quad \text{ab jerry} \leftarrow \text{ctxt penguin jerry}, \\ \text{bird tweety} \leftarrow \top, \quad \text{bird jerry} \leftarrow \top \}.$$

Weakly completing the ground program we obtain

$$\{ \text{fly tweety} \leftrightarrow \text{bird tweety} \wedge \neg ab tweety, \\ \text{fly jerry} \leftrightarrow \text{bird jerry} \wedge \neg ab jerry, \\ \text{ab tweety} \leftrightarrow \text{ctxt kiwi tweety} \vee \text{ctxt penguin tweety}, \\ \text{ab jerry} \leftrightarrow \text{ctxt kiwi jerry} \vee \text{ctxt penguin jerry}, \\ \text{bird tweety} \leftrightarrow \top, \quad \text{bird jerry} \leftarrow \top \}.$$

whose least model maps *kiwi jerry*, *kiwi tweety*, *penguin jerry* and *penguin tweety* to unknown, *ab jerry* and *ab tweety* to false (thanks to the *ctxt* operator), and *bird jerry* and *bird tweety* as well as *fly jerry* and *fly tweety* to true. Reasoning with respect to this least model allows to conclude that *Jerry flies* by *default* without even considering the exceptional cases of birds which do not fly.

If we learn, for example, that

Jerry is a penguin,

then the situation changes again. Our background knowledge is extended to

$$\{ \text{fly } X \leftarrow \text{bird } X \wedge \neg ab X, \\ \text{ab } X \leftarrow \text{ctxt kiwi } X, \quad \text{ab } X \leftarrow \text{ctxt penguin } X, \\ \text{bird tweety} \leftarrow \top, \quad \text{bird jerry} \leftarrow \top \\ \text{penguin jerry} \leftarrow \top \}.$$

The least model of the weak completion of this program maps *kiwi jerry*, *kiwi tweety* and *penguin tweety* to unknown and *penguin jerry*, *bird jerry* and *bird tweety* to true.

Consequently, $ab\ jerry$ is mapped to true, whereas $ab\ tweety$ is still mapped to false. Hence, $fly\ tweety$ is mapped to true, whereas $fly\ jerry$ is mapped to false. Thus, we can conclude that *Tweety flies, but Jerry does not fly*.

Let us consider the following characterizations about kiwis and penguins:

If a bird has feathers like hair then it is likely to be a kiwi.

If a bird is black and white then it is likely to be a penguin.

Consequently, let us extend the above program with the following clauses:

$$\{ \text{kiwi } X \leftarrow \text{featherslikeHair } X \wedge \neg ab_k X, \quad ab_k X \leftarrow \perp, \\ \text{penguin } X \leftarrow \text{blackAndWhite } X \wedge \neg ab_p X, \quad ab_p X \leftarrow \perp \}.$$

Let us assume that we observe

Tweety has feathers like hair and does not fly.

The most plausible explanation in the context of knowing that *Tweety has feathers like hair*, is, that *Tweety is a kiwi*, and not, that *Tweety is a penguin*.

Hypothesis 5: Humans prefer certain explanations depending on the context.

In order to provide a preference among explanations, [10] provides an extension of abduction, which refines the notion of strong dependency among literals in a program, i.e. when no *ctxt* operator is involved in the dependency. Accordingly, only abducibles that strongly depend on the observation, are those that are allowed to serve as explanations. In the above example, $\text{penguin } tweety \leftarrow \top$ cannot serve as an explanation,

IT IS NOT EVEN AN ABDUCIBLE AND, HENCE, IT CAN NEVER BE AN EXPLANATION. CAN YOU EXPLAIN THIS BETTER?

because the observation that *Tweety does not fly*, does not strongly depend on *penguin tweety*. However the observation that *Tweety has feathers like hair*, can be explained by itself, which maps *kiwi tweety* to true and is at the same time an explanation for *Tweety does not fly*.

1.5 Obligation versus Factual Conditionals

This example is taken from [21] with minor modifications. Consider the background knowledge:

If it rains then the roofs are wet and she takes her umbrella.

In fact, these are two conditionals with the same antecedent which can be represented by the program

$$\{wet_roofs \leftarrow rain \wedge \neg ab_w, ab_w \leftarrow \perp, umbrella \leftarrow rain \wedge \neg ab_u, ab_u \leftarrow \perp\},$$

where *wet_roofs* denotes that *the roofs are wet*, *rain* denotes that *it rains*, *umbrella* denotes that *she takes her umbrella* and ab_w and ab_u are abnormality predicates. Weakly completing this program we obtain

$$\{wet_roofs \leftrightarrow rain \wedge \neg ab_w, ab_w \leftrightarrow \perp, umbrella \leftrightarrow rain \wedge \neg ab_u, ab_u \leftrightarrow \perp\},$$

whose least model maps ab_w and ab_u to false and all other atoms to unknown.

Given this background knowledge we would like to evaluate the four conditionals

If the roofs are not wet then it did not rain.

If she did not take her umbrella then it did not rain.

If the roofs are wet then it did rain.

If she took her umbrella then it did rain.

To this end we apply minimal revision followed by abduction as presented in Section 1.3. All conditionals are mapped to true. Because the antecedents of the conditionals are unknown given the least model of the weak completion of the background knowledge, abduction is applied yielding the minimal explanation $\{rain \leftarrow \perp\}$ in case of the first and second conditional and $\{rain \leftarrow \top\}$ in the third and fourth conditional. Hence, the consequences of the four conditionals will be true.

This is a bit surprising. Consider the first two conditionals, where the negative consequence of the background knowledge is abductively affirmed, i.e. the consequence is denied. It is not clear that given the information that *the roofs are not wet*, humans would conclude *it did not rain* in the same way as they would conclude *it did not rain* in case the given information was *she did not take her umbrella*. A similar difference holds regarding the third and fourth conditionals, where the consequent of the background knowledge is confirmed. Given the information that *she took her umbrella*, it is again not clear that humans would conclude *it did rain* in the same way as they would conclude *it did rain* given the information that *the roofs are wet*.

1.5.1 Obligation and Factual Conditionals

It appears that the two conditionals stated as background knowledge should be semantically interpreted in two different ways. Such a semantic interpretation will be developed in the remainder of this section. Consider the conditionals

if it rains then the roofs are wet

and

if it rains then she takes her umbrella.

The consequence of the first conditional is obligatory. We cannot easily imagine a case, where the antecedent is true and the consequence is not. On the other hand, we can easily imagine a situation, where the antecedent of the second conditional is true and the consequence is not. The consequence of the second conditional is not causally obligatory. We will call the first conditional an *obligation conditional* and the second one a *factual conditional*.

As explained in [2], a conditional whose consequence is denied is more likely to be evaluated to true if it is an obligation conditional. This happens because for this type of conditional there is a forbidden or unlikely possibility where antecedent and not consequent happen together and, in this case, where the consequent is known to be false, it cannot be the case that the antecedent is true as, otherwise, the forbidden possibility is violated. Thus, not antecedent is concluded. Because in the case of a factual conditional this forbidden possibility does not exist, a conditional whose consequence is denied should be evaluated as unknown.

Hypothesis 6: Humans may classify conditionals as obligation or factual conditionals.

This is an informal and pragmatic classification. It depends on the background knowledge and experience of a subject as well as on the context in which a conditional is stated.

One should observe that the conditional

Usually birds can fly

considered in Section 1.4 is usually classified as a factual conditional. For most humans, its consequence is not obligatory. They can imagine cases, where its antecedent *X is a bird* is true and its consequent *X can fly* is not.

1.5.2 Necessary and Sufficient Antecedents

The antecedent of the first conditional in the background knowledge is *necessary*. The consequent cannot be true unless the antecedent is true. The antecedent of the second conditional in the background knowledge does not appear to be necessary. There are many different reasons for taking an umbrella like, for example, that the sun is shining. The antecedent of the second conditional is *sufficient*.

Hypothesis 7: Humans may classify antecedents as necessary or sufficient.

The classification is informal and pragmatic. It depends on the background knowledge and experience of a subject as well as on the context in which the condition is stated.

1.5.3 Reasoning with Obligation and Factual Conditionals

Obligation and factual conditionals are represented by programs as before. Thus, the program representing the background knowledge remains unchanged. However, the semantics of conditionals is taken into account by modifying the set of abducibles for a given program. As usual, undefined relations may be abduced. In addition we allow to abduce expressions of the form $A \leftarrow \top$ if A is the head of a conditional with sufficient antecedent. The consequent A may be true even if the antecedent is not true. And we allow to abduce defeaters of abnormalities if the abnormality occurs in the body of a factual conditional. (One should observe that this is consistent with the use of defeaters in the examples discussed in Section 1.4.) In our example we obtain the set

$$\{rain \leftarrow \top, rain \leftarrow \perp, ab_u \leftarrow \top, umbrella \leftarrow \top\}$$

of abducibles.

Given this modified set of abducibles we may now evaluate the four conditionals using the principle *minimal revision followed by abduction*. Evaluating the conditional

if the roofs are not wet then it did not rain

yields the same result as before because the explanation $\{rain \leftarrow \perp\}$ is the only minimal explanation for the observation that the *roofs are not wet*. Hence, the conditional is true.

Evaluating the conditional

if she did not take her umbrella then it did not rain

we now find two minimal explanations for the observation that *she did not take her umbrella*, viz. $\{rain \leftarrow \perp\}$ and $\{ab_u \leftarrow \top\}$. If we add the first explanation to the background knowledge, then *rain* is mapped to false. However, if we add the second explanation to the background knowledge, then *rain* remains unknown. Reasoning skeptically the conditional is evaluated as unknown.

Evaluating the conditional

if the roofs are wet then it did rain

yields the same result as before because the explanation $\{rain \leftarrow \top\}$ is the only minimal explanation for the observation that the *roofs are wet*. Hence, the conditional is true.

Evaluating the conditional

if she did take her umbrella then it did rain

we now find two minimal explanations for the observation that *she did take her umbrella*, viz. $\{rain \leftarrow \top\}$ and $\{umbrella \leftarrow \top\}$. If we add the first explanation to the background knowledge, then *rain* is mapped to true. However, if we add

the second explanation to the background knowledge, then *rain* remains unknown. Reasoning skeptically the conditional is evaluated as unknown.

According to this classification, [21] can solve the abstract as well as the social case of the selection task: In the abstract case the conditional is interpreted as a factual conditional with necessary antecedent and in the social case the conditional is interpreted as an obligation conditional with sufficient antecedent.

1.6 Conclusions

We have shown that different human reasoning episodes can be adequately modeled with abductive reasoning. Interestingly, different forms of abduction seem to appear in the different cases. As exemplified by the suppression task, reasoning within conditionals requires skeptical abduction. On the other hand, when modeling reasoning within conditionals and counterfactuals, a form of minimal revision prior to abduction seems to occur. Finally, the last two cases show that the plausibility of explanations can be determined by the context and by the types of conditions and conditionals.

This leads us to two related observations: (1) From a complexity point of view, computing skeptical conclusions is quite expensive [14, 11], and in case the set of abducibles become larger, skeptical reasoning appears to be infeasible. (2) From a cognitive perspective, it does not seem adequate that humans consider all explanations as equally plausible, but rather, that they only consider a *relevant* subset *bounded* by the context, other observations and the type of conditional.

In order to investigate how these explanations have to be bounded, we will need to know more about how humans reason. In general, we need findings from other psychological experiments that can tell us whether the underlying assumptions of the *Weak Completion Semantics* are adequate. In particular, whether humans do *always* prefer skeptical over credulous abduction and how exactly do they apply *revision* when reasoning about counterfactuals.

From our perspective the following questions are desirably worth evaluating.

- *Do humans reason with multi-valued logics and, if they do, which multi-valued logic are they using? Can an answer 'I don't know' be qualified as a truth value assignment or is it a meta-remark?*
- *What do we have to tell humans such that they fully understand the background information?*
- *Do humans apply abduction and/or revision if the condition of a conditional is unknown and, if they apply both, do they prefer one over the other? Do they prefer skeptical over credulous abduction? Do they prefer minimal revisions?*
- *How important is the order in which multiple conditions of a conditional are considered?*
- *Do humans consider abduction and/or revision steps that turn an indicative conditional into a subjunctive one?*

We surmise that humans do reason with a third truth value; we have shown that the *Suppression* and the *Selection Tasks* can be adequately modeled under *Weak Completion Semantics* and, moreover, in these tasks skeptical abduction had to be applied [6]. However, we also believe that they only certain explanations as plausible. These are determined by the context, the additional observations and the type of conditional in consideration. Finally, we venture that minimal revision followed by abduction are applied if the conditions of a conditional are *unknown*.

Acknowledgments

Luís Moniz Pereira acknowledges support from FCT/MEC NOVA LINCS PEst UID/CEC/04516/2019.

References

1. Ernest W. Adams. Subjunctive and indicative conditionals. *Foundations of Language*, 6(1):89–94, 1970.
2. R. M. J. Byrne. *The Rational Imagination: How People Create Alternatives to Reality*. MIT Press, Cambridge, MA, USA, 2007.
3. Ruth M.J. Byrne. Suppressing valid inferences with conditionals. 31:61–83, 1989.
4. Keith Leonard Clark. Negation as failure. In H. Gallaire and J. Minker, editors, *Logic and Data Bases*, volume 1, pages 293–322. Plenum Press, New York, NY, 1978.
5. Ana Costa, E.-A. Dietz Saldanha, Steffen Hölldobler, and Marco Ragni. A computational logic approach to human syllogistic reasoning. In G. Gunzelmann, A. Howes, T. Tenbrink, and E.J. Davelaar, editors, *Proceedings of the 39th Annual Conference of the Cognitive Science Society, (CogSci 2017)*, pages 883–888. Austin, TX: Cognitive Science Society, 2017.
6. E.-A. Dietz, S. Hölldobler, and M. Ragni. A computational logic approach to the suppression task. In N. Miyake, D. Peebles, and R. P. Cooper, editors, *Proceedings of the 34th Annual Conference of the Cognitive Science Society*, pages 1500–1505. Cognitive Science Society, 2012.
7. E.-A. Dietz, Steffen Hölldobler, and Raphael Höps. A computational logic approach to human spatial reasoning. In *IEEE Symposium Series on Computational Intelligence, (SSCI 2015)*, pages 1627–1634. IEEE, 2015.
8. Emmanuelle-Anna Dietz, Steffen Hölldobler, and Luís Moniz Pereira. On conditionals. In G. Gottlob, G. Sutcliffe, and A. Voronkov, editors, *Global Conference on Artificial Intelligence (GCAI)*, Epic Series in Computing. EasyChair, 2015.
9. Emmanuelle-Anna Dietz, Steffen Hölldobler, and Marco Ragni. A computational logic approach to the abstract and the social case of the selection task. In *Proceedings of the 11th International Symposium on Logical Formalizations of Commonsense Reasoning, (COMMON-SENSE)*, 2013.
10. Emmanuelle-Anna Dietz Saldanha, Steffen Hölldobler, and Luís Moniz Pereira. Contextual reasoning: Usually birds can abductively fly. In *Proceedings of 14th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR)*, pages 64–77. Springer International Publishing, 2017.
11. Emmanuelle-Anna Dietz Saldanha, Steffen Hölldobler, and Tobias Philipp. Contextual abduction and its complexity issues. In *Proceedings of the 4th International Workshop on Defeasible and Ampliative Reasoning (DARe) co-located with the 14th International Conference on*

- Logic Programming and Nonmonotonic Reasoning (LPNMR)*, volume 1872 of *CEUR Workshop Proceedings*, pages 58–70. CEUR-WS.org, 2017.
12. Steffen Hölldobler. *Logik und Logikprogrammierung 1: Grundlagen*. Kolleg Synchron. Synchron, 2009.
 13. Steffen Hölldobler and Carroline Dewi Puspa Kencana Ramli. Logics and networks for human reasoning. In Cesare Alippi, Marios M. Polycarpou, Christos G. Panayiotou, and Georgios Ellinas, editors, *International Conference on Artificial Neural Networks (ICANN), Part II*, volume 5769 of *Lecture Notes in Computer Science*, pages 85–94, Heidelberg, 2009. Springer Verlag.
 14. Steffen Hölldobler, Tobias Philipp, and Christoph Wernhard. An abductive model for human reasoning. In *Logical Formalizations of Commonsense Reasoning (Commonsense 2011)*, AAAI Spring Symposium Series Technical Reports, pages 135–138, Cambridge, MA, 2011.
 15. Antonis Kakas, Robert Kowalski, and Francesca Toni. Abductive logic programming. 2(6):719–770, 1993.
 16. Stephen Cole Kleene. *Introduction to Metamathematics*. North-Holland, Amsterdam, 1952.
 17. John Wylie Lloyd. *Foundations of Logic Programming*. Springer-Verlag New York, Inc., New York, NY, USA, 1984.
 18. Jan Łukasiewicz. O logice trójwartościowej. *Ruch Filozoficzny*, 5:169–171, 1920. English translation: On three-valued logic. In: Łukasiewicz J. and Borkowski L. (ed.). (1990). *Selected Works*, Amsterdam: North Holland, pages 87–88.
 19. J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York, USA, 2000.
 20. Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1-2):81–132, 1980.
 21. Emmanuelle-Anna Dietz Saldanha, Steffen Hölldobler, and Isabelly Lourêdo Rocha. Obligation versus factual conditionals under the weak completion semantics. In Steffen Hölldobler, Andrey Malikov, and Christoph Wernhard, editors, *YSIP2 - Proceedings of the Second Young Scientist's International Workshop on Trends in Information Processing, Dombai, Russian Federation, May 16-20, 2017.*, volume 1837 of *CEUR Workshop Proceedings*, pages 55–64. CEUR-WS.org, 2017.
 22. Keith Stenning and Michiel van Lambalgen. Semantic interpretation as computation in non-monotonic logic: The real meaning of the suppression task. *Cognitive Science*, 6(29):916–960, 2005.
 23. Keith Stenning and Michiel van Lambalgen. *Human Reasoning and Cognitive Science*. A Bradford Book. 2008.