A Formal Approach to Relevance: Extended Abstract

James P. Delgrande
School of Computing Science,
Simon Fraser University,
Burnaby, B.C.,
Canada V5A 1S6

Francis Jeffry Pelletier
Department of Computer Science,
University of Alberta,
Edmonton, Alberta
Canada T6G 2H1

Abstract

We present a précis of an approach to relevance as it pertains to weak conditionals and defeasible reasoning. The notion of relevance is taken as a relation between a property (such as, being green) and a conditional (such as, birds fly). A series of principles characterising a minimal, parsimonious notion of relevance is developed. Lastly, an explicit definition that agrees with the postulates is given.

Introduction

Naïvely, a property is relevant to another property if, perhaps with other information, it leads us to change our mind concerning whether the second property holds. Consider the following assertions:

Birds (normally) fly. Green birds fly. Birds with broken wings do not fly. Green birds with broken wings do not fly.

Intuitively, being green is irrelevant to whether or not a bird flies, whereas having a broken wing is relevant. However, things are not quite so clear as this. Thus while having a broken wing is relevant to whether a green bird flies, it is not apparent that having a broken wing *and* being green is relevant to a bird's being able to fly – at least, if it *is* relevant, it is not particularly parsimonious.

The same notion of relevance is present in counterfactual statements, deontic assertions, hypothetical assertions, and other subjunctive assertions. For example, if I am trying to get to my dentist's office, we might have the assertions:

If the office is above the third floor, I will take the elevator. But, if the elevator is broken, I will take the stairs.

Clearly the working order of the elevator is relevant; the colour of the elevator is not.

Our intent is to specify what it means for some thing to be relevant in settings such as the preceding. In the following section we introduce the formal framework and review previous approaches to defining relevance. After this, we briefly motivate our overall approach and explore the notion of relevance by proposing a series of principles implicitly characterising this term. Lastly, an explicit definition agreeing with these principles is given. Further details may be found in (DP94).[1]

Background

A common feature of the examples in the previous section is that we have conditionals of the form $\alpha \Rightarrow \gamma$ and $\alpha \wedge \beta \Rightarrow \neg\gamma$ which are simultaneously and nontrivially satisfied. In addition such conditionals do not support transitivity, for example: *Birds fly. Penguins are birds.* yet *Penguins don't fly.* The class of *conditional logics* (Sta68; Lew73; Nut80) is intended to formally capture such weak (subjunctive) conditionals. The semantic theory for these logics is usually expressed in terms of a possible worlds framework. The general idea is that the truth value at a world w of a conditional $\alpha \Rightarrow \beta$ depends both on w and on some subset of the worlds in which α is true: $\alpha \Rightarrow \beta$ is true at a world w just in case the relevant subset of the worlds (which depends both on w and on α) requires that β is true also. (The idea is that we will want to consider a subset of the worlds in which α is true, but the correct subset will vary from world to world.) In this sense the conditional can be regarded as a necessity operator on β but where the subset of "pertinent" worlds depends in some way on α as well as the world at which the sentence is being evaluated; for this reason the operator $\Rightarrow$ is referred to as a *variably strict conditional* or simply a *variable conditional*. Thus in Lewis's approach, $\alpha \Rightarrow \beta$ is true if the least set of worlds (or *sphere*) that have α true also have β true. In these logics the sentences $\{\alpha \Rightarrow \gamma, \alpha \wedge \beta \Rightarrow \neg\gamma\}$ and $\{\alpha \Rightarrow \beta, \beta \Rightarrow \gamma, \alpha \Rightarrow \neg\gamma\}$ may be simultaneously and nontrivially satisfied. For this reason, the conditional $\Rightarrow$ is also referred to as a *weak* conditional as well as a variable conditional. These logics have been used to characterise a wide class of subjunctive conditionals. In the following we will assume some such underlying logic, the further details of which need not con-

[1] As a result of space limitations, we leave a number of important issues to the full paper, including a discussion of the appropriate *type* of definitions: whether implicit or explicit, at the object- or meta-level, and so on.

cern us here (but see (DP94)). The connective $\Rightarrow$ will represent a weak conditional, corresponding to Lewis's "would" conditional (wherein a conditional cannot be vacuously true); $\supset$ will be used for standard material implication.

Related Work (SG87) explicitly addresses the problem of irrelevance in problem-solving systems (SG87). There the idea is to explicate the notion of a fact f being irrelevant to fact g given a knowledge base M. In the full paper we argue that these notions are syntactic, depending on the *form* of sentences. This is inappropriate to the task at hand because it means a sentence could be relevant to something, but a logically equivalent one not; and that a sentence could be relevant to something but not relevant to a logically equivalent way of stating it. More closely-related approaches are given in (Gär78) and (DP93). Gärdenfors specifies what it means for a property p to be relevant to r on evidence e (and with respect to some background theory). Darwiche and Pearl describe what it means for a database Δ to find X (logically) *independent_of* Y, given Z. Space limitations preclude a lengthy discussion, but we note that neither approach is appropriate here: Gärdenfors deals with a probabilistic interpretation of relevance, and Darwiche/Pearl deal with a restricted interpretation in classical databases. In contrast, we are concerned with a general specification of this notion, in the class of conditional logics.

Lastly, the relevance logics of (AB75) and subsequent researchers, deal with a notion with (arguably) the restricted interpretation "deductively-relevant-to".

Initial Considerations

There are at least two distinct senses of the term *relevant.* First, relevance may be taken as a relation between properties. For example, we might say that being feathered is relevant to birdness. This sense is treated, for example, in (Goo72). Following Goodman, we call this sense of *relevant* 'about' (and so might assert for example that Vancouver is *about* British Columbia). However as we show in the full paper, this notion is inappropriate here.

Second, relevant may be taken as a relationship between a property and a proposition represented by a conditional; this is the sense that we adopt. Thus, raather than saying that albinoism is relevant to being black, we say that albinoism is relevant to ravens being black. So the goal is to specify conditions for a property α being relevant to $\beta \Rightarrow \gamma$.

Intuitively, α is relevant to a conditional $\beta \Rightarrow \gamma$ if the addition of α changes our belief (confidence, whatever) in γ. We have the following tentative definition:

Definition 1 *α is relevant to $\beta \Rightarrow \gamma$ iff:*

$$\mathcal{T} \models \beta \Rightarrow \gamma \ \textit{ and } \ \mathcal{T} \models \alpha \wedge \beta \Rightarrow \neg\gamma, \qquad \textit{or}$$
$$\mathcal{T} \models \beta \Rightarrow \neg\gamma \ \textit{ and } \ \mathcal{T} \models \alpha \wedge \beta \Rightarrow \gamma.$$

This seems to work relatively well. Consider a default theory where birds fly; shorebirds fly; Australian birds do not fly; Australian shorebirds fly; and Australian shorebirds with broken wings do not fly. Hence:

$$B \Rightarrow F,\ ShB \Rightarrow F,\ AuB \Rightarrow \neg F,\ AuB \wedge ShB \Rightarrow F,\ AuB \wedge ShB \wedge BW \Rightarrow \neg F.$$

In addition suppose Australian birds and shorebirds must be birds, so $\Box(AuB \supset B)$ and $\Box(ShB \supset B)$. We obtain:

ShB	is not relevant to	$B \Rightarrow F$.
ShB	is relevant to	$AuB \Rightarrow \neg F$.
ShB	is not relevant to	$B \wedge BW \Rightarrow \neg F$.

This notion of relevance seems to be heading in the right direction. However, according to the tentative definition, if something is relevant to a conditional, then so is that thing conjoined with any other property, and we find this objectionable because we want to know *exactly what* brought about the relevance. We want a notion of "minimal relevance", and so we reject the following "augmentation" principle:

AUG: If α is relevant to $\beta \Rightarrow \gamma$ then $\alpha \wedge \delta$ is relevant to $\beta \Rightarrow \gamma$.

This would not be a reasonable relation, since it makes too many things relevant. For example, if all we know is that birds normally fly, but birds with broken wings do not, then it seems that the only thing relevant to birds flying is having a broken wing. Having a broken wing and being green is not relevant; nor is having a broken wing together with the current state of the Tokyo stock market. There are two reasons why such assertions are not relevant. First, if we were making a default inference concerning whether a bird flies or not, we would want to consider only whether it had a broken wing or not; we clearly would not want to also consider the Tokyo stock market. Second, intuitively, a relevant condition should itself be *parsimonious* in that it ought not contain as part an irrelevant condition.[2]

Principles of Relevance

This section proposes a number of principles governing relevance with respect to weak conditionals. First, *relevant* and *irrelevant* are taken as complementary notions:

Definition 2 *α is relevant to $(\beta \Rightarrow \gamma)$ iff α is not irrelevant to $(\beta \Rightarrow \gamma)$.*

Relevance in the sense under discussion is clearly a semantic notion, so no formalisation of relevance should depend on syntactic form:

[2]This is precisely the stance taken in diagnostic reasoning, where one is interested in a *minimal* diagnosis. Thus having influenza could be an adequate diagnosis accounting for fever, muscular pain, etc. However having influenza and a lawn that needs cutting, while equally well accounting for the symptoms, would not constitute an adequate diagnosis.

P1: If α is relevant to $\beta \Rightarrow \gamma$ and

$$\mathcal{T} \models \alpha \equiv \delta \text{ and } \mathcal{T} \models \beta \equiv \eta \text{ and } \mathcal{T} \models \gamma \equiv \zeta$$

then δ is relevant to $\eta \Rightarrow \zeta$.

The conditional logics we are considering allow substitutivity under equivalents with respect to $\Rightarrow$ in the antecedent of a conditional, and so we could replace principle **P1** with the stronger version.

P1a: If α is relevant to $(\beta \Rightarrow \gamma)$ and $\mathcal{T} \models (\alpha \Leftrightarrow \delta)$ and $\mathcal{T} \models (\beta \Leftrightarrow \eta)$ and $\mathcal{T} \models (\gamma \equiv \zeta)$ then δ is relevant to $\eta \Rightarrow \zeta$.

On the other hand, no weakening of the antecedent of a conditional should be relevant:

P2: If $\mathcal{T} \models \beta \supset \alpha$ then α is not relevant to $\beta \Rightarrow \gamma$.

If $\top$ is a designated propositional sentence true in all models, then if we take α as $\top$ in **P2**, we obtain the result that:

$$\top \text{ is not relevant to } \beta \Rightarrow \gamma. \quad (1)$$

If α and β are mutually contradictory, then α has nothing to say about whether $\beta \Rightarrow \gamma$:

P3: If $\mathcal{T} \models \Box\neg(\alpha \wedge \beta)$ then α is not relevant to $\beta \Rightarrow \gamma$.

We obtain as corollaries of this principle that:

$$\neg\beta \text{ is not relevant to } \beta \Rightarrow \gamma. \quad (2)$$

$$\neg\top \text{ is not relevant to } \beta \Rightarrow \gamma. \quad (3)$$

Our criterion of parsimony leads to the next principle, which states that there are irrelevant properties:

P4: If $\models \gamma_1 \vee \ldots \vee \gamma_n$ then for some k, where $1 \leq k \leq n$, we have that γ_k is not relevant to $\beta \Rightarrow \gamma$.

That is, if $\gamma_1, \ldots, \gamma_n$ essentially cover all possibilities, then for any given conditional, some γ_i is irrelevant. For $n = 2$ in **P4** we obtain:

$$\text{If } \delta \text{ is relevant to } \beta \Rightarrow \gamma \text{ then } \neg\delta \text{ is irrelevant to } \beta \Rightarrow \gamma. \quad (4)$$

Letting $n = 1$ in **P4**, also yields (1).

The next postulate is one which may or may not be adopted, depending on one's intuitions.

P5: If α is relevant to $\beta \Rightarrow \gamma$ then α is relevant to $\beta \Rightarrow \neg\gamma$.

The following relation is appealing, but isn't quite right:

If $\mathcal{T} \models \alpha \Rightarrow \gamma$, $\mathcal{T} \models \beta \Rightarrow \neg\gamma$ and $\mathcal{T} \models \Diamond(\alpha \wedge \beta)$ then β is relevant to $\alpha \Rightarrow \gamma$ or α is relevant to $\beta \Rightarrow \neg\gamma$.

Thus, if $\alpha \Rightarrow \gamma$ and $\beta \Rightarrow \neg\gamma$ are true in $\mathcal{T}$, then given $\alpha \wedge \beta$ we will either not conclude γ by default or else won't conclude $\neg\gamma$ by default. The difficulty however is that we don't know whether the conclusion γ or $\neg\gamma$ is blocked; and this goes against our interpretation of "relevant" is that of "provably relevant".

As we discuss below, Definition 1 leads to contentious results, in that strengthenings or weakenings of relevant conditions are not reasonably handled. However, we want to admit *some* strengthenings and weakenings of relevant conditions, and it is to this issue that we now turn. Consider strengthenings first; according to Definition 1, we would have that:

$$ShB \wedge BW \text{ is not relevant to } AuB \Rightarrow \neg F \quad (5)$$

$$BW \wedge Gr \text{ is relevant to } B \Rightarrow F. \quad (6)$$

Thus, in the first case, being a shorebird with a broken wing is not relevant to Australian birds flying, since Australian birds don't fly, nor do Australian shorebirds with broken wings. However, if we consider the individual conjuncts ShB and BW, we see that:

ShB is relevant to $AuB \Rightarrow \neg F$ and BW is relevant to $AuB \wedge ShB \Rightarrow F$.

Arguably then $ShB \wedge BW$ should be relevant to $AuB \Rightarrow \neg F$ since, given AuB along with $ShB \wedge BW$, we would want to conclude $\neg F$ by default, but for different reasons than if we were just given AuB alone. Thus $ShB \wedge BW$ *should* be relevant to $AuB \Rightarrow \neg F$, since ShB is. We add the condition:

P6: If α is relevant to $\beta \Rightarrow \gamma$ and δ is relevant to $\alpha \wedge \beta \Rightarrow \gamma$ then $\alpha \wedge \delta$ is relevant to $\beta \Rightarrow \gamma$.

This means that $ShB \wedge BW$ is now relevant to $AuB \Rightarrow \neg F$, and hence the problem with (5) is solved; but we are still left with (6) above, the fact that having a broken wing and being green is relevant to a bird flying. Someone might wish to argue that $BW \wedge Gr$ *is* relevant to $B \Rightarrow F$, since clearly birds fly whereas green birds with broken wings do not. However *this* argument, relying on augmentation, seems to confuse (6) with the assertion:

BW is relevant to $B \wedge Gr \Rightarrow F$

This last assertion is clearly acceptable, since green birds normally fly whereas green birds with broken wings normally do not. Note that if we were to take the position that $BW \wedge Gr$ is relevant to $B \Rightarrow F$ then **P6** could be replaced by an augmented variant, which we have already rejected:

P6′: If α is relevant to $\beta \Rightarrow \gamma$ then $\alpha \wedge \delta$ is relevant to $\beta \Rightarrow \gamma$.

For similar reasons we reject another variant of **P6**:

P6″: If α is relevant to $\beta \Rightarrow \gamma$ and δ is relevant to $\beta \Rightarrow \gamma$ then $\alpha \wedge \delta$ is relevant to $\beta \Rightarrow \gamma$.

For, if α and δ are both relevant to $\beta \Rightarrow \gamma$ then it must be that $\beta \wedge \alpha \Rightarrow \gamma$ and $\beta \wedge \delta \Rightarrow \gamma$ have the same truth value in $\mathcal{T}$.[3] Gärdenfors (Gär78) rejects a similar principle, arguing that it leads to unintuitive results.

[3] Again this is the stance taken in diagnostic reasoning: having influenza and having allergies might both account the same set of symptoms; however having-influenza-and-allergies would not be a distinct diagnosis, but rather conflates two individual diagnoses.

Such concerns do not apply to *weakenings* of relevant conditions. Consider for example where we have a theory in which birds fly and (as usual) birds with broken wings do not fly, nor do Australian birds. If we learned that a particular individual either had a broken wing, or was from Australia, then we would nonetheless want to conclude that it did not fly. Consequently we adopt the following:

P7: If α is relevant to $\beta \Rightarrow \gamma$ and δ is relevant to $\beta \Rightarrow \gamma$ then $\alpha \vee \delta$ is relevant to $\beta \Rightarrow \gamma$.

Thus having a broken wing or being an Australian bird is relevant to a bird's flying.

An Explicit Definition of Relevance

In this section we present an explicit definition of relevance that satisfies (most of) the postulates of the preceding subsections. We begin with Definition 1 which supplies a base case, and from this define relevance. A condition is relevant if it is naïvely relevant and not composed of two "independent" relevant conditions, or is formed by iterated notions of relevance:

Definition 3
α *is* naïvely relevant *to* $\beta \Rightarrow \gamma$ *iff:*

$\mathcal{T} \models \beta \Rightarrow \gamma$ *and* $\mathcal{T} \models \alpha \wedge \beta \Rightarrow \neg\gamma$, *or*
$\mathcal{T} \models \beta \Rightarrow \neg\gamma$ *and* $\mathcal{T} \models \alpha \wedge \beta \Rightarrow \gamma$,

Definition 4
α *is* relevant *to* $\beta \Rightarrow \gamma$ *if and only if*

1. *for every* δ *and* δ' *where* $\mathcal{T} \models \alpha \equiv (\delta \wedge \delta')$ *and where* δ *and* δ' *are naïvely relevant to* $\beta \Rightarrow \gamma$ *we have* $\mathcal{T} \models \alpha \equiv \delta$ *or* $\mathcal{T} \models \alpha \equiv \delta'$, *or*
2. *there is* α' *such that* $(\alpha \wedge \beta) \supset (\alpha' \wedge \beta)$ *where* α *is relevant to* $\alpha' \wedge \beta \Rightarrow \gamma$ *and* α' *is relevant to* $\beta \Rightarrow \gamma$.

The notion of relevance expressed in Definition 4, Part 2 ensures that we have a chain of successively stronger propositions, each directly relevant to the next in the "context" of β. Thus, AuB is relevant to $B \Rightarrow F$, and so $AuB \wedge ShB$ is relevant to $B \Rightarrow F$ even though both $B \Rightarrow F$ and $AuB \wedge ShB \wedge B \Rightarrow F$ are true. We also obtain the result:

Theorem 1 *Definition 4 fulfills* **P1 – P7**.

Arguably our definition captures a reasonable notion of *relevant*. For example, we obtain the following:

BW	is relevant to	$B \Rightarrow F$
$BW \wedge Gr$	is not relevant to	$B \Rightarrow F$
BW	is relevant to	$B \wedge Gr \Rightarrow F$
BW	is not relevant to	$AuB \Rightarrow \neg F$
$ShB \wedge BW$	is relevant to	$AuB \Rightarrow \neg F$
$BW \vee AuB$	is relevant to	$B \Rightarrow F$
$BW \vee Gr$	is not relevant to	$B \Rightarrow F$

Conclusion

We have presented an investigation of the notion of relevance with respect to weak or defeasible conditionals. The intent has been to characterise this notion with respect to the class of "commonsense" or weak conditionals, as expressed within the general framework of conditional logics (encompassing logics of counterfactuals, default properties, obligation, and other subjunctive conditionals). A series of principles characterising relevance was presented and from this an explicit definition is given.

This approach may have some practical consequences. Typically in conditional logic one does not have modus ponens for the weak conditional. For example, given that birds normally fly, we cannot conclude that a particular individual flies. Yet the consequent of the conditional constitutes a plausible *default* conclusion. Hence, all other things being equal, a bird may (pragmatically) be concluded to fly. What the present work may allow is a means of sanctioning default inferences in a wide class of logics. Thus, given that an individual is a green bird, we may be able to argue there are no relevant conditions "blocking" the consequent, and so conclude that a bird flies. Hence default reasoning is *reduced* to notions of relevance. The advantage of this programme, clearly, is that it would provide a uniform, justified approach to default reasoning in a wide class of logics, and consequently for a wide range of applications.

Acknowledgments We thank Fahiem Bacchus, Craig Boutilier, Randy Goebel, Hector Levesque, Bob Mercer, Eric Neufeld, David Poole, and Ray Reiter: attendees at a workshop on relevance held at Kananaskis, Alberta, November 1992. This research was supported in part by the NSERC of Canada and the IRIS Program of the Networks of Centres of Excellence.

References

A.R. Anderson N.D. Belnap Jr. *Entailment: The Logic of Relevance and Necessity, Vol. I.* Princeton UP, 1975.

Adnan Darwiche and Judea Pearl. Embracing causality in classical logic. Published?, 1993.

James P. Delgrande and Francis Jeffry Pelletier. A formal approach to relevance. In preparation, 1994.

P. Gärdenfors. On the logic of relevance. *Synthese*, 37(3):351–367, 1978.

N. Goodman. About. In *Problems and Projects.* Bobbs Merrill Co., Inc., 1972.

D. Lewis. *Counterfactuals.* Harvard UP, 1973.

J. McCarthy. Circumscription – a form of non-monotonic reasoning. *Artificial Intelligence*, 13:27–39, 1980.

D. Nute. *Topics in Conditional Logic*, Reidel, 1980.

D. Subramanian and M. R. Genesereth. The relevance of irrelevance. In *Proc. IJCAI-87*, pp. 416–422, 1987.

R. Stalnaker. A theory of conditionals. In N. Rescher, ed., *Studies in Logical Theory*, pp. 98–112. Blackwell, 1968.