

A review of proposed principles of causal non-monotonic reasoning

Patrick Marchisella
School of Environment and Science
Griffith University
Brisbane, Australia

Abstract

Within Non-monotonic Reasoning, numerous principles of causal reasoning have been proposed. Many of these principles have been viewed as desirable in formalisms that reason with causality, and have been widely adopted throughout the literature. We provide a critique of these principles, evaluate their suitability for characterising and formulating causal non-monotonic reasoning, and find that most are unsuitable. Further, we discuss a new approach to causal non-monotonic reasoning motivated by how humans typically reason with causality.

1 Introduction

In day-to-day life, we often make conclusions from information that is incomplete or vague, at no point quantify any vagueness or uncertainty in the information, and retract or change our conclusions (if necessary) in light of new information. As an example, consider the following simple scenario: Bruce is a dog, and dogs usually like to play fetch. From this information, we naturally conclude that Bruce probably likes to play fetch. Suppose we then add a new sentence to the scenario: Bruce does not like to play fetch. We retract our previous conclusion, and instead conclude that Bruce certainly does not like to play fetch. In doing so, we have reasoned *non-monotonically*: retracted our earlier conclusion and made a different conclusion due to new information, without quantifying any uncertainty.

Non-monotonic Reasoning (NMR) is an attempt to formally capture *at least* the non-monotonic reasoning capability native to humans. It is an important area of Artificial Intelligence, in particular of Knowledge Representation, and is

motivated by the ability to endow intelligent systems with the capability to reason non-monotonically using no quantification of uncertainty [29].

For some time, understanding cause and effect has been a goal of various sciences. For example: What causes erroneous cell mitosis? What causes atomic decay? It has also been a goal of philosophy. For example: What causes humans to think about causality?

Formalising how humans reason with causality is known to be a very difficult problem [34, 18, 22, 23, 19, 35, 30, 12]. Firstly, we often don't know the cause of things with certainty. So we need to permit uncertainty. For example: turning the key to "ON" causes the car to start. But what if the battery is dead? Or not connected? Or the starter motor is broken? Secondly, many familiar logical and mathematical tools cannot be used to correctly represent causality: material implication, reflexivity, symmetry, transitivity, and others. Hence formalising how humans reason with causality quickly becomes a non-trivial task.

Researchers within NMR have noticed that humans very often reason non-monotonically with cause and effect [31, 20, 23, 19, 35]. As an example, consider the following simple scenario.

Example 1.1

1. Fred is driving recklessly.
2. Reckless driving usually causes accidents.
3. Fred is later found involved in an accident.

We naturally conclude that Fred's reckless driving likely caused the accident. Suppose we then add a new sentence to the scenario, as follows.

1. Fred is driving recklessly.
2. Reckless driving usually causes accidents.
3. Fred is later found involved in an accident.
4. Fred was rear-ended while he was stopped at a red light.

We naturally retract our previous conclusion, and instead conclude that the accident was caused by a rear-end collision (and not by Fred's reckless driving).

Similarly to how we reasoned non-monotonically about Bruce playing fetch, in Example 1.1 we reasoned non-monotonically about the cause of the accident. From this simple example we see that NMR is well suited to the task of formalising how humans reason with causality in day-to-day life.

Lists of principles or postulates for certain types of reasoning are helpful for many reasons. As noted in [2], some of those reasons are the following: they provide characterisations for certain types of reasoning, they facilitate the evaluation of formalisms that claim to do a certain type of reasoning, they provide guidelines that formalisms can follow in order to do a certain type of reasoning, and they

can expose differences (which might not have been previously recognised) between certain types of reasoning. So lists of principles provide a valuable contribution to the development of formal reasoning systems. When several lists of postulates or principles have been proposed for a particular type of reasoning, it is often very helpful to review those lists and see which principles are commonly agreed upon. Such principles can often be seen as ideas that characterise a certain type of reasoning in a general sense.

Our aim in this paper is twofold. Firstly, we present a list of principles of causal non-monotonic reasoning that have been commonly adopted within Non-monotonic Reasoning (NMR). In doing so, we also provide a survey of the current literature on causal non-monotonic reasoning. Secondly, we determine whether or not these principles sensibly characterise causal non-monotonic reasoning in a general sense; that is, causal non-monotonic reasoning as displayed by humans in day-to-day life. Interestingly, the second of our aims presents us with the following problem. The principles suggested in [38], [3], and [12] are, in each case, stated only in the context of the logics (or mathematical relations, in the case of [3]) proposed in those papers. This is also the case in [30], but to a lesser degree: in that paper, some of the principles are stated in the context of propositional logic, whereas others are context free. The problem arises in that for some of these context-bound principles, it is difficult to determine their meaning when they are removed from their particular contexts. For example: some of the axiomatic postulates proposed in [12] concern how a particular modal operator should be used and understood in the context of the conditional logic proposed in that paper. So when it comes to a principle concerning the usage and meaning of a modal operator in a conditional logic, how are we to understand this principle in a general sense? Can it be sensibly translated? If so, how?

When deciding which principles or postulates sensibly characterise general causal non-monotonic reasoning, we would like our decision to be as well-informed as possible. Hence for any proposed principle that is context-bound, instead of disregarding that principle, we will try to find the motivation or idea that underpins it. If we can, then we will evaluate that motivation or idea for general causal non-monotonic reasoning. If not, then we will not consider that principle for now.

2 Proposed principles of causal non-monotonic reasoning

Here we present a list of proposed principles of causal non-monotonic reasoning that have been commonly adopted within NMR, and evaluate their suitability for characterising causal non-monotonic reasoning in a general sense. Our viewpoint

here will be an intuitive one: we wish to review and evaluate the proposed principles in as general a manner as possible, and hence we will not assume any specific context or domain of causality. Interestingly, it appears as though most of the formalisms within NMR that reason with causality have not taken this viewpoint: rather, each appears to have been motivated firstly by a desire to use familiar mathematics, and secondly by an interest in dealing with causality only within a specific domain. We comment further on this in Section 3.

In what follows, we will use the language of classical propositional logic. In particular, let a, b, c, f , and g be propositional logic formulas, and \mathfrak{t} be any tautology. We will also need to refer to reasoning scenarios that contain knowledge of the form “ f causes g ”. We call such knowledge *causal information*.

2.1 Universal causality

The Principle of Universal Causality states that anything that is true must have a cause. It has been present in western Philosophy for some time, with its first appearance arguably occurring in Plato’s *Timaeus* [6]. Since the time of post-Aristotelian Philosophy, the principle has been the subject of some debate - a discussion of notable arguments for and against the idea of universal causality is given in [21]. Indeed, philosophers have been suspicious of the principle for more than two-hundred years, with doubts first being cast by Hume in his important work of 1739 titled *A Treatise of Human Nature* [14].

Within NMR, there seems to be common agreement that the principle should be obeyed by formalisms that reason with causality [31, 23, 35, 5, 3, 13, 4, 36]. In many cases, the agreement appears to rest upon the observation that adoption of the principle allows the use of familiar mathematics - namely a fixed-point semantics. We show that we do not want to adopt the Principle of Universal Causality.

Firstly, there exist physical phenomena for which a cause is unknown, and seems impossible to determine. One such example is radioactive decay. The cause of radioactive decay is known - an unstable atomic nucleus will emit nucleons or elementary particles in order to reach the lowest possible energy state. However, what causes the emission to occur at some particular instant of time is unknown, and research suggests that the event is entirely stochastic [27, 33, 1]. Secondly, adopting the Principle of Universal Causality presents the following problems. Consider any tautology. Regardless of whether or not the tautology has a cause, it is always true. That is, we do not need a tautology to have a cause in order to accept its truth. Now consider any entity. If we adopt the Principle of Universal Causality, then the entity has a cause. Also, the cause of the entity has a cause, and the cause of the cause of the entity has a cause, and so on until we end up with a chain of cause and effect. So where does the chain end? Either it has no

end, and hence is infinitely long, or it is circular. In both cases, it is very difficult to see how we might make sense of such a chain.

When attempting to formalise causal non-monotonic reasoning, we certainly would not like to adopt a principle that is defied by an observable physical phenomenon, and which presents problems concerning tautologies and infinite causal regression. Hence we do not think that the Principle of Universal Causality should be a principle of causal non-monotonic reasoning.

2.2 Non-material implication

The Principle of Non-Material Implication states that material implication must not be used to represent causality. The principle rests on the observation that material implication cannot always correctly represent causal information, and hence should not be used to do so. As a simple example, consider the sentence “If I go to the beach then the weather is fine”. Whilst being a harmless instance of material implication, my presence at the beach certainly does not cause the weather to be fine, and hence the sentence makes very little sense if we interpret it as causal information. More examples are given in [31] and [30], where it is shown that contraposition and conjunctive antecedents are also problematic when attempting to use material implication to represent causal information. Hence we agree that the Principle of Non-material Implication should be a principle of causal non-monotonic reasoning.

2.3 Non-monotonicity

The Principle of Non-monotonicity states that if a formalism claims to reason with cause and effect, then the formalism must reason non-monotonically. This principle is motivated by the observation that humans often reason non-monotonically with cause and effect [31, 20, 19, 23, 35]. This was shown in Example 1.1 above, and hence we adopt the Principle of Non-monotonicity.

2.4 Context sensitivity

The Principle of Context Sensitivity states that when reasoning with cause and effect, we do not need to know everything about the scenario at hand. Instead, we need to know only what is usually, or typically true in that scenario. The motivation behind this principle is the simple observation that we often make conclusions about cause and effect based on information that is incomplete or vague. This was shown in Example 1.1 above, and hence we adopt the Principle of Context Sensitivity.

Interestingly, the same observation, but for more general reasoning (not restricted to cause and effect), underpins an active area of research within Philosophy known as “Defeasible Reasoning”. Within NMR, there have been some very well-known and popular attempts at formalising Defeasible Reasoning, such as Default Logic [28], Circumscription [24], Autoepistemic Logic [25], and Defeasible Logic [26]. Indeed, formalising Defeasible Reasoning is one of the underpinning ideals of NMR. So it appears as though The Principle of Context Sensitivity is simply another name for Defeasible Reasoning.

2.5 Temporality

The Principle of Temporality states that objects or entities involved in causal reasoning must have some kind of temporal aspect. That is, the objects or events should be indescribable without relying on at least a reference to time. It is proposed in [31], and obeyed by many formalisms within NMR that reason with causality [8, 31, 19, 23, 35, 13, 36]. These formalisms appear to be motivated by dealing with causality only in the domain of “actions and change” - a task which the authors believe is well suited to the adoption of temporality.

Causality occurs naturally in many contexts. In particular, it occurs in contexts that do not require temporality in order to be understood. For example: the sentence “Thoughts of home cause Fred to feel homesick” describes a causal relationship that does not require the involvement of time in order to be understood.

Whilst temporality may be well suited to reasoning with causality in the domain of actions and change, we do not want to force domain-specific principles upon all formalisms that reason non-monotonically with causality. Thus we do not think that the Principle of Temporality should be a principle of causal non-monotonic reasoning.

2.6 Irreflexivity

The Principle of Irreflexivity states that when representing and reasoning with causal information, irreflexivity should hold. That is, we should not have a causes a . It is proposed in [31], and obeyed by many formalisms within NMR that reason with causality [31, 20, 9, 10, 5, 12, 3, 4].

Causality naturally brings to mind the idea of what causes what. That is, we naturally understand the sentence “ a causes b ” to mean that b is brought about by the occurrence of a (or that b occurs due to the occurrence of a). To determine whether or not the Principle of Irreflexivity might be sensibly adopted when reasoning with causality, we can consider the question “Can something bring about itself?” An example of such a thing certainly does not spring to mind easily. Indeed, dismissal of the idea of something causing itself can be found in the

theoretical foundations of modern philosophy: Immanuel Kant, in his well known 1781 thesis titled *A Critique of Pure Reason*, strongly refutes the possibility of an entity causing itself [32]. Surprisingly, some authors in modern Philosophy believe that this idea is still worthy of debate [37].

We are not comfortable permitting formalisms to reason with causal information that is reflexive. Hence we think that the Principle of Irreflexivity should be a principle of causal non-monotonic reasoning.

2.7 Antisymmetry

The Principle of Antisymmetry states that when representing and reasoning with causality, antisymmetry should hold. It is proposed in [31], and motivated by examples in which symmetry does not make sense. One such example involves the sentences “Smoking caused John’s death”, and “John’s death caused smoking.” Clearly the former makes sense, while the latter doesn’t.

Examples like this might appear to be enough to persuade us to adopt the Principle of Antisymmetry. But there exists examples in which symmetry is not problematic. Consider the sentence “Smoking cigarettes usually causes a nicotine addiction”. Applying symmetry here, we get “A nicotine addiction usually causes smoking”, which makes sense. So there exists examples that support this principle and examples that oppose it. In deciding which principles to adopt when characterising causal non-monotonic reasoning, we do not want to adopt any principle which can be shown to be questionable. Hence we do not adopt the Principle of Antisymmetry.

2.8 Transitivity

The Principle of Transitivity, in simplest form, states that from a chain of sentences such as “ a causes b ” and “ b causes c ”, “ a causes c ” should follow. It is discussed in [30], and obeyed by many formalisms within NMR that reason with causality [20, 23, 38, 5, 3, 4]. We show that we do not want to adopt this principle.

In the NMR literature on reasoning with causality, there are both examples that support this principle and examples that oppose it [30]. A simple example that supports transitivity involves a sequence of dominoes: when pushed, the first domino causes the second domino to fall, which in turn causes the third domino to fall, and so on until the final domino falls. Here it is reasonable to say that pushing the first domino causes, via transitivity, the final domino to fall. On the other hand, an example in which transitivity is not sensible is the following: thinking about causality causes Fred to become tired, tiredness causes Fred to be unable to concentrate, and the inability to concentrate causes Fred to be unable to

think about causality. Applying transitivity, we get that thinking about causality causes Fred to be unable to think about causality, which does not make sense.

However, here we make an important observation concerning a relationship between transitivity and temporality. Assuming a temporal displacement between each of the events in the latter example above, we get the following: thinking about causality at time t causes Fred to become tired at time $t + 1$, tiredness at time $t + 1$ causes Fred to be unable to concentrate at time $t + 2$, and the inability to concentrate at time $t + 2$ causes Fred to be unable to think about causality at time $t + 3$. Applying transitivity, we get that thinking about causality at time t causes Fred to be unable to think about causality at time $t + 3$. That is, thinking about causality at some point in time causes Fred to be unable to think about causality at a later point in time, which seems sensible.

Importantly, we note that the introduction of temporal displacement to our example allowed us to use transitivity without reaching a counter-intuitive conclusion. So it appears as though in order to avoid undesirable conclusions when using transitivity, some kind of temporality is required. Interestingly, this observation appears to be supported by the fact that many of the formalisms within NMR that reason with causality and adopt transitivity are also temporal.

We have already shown that we do not want to force temporality on formalisms that reason non-monotonically with causality. Hence we are not comfortable forcing transitivity upon the same formalisms, since we would then risk those formalisms making conclusions that are counter-intuitive. Thus we do not think that the Principle of Transitivity should be a principle of causal non-monotonic reasoning.

2.9 Causal conjunction

The Principle of Causal Conjunction states that from “ a causes b ” and “ a causes c ”, “ a causes $b \wedge c$ ” should follow. It is proposed in [30], and obeyed by many formalisms within NMR that reason with causality [20, 23, 35, 9, 10, 38, 5, 12, 3, 13, 17, 4, 7, 15, 16].

In each of the formalisms mentioned so far that deal with causality, only causal information that is certain is represented and reasoned with. That is, the information does not contain any inherent uncertainty. But what if our scenario contains causal information that is inherently uncertain? The following example shows that this type of information can behave very differently to causal information that is certain. Consider a switch that causes a light bulb to illuminate in exactly one of three colours: red, green, or blue. Then for each element in $\{red, green, blue\}$, it is unlikely that the switch will cause the bulb to illuminate in that colour. Also, if c_1 and c_2 are different elements of $\{red, green, blue\}$, then it is likely that the switch will cause the bulb to illuminate in c_1 or c_2 .

Now, it is likely that the switch will cause the lightbulb to illuminate in a colour that is not red. It is also likely that the switch will cause the lightbulb to illuminate in a colour that is not green. This information would be enough for a formalism that obeys the Principle of Causal Conjunction to conclude that a colour which is not red and not green is likely. However, since the lightbulb must illuminate in exactly one of the three colours, the likelihood of the switch causing the lightbulb to illuminate in a colour which is not red and not green is the same as the likelihood of the switch causing the lightbulb to illuminate in blue. But it is unlikely that the switch will cause the lightbulb to illuminate in blue. So although a colour which is not red is likely and a colour which is not green is likely, a colour which is not red and not green is unlikely.

When asking a formalism to reason with this scenario, we would not want that formalism making conclusions that are unlikely. Hence we do not think that the Principle of Causal Conjunction should be a principle of causal non-monotonic reasoning.

2.10 The “Or” rule

The “Or” rule states that from “ a causes c ” and “ b causes c ”, “ $a \vee b$ causes c ” should follow. Like the Principle of Causal Conjunction, it is proposed in [30] and obeyed by many formalisms within NMR that reason with causality [20, 23, 9, 10, 38, 5, 12, 3, 4].

We again show that in scenarios that contain causal information and inherent uncertainty, the uncertainty can lead to the causal information behaving in an unexpected manner. We will begin with the 7-lottery example given in [2]. Consider a lottery that contains exactly seven tickets. Let t_i be the i 'th ticket, and $T_7 = \{t_1, t_2, \dots, t_7\}$ be the set of tickets in the lottery. Now, exactly one ticket in T_7 will be the winning ticket; that is, the outcome of the lottery must satisfy $\vee T_7$. Also, since only one ticket can win the lottery, we have for each i and j such that $1 \leq i < j \leq 7$, the outcome of the lottery must satisfy $\neg \wedge \{t_i, t_j\}$. Let L_7 be the formula that characterises the 7-lottery. Then L_7 is the conjunction of $\vee T_7$ and the formulas $\neg \wedge \{t_i, t_j\}$ such that $1 \leq i < j \leq 7$.

Suppose that the outcome of the lottery being t_5 , t_6 , or t_7 causes Fred to be happy. Consider the formulas $f = \wedge \{L_7, \neg t_1, \neg t_2\}$, $g = \wedge \{L_7, \neg t_3, \neg t_4\}$, and $h = \vee \{t_5, t_6, t_7\}$. So f restricts the winning ticket to being exactly one of t_3 , t_4 , t_5 , t_6 , or t_7 , while g restricts the winning ticket to being exactly one of t_1 , t_2 , t_5 , t_6 , or t_7 . Now if f is true, then it is likely that the winning ticket will be t_5 , t_6 or t_7 ; that is, it is likely that h will be true. This is also the case for the formula g : if g is true then it is likely that h will be true. Hence it is reasonable to say that if f is true then it is likely that Fred will be caused to be happy, and if g is true then it is likely that Fred will be caused to be happy. This information would

be enough for a formalism that adopts the “Or” rule to conclude that if $f \vee g$ is true, then it is likely that Fred will be caused to be happy. But $f \vee g$ places no restriction on the winning ticket, and the winning ticket being t_5 , t_6 , or t_7 is not a likely outcome of a 7-lottery. So it is not the case that if $f \vee g$ is true then h is likely to be true. Hence it is not correct to say that if $f \vee g$ is true then it is likely that Fred will be caused to happy.

When asking a formalism to reason with this scenario, we would not want that formalism making conclusions that are incorrect. Hence we do not think that the “Or” rule should be a principle of causal non-monotonic reasoning.

2.11 Right weakening

The Principle of Right Weakening states that if a causes b and $b \models c$, then “ a causes c ” should follow. Along with the Principle of Causal Conjunction and the “Or” rule, it is proposed in [30] and obeyed by many formalisms within NMR that reason with causality [23, 9, 10, 38, 5, 12, 3, 4]. We show that we do not want to adopt this principle.

When dealing with causality, tautologies, together with the Principle of Right Weakening, can lead to counter-intuitive results. To see this, let *RecklessDriving* and *Accident* be propositional logic atoms representing reckless driving and a road accident, respectively. Suppose that in a scenario S , *Accident* is caused by *RecklessDriving*. Suppose also that we adopt the Principle of Right Weakening. Now since \top is a tautology, we have *Accident* $\models \top$. Hence by the Principle of Right Weakening, we should have that within S , \top is also caused by *RecklessDriving*. But \top can be any tautology. For example, \top could represent the sentence “The grass is green or not green”. Then the Principle of Right Weakening would lead us to conclude that within S , reckless driving caused the grass to be green or not green. Clearly, when it comes to reasoning about cause and effect, this conclusion is not sensible regardless of what scenario is being reasoned with.

When asking a formalism to make conclusions about cause and effect, we would not want that formalism making conclusions that are not sensible. Hence we do not adopt the Principle of Right Weakening.

2.12 Left logical equivalence

The Principle of Left Logical Equivalence states that if a causes b and $a \equiv c$, then “ c causes b ” should follow. Along with the Principle of Causal Conjunction, the “Or” rule, and the Principle of Right Weakening, it is proposed in [30] and obeyed by many formalisms within NMR that reason with causality [20, 23, 9, 35, 10, 38, 12, 36]. We think it is reasonable to say that logically equivalent causes lead to the

same effect. Hence we agree that the Principle of Left Logical Equivalence should be a principle of causal non-monotonic reasoning.

2.13 Strengthening

The Principle of Strengthening states that if $a \models b$ and b causes c , then “ a causes c ” should follow. It is proposed in [38] and adopted in [38, 5, 3, 4]. We show that we do not want to adopt this principle.

It is not difficult to see that this principle leads to counter-intuitive results. For example, consider the sentences “ a causes b ”, and “ c prevents b ”. Then by strengthening, we may conclude “ $a \wedge c$ causes b ”. Clearly, this conclusion is not sensible.

As was the case for the Principle of Right Weakening, when asking a formalism to make conclusions about cause and effect, we would not want that formalism making conclusions that are not sensible. Hence we do not think that the Principle of Strengthening should be a principle of causal non-monotonic reasoning.

2.14 Cut

The Principle of Cut states that if a causes b and $a \wedge b$ causes c , then “ a causes c ” should follow. It is proposed in [5], and adopted in [5], [3] and [4] (it is also shown to be a property of the formalism proposed in [38]). We show that we do not want to adopt this principle.

As was the case for the Principle of Transitivity, in the absence of temporality the Principle of Cut can lead to counter-intuitive conclusions. To see this, consider the following example: thinking causes Fred to become tired; thinking and tiredness causes Fred to be unable to think clearly. Applying Cut, we get that thinking causes Fred to be unable to think clearly, which does not make sense.

However, if we assume a temporal displacement between each of the above events, we get the following: thinking at time t causes Fred to become tired at time $t + 1$; thinking at time t and becoming tired at time $t + 1$ causes Fred to be unable to think clearly at time $t + 2$. Applying Cut, we get that thinking at time t causes Fred to be unable to think clearly at time $t + 2$. That is, thinking at some point in time causes Fred to be unable to think clearly at a later point in time, which seems reasonable.

Importantly, we note that the introduction of temporal displacement to our example allowed us to use Cut without reaching a counter-intuitive conclusion. So it appears as though in order to avoid undesirable conclusions when using Cut, some kind of temporality is required.

We have already shown that we do not want to force temporality on formalisms that reason non-monotonically with causality. Hence we are not comfortable forc-

ing Cut upon the same formalisms, since we would then risk those formalisms making conclusions that are counter-intuitive. Thus we do not adopt the Principle of Cut.

2.15 Arbitrary conjunction

The Principle of Arbitrary Conjunction states that if a causes b , then for any propositional logic formula c , “ $a \wedge c$ causes b ” should follow. It is proposed and adopted in [38], where it is referred to as “The Law of AND”. It is not difficult to see that when it comes to reasoning with what causes what, this principle is not sensible. For example, consider the sentence “An approaching person causes the door to automatically open”. Suppose we adopt the Principle of Arbitrary Conjunction. Suppose we then learn that the door’s automatic opening device is not working. Then the Principle of Arbitrary Conjunction lets us conclude that an approaching person, together with the fact that the door’s automatic opening device is not working, causes the door to automatically open. Clearly, when it comes to reasoning with what causes what, this is not a sensible conclusion. Hence we do not adopt the Principle of Arbitrary Conjunction.

2.16 Arbitrary causation

The Principle of Arbitrary Causation states that propositional logic tautologies can be caused by any proposition. It is proposed in [38], and adopted in [38, 12]. It is not difficult to see that when it comes to reasoning with causality, this principle is not sensible. For example, let the propositional logic atoms d and g represent the sentences “Fred is fond of dogs” and “The grass is green”, respectively. Suppose we adopt the Principle of Arbitrary Causation. Then since $g \vee \neg g$ is a propositional logic tautology, we can choose d to be the cause of $g \vee \neg g$. That is, we may conclude that Fred’s fondness of dogs causes the grass to be green or the grass to be not green. Clearly, such a conclusion is not sensible. Hence we do not adopt the Principle of Arbitrary Causation.

As was the case for the Principle of Right Weakening and the Principle of Strengthening, when asking a formalism to make conclusions about cause and effect, we would not want that formalism making conclusions that are not sensible. Hence we do not think that the Principle of Arbitrary Causation should be a principle of causal non-monotonic reasoning.

2.17 Weakened monotonicity

The Principle of Weakened Monotonicity states that if a causes b , and it is not the case that a causes $\neg c$, then “ $a \wedge c$ causes b ” should follow. It is proposed in [38],

and adopted in [38, 11, 12]. It is not difficult to see that this principle leads to counter-intuitive results. For example, let a , b and c represent the sentences “Fred spends one hour in the sun”, “Fred gets sunburnt”, and “Fred applies a sufficient amount of sunscreen before spending one hour in the sun”, respectively. Now, it is reasonable to say that a causes b . It is also reasonable to say that a does not cause $\neg c$. Then by weakened monotonicity, we may conclude “ $a \wedge c$ causes b ”; that is, Fred spending one hour in the sun and applying a sufficient amount of sunscreen before spending one hour in the sun causes Fred to get sunburnt. Clearly, this conclusion is not sensible.

As was the case for the Principle of Right Weakening, the Principle of Strengthening, and the Principle of Arbitrary Causation, when asking a formalism to make conclusions about cause and effect, we would not want that formalism making conclusions that are not sensible. Hence we do not think that the Principle of Weakened Monotonicity should be a principle of causal non-monotonic reasoning.

3 Summary and conclusion

In this paper, we have reviewed proposed principles of causal non-monotonic reasoning that have been commonly adopted within NMR, and evaluated their suitability for characterising causal non-monotonic reasoning in a general sense. Of those principles, Non-material Implication, Non-monotonicity, Context Sensitivity, Irreflexivity, and Left Logical Equivalence appear to be suitable. On the other hand, although Universal Causality, Temporality, Antisymmetry, Transitivity, Causal Conjunction, the “Or” rule, Right Weakening, Strengthening, Cut, Arbitrary Conjunction, Arbitrary Causation and Weakened Monotonicity have been proposed and very often adopted in the literature, we have shown via simple examples that each is either unnecessary or not suitable for characterising causal non-monotonic reasoning. In particular: Causal Conjunction, the “Or” rule, Right Weakening, Strengthening, Arbitrary Conjunction, Arbitrary Causation and Weakened Monotonicity can each result in formalisms giving counter-intuitive or incorrect answers to simple reasoning scenarios.

When asking a formalism to reason with causal information, we would not want the formalism giving counter-intuitive or incorrect answers to simple reasoning scenarios. One reason for this is that it would make us justifiably suspicious of the formalism’s ability to reason sensibly with “real world” scenarios, which are typically much more complicated. It is here that we make an important observation about the causal non-monotonic formalisms that have been proposed so far in the literature: it appears as though each has been motivated primarily by the re-use of familiar mathematics, and less by a desire to formalise the way humans typically reason about cause and effect. For instance: amongst the formalisms mentioned

in Section 2, we find the usage of well-known mathematics such as fixed-point semantics, Circumscription [24], and conditional logics. Unfortunately, each of the same formalisms adopts at least one of the principles that we have shown to be either unnecessary or not suitable for characterising causal non-monotonic reasoning.

The above observation exposes a significant gap in the work that has been done on formalising causal non-monotonic reasoning within NMR so far. What is needed is an approach to causal non-monotonic reasoning whose primary motivation is the formal representation of the way in which humans typically reason with cause and effect. Such an approach would focus firstly on the accurate representation of how humans reason with causality in day-to-day life, and secondly on the selection or construction of suitable mathematics to facilitate formalising this reasoning. For instance: such an approach could begin with the simple observation that humans very naturally make conclusions about cause and effect based on information that is incomplete or vague. Such information contains inherent uncertainty, which is often not quantified, and is instead represented by the use of words such as “usually”, “typically”, or “likely”. Future work could then involve the construction of a formalism that correctly captures this reasoning, which could be guided by the review of proposed principles that we have provided in this paper. Such a formalism would initiate a new approach to the formalisation of causal non-monotonic reasoning within NMR: one that is more closely aligned to the way in which humans reason about cause and effect. In this respect, the formalism would be a highly valuable contribution to ongoing work in the area.

Acknowledgements

The author wishes to thank David Billington for helpful discussions regarding the topic and content of this paper, and in particular for very helpful suggestions concerning how certain thoughts might be communicated in a way that is understandable by humans.

References

- [1] L. Best, G. Rodrigues, and V. Velker. *Radiation Oncology Primer and Review*. Demos Medical Publishing, New York, NY, USA, 2013.
- [2] D. Billington. *Factual and Plausible Reasoning*. College Publications, London, UK, 2019.
- [3] A. Bochman. A causal approach to nonmonotonic reasoning. *Artificial Intelligence*, 160(1-2):105–143, 2004.
- [4] A. Bochman. A causal theory of abduction. *Journal of Logic and Computation*, 17(5):851–869, 2007.
- [5] A. Bochman. A logic for causal reasoning. In *IJCAI’03: Proceedings of the 18th International Joint Conference on Artificial Intelligence*, pages 141–146, San Francisco, CA, USA, 2013. Morgan Kaufmann Publishers.
- [6] F. M. Cornford. *Plato’s cosmology: the Timaeus of Plato*. Routledge, 2014.
- [7] P. Ferraris. A logic program characterization of causal theories. In *IJCAI’07: Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 366–371, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers.
- [8] H. Geffner. Causal theories for nonmonotonic reasoning. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, pages 524–530, Palo Alto, CA, USA, 1990. AAAI Press.
- [9] L. Giordano, A. Martelli, and C. Schwind. Dealing with concurrent actions in modal action logics. In *ECAI’98: Proceedings of the 13th European Conference on Artificial Intelligence*, pages 537–541, New Jersey, USA, 1998. John Wiley and Sons.
- [10] L. Giordano, A. Martelli, and C. Schwind. Ramification and causality in a modal action logic. *Journal of Logic and Computation*, 10(5):625–662, 2000.
- [11] L. Giordano and C. Schwind. Towards a conditional logic of actions and causation. In S. Flesca, S. Greco, G. Ianni, and N. Leone, editors, *JELIA 2002: Logics in Artificial Intelligence*, pages 382–393, Berlin, Germany, 2002. Springer.
- [12] L. Giordano and C. Schwind. Conditional logic of actions and causation. *Artificial Intelligence*, 157(1-2):239–279, 2004.

- [13] E. Giunchiglia, J. Lee, V. Lifschitz, N. McCain, and H. Turner. Nonmonotonic causal theories. *Artificial Intelligence*, 153(1-2):49–104, 2004.
- [14] D. Hume. *A treatise of human nature*. Courier Corporation, 2003.
- [15] J. Ji and F. Lin. *Lecture Notes in Computer Science, Volume 7265: Correct Reasoning*, volume 7265 of *Lecture Notes in Computer Science*, chapter From Turner’s Logic of Universal Causation to the Logic of GK, pages 380–385. Springer, Berlin, Germany, 2012.
- [16] J. Ji and F. Lin. Turner’s logic of universal causation, propositional logic, and logic programming. In P. Cabalar and T. Son, editors, *LPNMR 2013: International Conference on Logic Programming and Nonmonotonic Reasoning*, pages 401–413, Berlin, Germany, 2013. Springer.
- [17] J. Lee. Nondefinite vs. definite causal theories. In V. Lifschitz and I. Niemela, editors, *LPNMR 2004: International Conference on Logic Programming and Nonmonotonic Reasoning*, pages 141–153, Berlin, Germany, 2004. Springer.
- [18] D. Lewis. Causation. *Journal of Philosophy*, 70(17):556–567, 1973.
- [19] V. Lifschitz. On the logic of causal explanation. *Artificial Intelligence*, 96(2):451–465, 1997.
- [20] F. Lin. Embracing causality in specifying the indirect effects of actions. In *IJCAI’95: Proceedings of the 14th international joint conference on Artificial Intelligence*, volume 2, pages 1985–1991, San Francisco, CA, USA, 1995. Morgan Kaufmann Publishers.
- [21] J. Losee. *Theories of Causality*. Routledge, New York, USA, 2011.
- [22] J. L. Mackie. *The cement of the universe: A study of causation*. Oxford: Clarendon Press, 1974.
- [23] N. McCain and H. Turner. Causal theories of action and change. In *AAAI’97/IAAI’97 Proceedings of the fourteenth national conference on artificial intelligence and ninth conference on Innovative applications of artificial intelligence*, pages 460–465, Rhode Island, USA, 1997. AAAI Press.
- [24] J. McCarthy. Circumscription - a form of non-monotonic reasoning. *Artificial Intelligence*, 13(1-2):27–39, 1980.
- [25] R. Moore. Semantical considerations on nonmonotonic logic. *Artificial Intelligence*, 25(1):75–94, 1985.

- [26] D. Nute. Defeasible logic. In O. Bartenstein, U. Geske, M. Hannebauer, and Y. O., editors, *International Applications of Prolog 2001: Web Knowledge Management and Decision Support*, volume 2543 of *Lecture Notes in Computer Science*, pages 151–169. Springer, 2001.
- [27] S. Patel. *Nuclear physics: an introduction*. New Age International, Limited, New Delhi, India, 2000.
- [28] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1-2):81–132, 1980.
- [29] S. Russel, P. Norvig, et al. *Artificial intelligence: a modern approach*. Pearson Education Limited, 2013.
- [30] C. Schwind. Causality in action theories. *Linköping Electronic Articles in Computer and Information Science*, 4(4), 1999.
- [31] Y. Shoham. Nonmonotonic reasoning and causation. *Cognitive Science*, 14(2):213–252, 1990.
- [32] N. K. Smith. *Immanuel Kant's critique of pure reason*. Read Books Ltd, 2011.
- [33] M. Stabin. *Radiation Protection and Dosimetry*. Springer, New York, NY, USA, 2007.
- [34] P. Suppes. *A probabilistic theory of causality*. Amsterdam: North-Holland Publishing Company, 1973.
- [35] H. Turner. A logic of universal causation. *Artificial Intelligence*, 113(1):87–123, 1999.
- [36] H. Turner. *Handbook of Knowledge Representation*, chapter Nonmonotonic Causal Logic, pages 759–776. Elsevier Science, San Diego, CA, USA, 2008.
- [37] J. Van Cleve. *Problems from Kant*. Oxford University Press, New York, NY, USA, 1999.
- [38] D. Zhang and N. Foo. Epdl: A logic for causal reasoning. In *IJCAI'01: Proceedings of the 17th International Joint Conference on Artificial Intelligence*, pages 131–136, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers.