The Relevance of Irrelevance

Devika Subramanian
Michael R. Genesereth

Department of Computer Science
Stanford University, Stanford, CA 94305

Abstract

This research describes the formalization of statements of the form *fact f is irrelevant to fact g given theory M.* We motivate the need for representing and reasoning with such statements in problem-solving systems, and outline the semantics and properties of statements about irrelevance. We then describe a logic irrelevance that serves as a language for specifying irrelevance claims in the world, and present an associated calculus that allows us to draw new irrelevance conclusions from given ones. The utility of the formalization and the types of inferences it sanctions are demonstrated with examples from data interpretation, representation reformulation and experiment design.

I Introduction

Often, wisdom is knowing what to ignore. A resource-limited agent with a very detailed theory of the world needs to construct a simpler theory that allows it to make predictions at the required level of accuracy within the given resource constraints. Such an agent has to identify parts of its theory that are irrelevant to the class of predictions it is designed to make, and weaken its theory to make it computationally tractable.

This paper introduces a class of statements called irrelevance statements and identifies their special role in problem solving, both in representing certain types of knowledge, and in the construction of weaker theories with better computational properties. We present a theoretical framework for irrelevance and develop a hierarchy of logics that capture different senses of irrelevance. An axiom system for one member of this class is presented along with an irrelevance proof that uses it. Efficient graph-theoretic methods of detecting irrelevance in special cases are indicated.

Informally, a fact f is irrelevant to the fact g in a theory M, written as I(f,g,M), if there exists a weakening of M that allows us to establish the truth of g without committing ourselves on f. A system of logics can be developed by imposing different restrictions on the constructions of these weakenings. These logics permit the specification of domain-specific claims about irrelevance and their axioms allow the deduction of new claims. An irrelevance statement is a fact *about* the relationship between two facts in the theory M. Reasoning about irrelevance is thus a mode of meta-theoretic reasoning.

The paper is organized as follows. Section II motivates the problem of representing and reasoning with statements about irrelevance. The relationships to the notion of relevance are also outlined here. Section III presents some properties of irrelevance. Section IV formalizes an intuitive notion of irrelevance. We present the propositional version of two specializations of this definition and indicate extensions to the first-order case. The utility of this formalization is demonstrated with examples in Section V. Section VI briefly discusses issues in the acquisition of facts about irrelevance. Section VII concludes with a summary of the main points and lists future directions for this research.

II Motivations

Informally, we address the *what is relevant* question by considering the complementary question, since in some cases it is far easier to determine and specify what is irrelevant. Also useful classes of such statements focus the construction of weaker and often more computationally tractable theories. Both these points are discussed in some detail below.

An important reason for building the ability to reason about irrelevance into systems is that we would like to give advice to a problem-solving system (or bias to a learning system) in terms of irrelevance. For the missionaries and cannibals problem, we would like to tell the problem-solver that the names of the missionaries and cannibals are irrelevant, and have the system clump the missionaries and cannibals into sets and deal with the cardinalities of these sets. Amarel indicates this sort of reasoning in his well-known 1968 paper[Amarel 1968]. *Removing irrelevant facts and objects from a formulation is an important method of changing representations.*

Explicitly identifying what to ignore is easier in some frameworks than in others. For instance, in Dendral, there are two classes of mass spectrograph points that are ig-

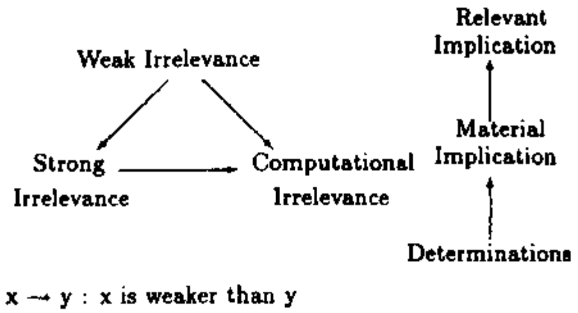

Figure 1: Logics of Relevance and Irrelevance

nored for the purposes of the structure interpretation task. Scientists have much sharper criteria for the data points to ignore rather than the data points to include. *Specification of irrelevance claims is then a valuable complementary mode of expressing knowledge about a domain.*

Yet another motivation for reasoning about irrelevance is the need for problem-solving systems to reason flexibly at varying grain sizes [Hobbs 1985]. These systems require the ability *to recognize and ignore detail irrelevant to their current goals* in order to shift to a bigger grain size where those goals can be achieved more efficiently.

Reasoning about irrelevance can thus be used as a *basis for focusing attention in both inductive and deductive tasks.* In induction, certain irrelevance claims bias the learner towards the construction of approximate, simpler theories. In deductive tasks, irrelevance statements help focus search by identifying unfruitful or redundant paths. They also help *restructure the search space* by collapsing primitives in a formulation.

A formal study of irrelevance reveals that there is interesting substructure in the space of useful irrelevance inferences. There is a hierarchy of irrelevance logics that is similar to the hierarchy of logics that capture the notion of relevance [Anderson and Belnap 1975],[Davies and Russell 1986]. See Figure 1. These logics are not exact duals of each other; many aspects of the relationship between them remain to be investigated. The statement that f is *relevant* to g in M conveys the information that knowing f in M restricts the space of possibilities for g. The statement that f is *irrelevant* to g in M indicates that the value of g is insensitive to the value off, in that even if f were changed in M, g would not. In the missionaries and cannibals puzzle, the names of the missionaries are irrelevant to the scheduling of the boat trips. This means that even if the names were changed*, the solution (which does not refer to the missionaries by name) would not. This irrelevance statement gives us the justification to modify the formulation in which missionaries are named, to the more abstract one that uses the cardinality of the set of missionaries.

The two notions of relevance and irrelevance are complementary; one expresses a dependence between two facts and the other captures a one-sided lack of dependence (g being not dependent on f). However, the normal use of an irrelevance statement is to modify M while preserving g, the normal use of a relevance statement is to infer restrictions on g given f: thus the inferences they sanction are not duals.

III Properties of Irrelevance

Our notion of irrelevance has the properties of non-monotonicity, intransitivity and asymmetry.

A. Non-monotonicity

Irrelevance is non-monotonic in general. If the addition of a fact to M causes the truth of g to become contingent on f, then the irrelevance statement ceases to hold. Before Descartes, facts about algebra would have been irrelevant for proving a theorem in plane geometry. Descartes' discovery of analytic geometry connected these two theories together and each has since become relevant to the other.

B. Intransitivity

Irrelevance is not transitive. An example brings this out:

The blight on the tea crop in China is irrelevant to my writing this paper. My writing this paper is irrelevant to the price of tea in China.
Had irrelevance been transitive, we could deduce that:
The blight on the tea crop in China is irrelevant to the price of tea in China.
This makes it difficult to propagate irrelevance conclusions and derive useful irrelevance claims from given ones, except in special cases in Section IV.A.

C. Asymmetry

Irrelevance is also not symmetric in the general case. However, it is symmetric for a certain class of irrelevance statements I(p,q,M) where p and q are logically independent. Thus knowing p does not help in constraining the possible values of q in the context of M, and knowing q does not help in constraining the possible values of p. The statement *My writing this paper is irrelevant to the price of tea in China* is an example of a symmetric irrelevance statement. These statements capture independence between facts and thus provide a natural way of factoring a formulation. Thus p and q could be put in different parts

"as long as they remain distinct from each other!

of the formulation: a change in one part will not have repercussions in the other for any use of the formulation.

IV Theoretical Aspects of Irrelevance

We say that fact f is irrelevant to fact g modulo the set of sentences M, written as I(f,g,M), if f is inessential to the truth of g in M, that is, changing the truth value of f in M would not affect the truth of g. One possibility for a formal definition is:

Definition 1: f is irrelevant to g modulo M if $M \models g$ and we can construct a set M' of sentences with the following properties:

1. M' is weaker** than M. $M \Rightarrow M'$
2. M' entails g. $M' \models g$.
3. M' is non-committal on f. $M' \not\models f$ and $M' \not\models \neg f$

Consider this simple propositional example, l and t stand for lightning and thunder respectively.
$M_1 = \{l, t\}$ and $M_2 = \{l, l \Rightarrow t\}$.
According to the above definition, l is irrelevant to t in M_1 since we can construct $M_1' = \{t\}$ that has the required properties. Also, we can show that l is irrelevant to t in M_2 because we can construct $M_2' = \{t\}$. However, we would not like to conclude in the latter case that l is irrelevant to t. This anomaly occurs because we imposed no restrictions on the weakenings of M. This allowed us to construct M' that just contains g (since $M \models g$) and allowed us to conclude everything else in M irrelevant to it. Even though M_1 and M_2 are identical model-theoretically, the irrelevance judgements were different in the two sets. Thus irrelevance claims are very sensitive to M, which is usually not a closed set of sentences. Irrelevance is an inherently proof-theoretic notion and model theory fails to capture the lack of a dependency between truths of propositions except in the strongest sense, namely independence. For the purposes of this paper, we will adopt the following definition that constructs weakenings of M by subsetting M. We call this category of irrelevance statements, *weak irrelevance* statements. Section V.D. examines the case where weakenings of a theory are constructed by collapsing terms in the theory.

Definition 2: f is weakly irrelevant to g modulo M, written as WI(f,g,M) if $M \models g$ and we can construct the set M' with the following properties.

1. $M' \subseteq M$
2. M' is non-committal on f. $M' \not\models f$ and $M' \not\models \neg f$
3. M' is a maximal subset of M with the above two properties.
4. M' entails g.

**The set of models of M' are a superset of the set of models of M.

If there are multiple maximal subsets of M then g must hold in at least one of them. Now, / is no longer irrelevant to *t* in M2 because we cannot construct the set M2' with these properties. The judgement of irrelevance can be made without imposing the requirement that the subset M' be maximal. However, the maximality condition ensures that the weakening of M constructed is the most conservative one; i.e. we weaken M *minimally* so that it becomes non-committal on f and then show that g still holds in this set. Definition 2 is a specialization of the counterfactual construction proposed in [Ginsberg 1985]. This is because irrelevance statements are special cases of counterfactuals.

We investigate the properties of weak irrelevance (Definition 2) and then define two useful specializations of the definition: strong irrelevance and computational irrelevance.

A. Weak Irrelevance

If f is weakly irrelevant to g in M, then there exists *at least one* proof of g in M that does not use f essentially. This happens when either g is independent of f in M or there are multiple proofs of g in M, ones that use f and ones that don't, so that g can be established without f.

Example $M = \{f, h, f \Rightarrow g, h \Rightarrow g\}$

We have WI(f,g,M) and WI(h,g,M), however, we cannot conclude that WI(f∧h,g,M). Note that the implication holds in the other direction.

Some properties of this definition are given below. Proofs can be found in [Subramanian 1986].
Observation 1: WI is non-symmetric on its first two arguments, in general.
Observation 2: WI can be detected by checking if there exists a path to g that does not go through f in the inference graph representation of M.
Observation 3: Some axioms of WI are:

1. $WI(f,g,M) \equiv WI(\neg f,g,M)$
2. $WI(f1 \wedge f2,g,M) \Rightarrow WI(f1,g,M) \vee WI(f2,g,M)$
3. $WI(f,g1 \wedge g2,M) \Rightarrow WI(f,g1,M) \wedge WI(f,g2,M)$
4. $\neg WI(f,f,M)$
5. $f1 \equiv f2 \Rightarrow [WI(f1,g,M) \Rightarrow WI(f2,g,M)]$
6. $g1 \equiv g2 \Rightarrow [WI(f,g1,M) \Rightarrow WI(f,g2,M)]$
7. $g \in M \wedge f \not\equiv g \Rightarrow WI(f,g,M)$
8. $S = \{p \in M \mid f \Rightarrow p\} \wedge WI(\bigwedge_{p \in S} p,g,M) \Rightarrow WI(f,g,M)$ when $f \not\Rightarrow g$ in M
9. $S = \{p \in M \mid p \Rightarrow g\} \wedge WI(f,\bigwedge_{p \in S} p,M) \Rightarrow WI(f,g,M)$ when $g \not\Rightarrow f$ in M
10. $p \in M \wedge$ Derives-Only(f,p,M) $\Rightarrow WI(f,g,M)$ when $M \models g$

Axioms 1 through 6 are straightforward consequences of our definition. Axioms 7 through 10 identify four special cases of redundancy. Axiom 7 says that if g is present in M, then everything except itself is redundant to it. Axiom 8 states that if all facts that are derivable from f are redundant to g, then f itself is redundant to g in M, unless f directly implies g. Axiom 9 is the dual axiom that shows that f is redundant to g if f is redundant to all facts that derive g, unless g directly implies f. Axiom 10 is a generalization of axiom 7: if p is already in M, then any f that only derives p is redundant to g in M. To complete this axiomatization we need to identify *all* the base cases for proving redundancy. Since detection of redundancy is semi-decidable, there is good reason to believe that we cannot construct a complete axiomatization of WI.

Observation 4: One benefit of proving redundancy of some facts in a theory, is that we can optimize space requirements by simply removing the redundant facts.

An example of this is in Section VLB. A proof of redundancy using the lemmas of WI in the meta-theory of M captures an important property of the proofs of g in M, without exhaustively enumerating them. A problem solver that can prove this redundancy claim at the meta-level can prune a large class of inferences this way. Also if the redundancy statements were made available to the problem solver, it could compile it into the base level formulation by simply throwing away those facts that cause the redundancy. This leads to savings in space. To ensure that savings in time in *proving g* result from this pruning, we need to show that the proofs of g that fail as a result are derivation ally more complex than the ones that are retained. This is captured in our definition of computational irrelevance. In either case, the *deliberation time* for the problem solver will be reduced because the search space of proofs of g in M is reduced by the removal of redundancy.

B. Strong Irrelevance

We say that f is strongly irrelevant to g in M written as SI(f,g,M) if and only if f is not logically necessary for the truth of g and g is not logically necessary for the truth of f. This means that a proof of g in M can be constructed even if f were false and a proof of f in M can be constructed even if g were false.

Example: $M = \{f, h, h \Rightarrow g\}$. f is SI to g in M.

Syntactic characterization of this relation

SI(f,g,M) if and only if $WI(f,g,M) \wedge WI(g,f,M)$. Using Definition 2 we can see that M can be split into subsets M_1 and M_2 with the following properties.

- $M = M_1 \cup M_2$
- M_1 is a maximal subset of M such that $M_1 \not\models f$ and $M_1 \not\models \neg f$ where $M_1 \models g$.
- M_2 is a maximal subset of M such that $M_2 \not\models g$ and $M_2 \not\models \neg g$ where $M_2 \models f$.

M_1 is the required weakening of M. This definition extends to the first order case straightforwardly. Given the information that g is the only class of queries that need to be answered or the only predicate that needs to be considered for a certain problem solving task, an agent can restrict its attention to M_1 alone. An example where this leads to significant computational gain is in Section V.C.

Model-theoretic characterization of a superset of this relation

Consider all models of M, call the set Int_M. Project assignments to f from Int_M: this is the set $Int_M|_f$. Similarly, project out assignments to g from the set of all models of M, call the set $Int_M|_g$. When f and g are independent of each other in M, the following model-theoretic equation holds.

$$Int_M|_{(f,g)} = Int_M|_f \otimes Int_M|_g$$

When the set M can be factored into subsets M_1 and M_2 as above, the models of M have the Cartesian factorability property. Unfortunately, the converse doesn't hold.

We now describe some properties of this definition.

Observation 1: SI is a special case of Definition 2, i.e. $SI(f,g,M) \Rightarrow WI(f,g,M)$.

Observation 2: The syntactic characterization of SI captures a strict subset of the cases covered by the model-theoretic definition.

This is because the construction is strongly dependent on the form of M. In order to make the two definitions correspond, we have to rewrite M in a normalized form.

Observation 3: SI can be detected efficiently by representing the sentences of M in an inference graph and checking that f and g belong to different forests of this graph.

Observation 4: M_1 and M_2 are computationally better than the combination M for processes whose computational complexity is a function f of the size of the set they work on, where $f(|M|) > f(|M_1|) + f(|M_2|)$.

One such function is the exponential function. Section V.C. presents an example.

C. Computational Irrelevance

This is a refinement of weak irrelevance where the derivations of g in M are ordered by some measure of complexity. We say f is computationally irrelevant to g if we can prove g in M without the use of f, and also show that the proof has better complexity characteristics (it might be shorter or be easier to find). This version of irrelevance leads to the construction of weaker theories with better computational properties.

Syntactic characterization of computational irrelevance

f is CI with respect to g modulo M if we can split M into subsets M_1 and M_2 such that

- $M = M_1 \cup M_2$
- Mi is the maximal set such that Mi $\not\models f$ and
  $M_1 \not\models \neg f$
  $M_1 \models g$
- Deriv-Complexity(g,M_1) better-than Deriv-Complexity(g,Mj)

M_1 is the required weakening of M. This definition extends to the first order case straightforwardly.

Some properties of this definition are:
Observation 1: CI(f,g,M) $\Rightarrow$ WI(f,g,M)
Observation 2: SI(f,g,M) $\Rightarrow$ CI(f,g,M)
The proof uses the fact that the derivational complexity of g in M when M does not entail g is oo.
Observation 3: CI is non-symmetric on its first two arguments.
Observation 4: CI leads to the construction of computationally better subtheories for the class of queries g.
In the missionaries and cannibals problem in which both the individuals and the sets of individuals are present in the formulation, we can show that the existence of the individuals is CI with respect to the construction of the schedule, since the proof that uses the individuals (as opposed to the set of individuals) is derivationally more complex

V Applications

We show how the data filtering procedure in Dendral [Buchanan 1974] can be declaratively formulated in our logic. Our logic thus serves as a specification language for the data selection component of a theory formation system. The reasoning required to capture two types of representation reformulation are then presented. The first is an abstraction reformulation for space efficiency in which we identify and minimize redundancy in a problem formulation. We present a proof of the redundancy of a class of statements in the formulation and show how this can be used as a justification for jettisoning that class of facts from the formulation. The next example formalizes Quine's[Quine 1964] notion of *indisctrnibility of identicals* and shows how our irrelevance logic can be modified to sanction inferences that lead to the construction of theories of coarser granularity. Our last example explores the role of reasoning about irrelevance in factoring theories and automatic experiment design in learners [Subramanian and Feigenbaum 1986].

A. Data Selection in Dendral

In data interpretation problems, it is useful to declaratively represent heuristics for pruning data points. These heuristics are statements about the *computational irrelevance* of a certain class of data points with respect to the particular interpretation task at hand. In Dendral they have the following form:

CIClass of Data points, Structural Theory, Background Theory)

The background theory is the theory of covalent bonding in chemistry as well as the theory of operation of the spectrometer. Isotopic peaks in the mass spectrum are considered irrelevant because even if they were taken into account in the construction of the structural theory of the molecule, the predictions obtained from it would not be significantly different from the theory that didn't take these peaks into account. The irrelevance statement is a control heuristic in the search space of all structural theories of the molecule. These theories are ordered by the data points they cover and this heuristic identifies important data points by exclusion and thus focuses the search in the direction of weaker and simpler theories.

The advantage of specifying these irrelevance claims explicitly is that they have a clearly understood semantics. They can then be derived from other background knowledge. We can reason about the filtering criteria explicitly and do a comparative sensitivity analysis on them. We can use them in conjunction with other background knowledge to derive new computational irrelevance claims.

B. Reformulation: Removing Redundancy

The formulation of a problem has critical impact on its efficiency. The following example shows how we can use irrelevance reasoning to reformulate a kinship problem to obtain computational efficiency. Figure 2 shows a simplified family tree representing a set of *Father* relations. The *Ancestor* relation is defined as the transitive closure of the *Father* relation. The goal is to determine whether two people belong to the same family. The defined relation *Same-Family* is used for this: two people belong to the same family if they have a common ancestor. The time required to compute *Same-Family* in this formulation is proportional to the height of the family tree. If we are given the extra information that *Same-Family* is the only relation that will be queried, we could rewrite the family tree information in terms of the *Founding-Father* relation as in Figure 3. The *Founding-Father* of an individual is his maximal ancestor in the family tree. The time needed to solve a *Same-Family* question in the new formulation is a constant: we need to check if the two individuals have the same maximal common ancestor!

Formulation 2 is constructed from Formulation 1 by discarding information that was irrelevant to the computation of *Same-Family.* This is achieved by eliminating the *Father* relation on grounds of redundancy (captured by weak irrelevance) and then eliminating the intermediate *Ancestor* links by a computational irrelevance argument. What remains after this theory weakening process are the maximal ancestor links in the family tree.

The redundancy claim can be formally proved:

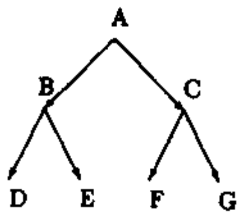

Father(A,B)
Father(A,C)
Father(B,D)
Father(B,E)
Father(C,F)
Father(C,G)
Father(x,y) $\Longrightarrow$ Ancestor(x,y)
Ancestor(x,z) $\wedge$ Ancestor(z,y) $\Longrightarrow$ Ancestor(x,y)
Ancestor(z,x) $\wedge$ Ancestor(z,y) $\Longrightarrow$ SameFamily(x,y)

Figure 2: The Given Formulation

FoundingFather(A,B)
FoundingFather(A,C)
FoundingFather(A,D)
FoundingFather(A,E)
FoundingFather(A,F)
FoundingFather(A,G)
FoundingFather(z,x) $\wedge$ FoundingFather(z,y) $\Longrightarrow$ SameFamily (x,y)

Figure 3: The New Formulation

A(x,y) $\in$ M $\wedge$ Derives-Only(F(x,y), A(x,y), M)

$\Rightarrow$ WI(F(x,y), SF(m,n), M)

$\Rightarrow$ WI(F(x,y), A(z,m) $\wedge$ A(z,n), M) by ax. WI6 & ax. M3

$\Rightarrow$ WI(F(x,y), A(z,m), M) $\wedge$ WI(F(x,y), A(z,n), M) by ax. WI3

Each conjunct is true because of ax. WI10.

C. Factoring of theories and experiment design

Strong irrelevance claims allow a theory to be factored into independent subtheories.

$$\frac{M(Shape,Colour),\ SI(Shape,Colour,M)}{M = M_1(Shape) \cup M_2(Colour)}$$

Providing the above irrelevance fact to a blocks-world learner allows it to factor the learning problem into two independent subproblems: one for learning shape and one for learning colour. The task of devising informative instances is exponential in the size of M, so that a learner that acquires information by active experimentation obtains significant computational advantages by this split [Subramanian and Feigenbaum 1986].

D. Reformulation: Quine's Principle

Quine'8 maxim of the identification of indiscemibles states that objects indistinguishable from one another in the context of a discourse should be construed as identical for that discourse. That is, the references to original objects should be reinterpreted in such a way that indistinguishable originals give way to the same new object.

The fact that the distinction between two objects is irrelevant for our purposes can be stated in our logic as:

$$Irrelevant(term_1 \neq term_2, Query, M)$$

The (perhaps, default) assumption that these 2 distinct terms denote distinct objects is now relaxed. The informal meaning of this is that we would be able to establish *Query* in a revision of *M* that did not treat $term_1$ and $term_2$ as distinct.

We call this term irrelevance and define it using fact irrelevance as follows:

$$Irrelevant(term_1 \neq term_2, Query, M) \equiv$$
$$\forall f.Contains(f, term_1) \wedge \neg True(f|_{term_1/term_2})$$
$$\Rightarrow WI(f, Query, M)$$

This says that all facts in M that cause us to infer that the two terms are distinct are irrelevant to the query.

To implement Quine's maxim, we sanction the following inference:

$$\frac{M(term_1, term_2), Irrelevant(term_1 \neq term_2, Query, M)}{M' = M[term_1/t, term_2/t] \text{ (where t is a new term)}}$$

We remove those facts in M that treat the two terms differently and then introduce a new term in the coarser theory to stand in interchangeably for two distinct terms in the finer theory.

This leads one to minimize the (objects) primitives in one's theory as opposed to minimizing the conclusions from the theory, which is characteristic of most work in circumscription. Such simplification occurs during the construction of Thevenin equivalents in circuits, and in general, in the construction of coarser theories from more detailed ones.

VI Acquisition of Facts about Irrelevance

The three examples above indicate how irrelevance claims can be used for the construction of weaker theories from stronger ones, but leave open the question of how the irrelevance claims are acquired. There are three principal ways of obtaining such facts

- By being told.
- By deriving them in the meta-theory of M using the calculus of irrelevance.
- By active experimentation in the world.

The computational irrelevance claims in Dendral were provided by the designer. The reformulator in Section VLB. formally proved the weak irrelevance of the *Father* relation with respect to the *Same Family* relation. The strong irrelevance between shape and colour could be determined empirically by correlation analyses on the values of these two attributes.

VII Conclusions and future directions

This paper introduced the need to represent and reason explicitly with statements about irrelevance. It identified a class of useful irrelevance statements that lead to the construction of logically weaker theories with better computational properties from stronger, more detailed theories. A theoretical framework for the study of irrelevance was proposed that allowed for the specification of irrelevance claims and the deduction of new irrelevance claims from given ones. The kinds of inferences sanctioned by an irrelevance claim were demonstrated by examples: automatic factorization of theories, minimization of redundancy in, as well as restructuring of search spaces for computational efficiency. We conclude that reasoning about irrelevance is a valuable mode of meta-theoretic reasoning that will enhance the functionality of present-day problem solving systems

Much remains to be done. We are presently exploring the mechanization of irrelevance proofs and the development of a better axiomatization for WI and CI. The characterization of the cases where reasoning about irrelevance is useful needs to be generalized. Identifying special-purpose methods for detecting SI, WI or CI efficiently will make the use of this kind of reasoning practical. Since irrelevance judgements are made in contexts delimited by relevance reasoning, the interweaving of relevance and irrelevance reasoning appears to be a fruitful direction to explore. The technical relationship between the logics presented here and the logics of relevance [Anderson and Belnap 1975][Davies and Russell 1987][Ginsberg 1985] remains to be fully explained. Certain points in the space of theory weakening operations were identified here: the discovery of useful classes of irrelevance statements that sanction further kinds of theory weakenings with provably better computational properties is a challenging open problem.

Acknowledgments

We would like to thank Professors Paul Rosenbloom and Bruce Buchanan for their valuable contributions to this work. Chris Fraley, David Smith, Matt Ginsberg, Stuart Russell, David Wilkins, Haym Hirsh, Jane Hsu and Jeff Finger provided technical help and encouragement. The first author is supported by an IBM doctoral fellowship.

References

1 Amarel, S. "On representation of problems of reasoning about actions", reprinted in *Readings m Artificial Intelligence,* edited by Nilsson and Webber, Kaufmann, 1982.

2 Anderson, A. and Belnap, J. *Entailment: the logic of relevance and necessity,* Princeton University Press, 1975.

3 Buchanan, B. G. "Scientific Theory Formation by Computer", HPP-74-3. Also in *Computer Aided Learning Processes, Nato Advanced Study Institutes Series, Series E: Applied Science, 14, 515, Leyden: Noordhoff,* edited by J.C. Simon, 1976

4 Ginsberg, M. L. "Counterfactuals", KSL-84-43, revised October 1985.

5 Hobbs, J. "Granularity", in *Proc. IJCAI-85,* Los Angeles, CA, pp 432-435.

6 Quine, W.V.O. *From a logical point of view,* Harvard University Press, Cambridge, MA, 2nd edition (revised).

7 Davies, T. and Russell, S.J. "A logical approach to reasoning by analogy", in *Proc. IJCAI-87,* Milan.

8 Subramanian, D. and Feigenbaum, J. "Factorization in Experiment Generation", in *Proc. AAAI-86,* Philadelphia, PA, pp 518-522.

9 Subramanian, D. "The Relevance of Irrelevance", Logic Group Memo 5, 1986.