

Proof theoretic reasoning in System P

Simon Parsons and Rachel A. Bourne

Department of Electronic Engineering,
Queen Mary and Westfield College
University of London,
London E1 4NS, UK.

Abstract

This paper investigates how the rules of System P might be used in order to construct proofs for default rules which take into account the bounds on the probabilities of the consequents of the defaults. The paper defines a proof system, shows that it is sound, and then discusses at length the completeness of the system, and the kind of proofs that it can generate.

Introduction

Default reasoning has been widely studied in artificial intelligence. One of the the most influential piece of work within this area is that of Kraus *et al.* (1990). Kraus *et al.* investigated the properties of different sets of Gentzen-style proof rules for non-monotonic consequence relations, and related these sets of rules to the model-theoretic properties of the associated logics. Their major result was that a particular set of proof rules generated the same set of consequences as a logic in which there is a preference order over models. This system of proof rules was termed System P by Kraus *et al.*, the P standing for “preferential”. System P has been the subject of much research, and is now widely accepted as the weakest interesting non-monotonic system; it sanctions the smallest acceptable set of conclusions from a set of default statements.

The reason that we are interested in the rules of System P is that, in addition to a semantics in terms of a preference order over models, they also have a semantics in terms of infinitesimal probabilities (Adams 1975; Pearl 1988). This semantics can be extended to deal with finite probabilities (Bourne & Parsons 1998), an extension which promises to make System P more useful in practice. However, the use of finite probabilities raises a new problem. In order to keep track of the probability attached to each default, it is necessary to establish a new mechanism for inferring new defaults along with their probabilities, and this is the subject of this paper.

Copyright ©1998, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

Entailment in System P

System P is concerned with conditional assertions of the form $\alpha \sim \beta$. In this context α and β are well-formed formulae of classical propositional logic, and \sim is a binary relation between pairs of formulae. The probabilistic semantics for System P assumes that the propositional variables are the basis of a joint probability distribution which is constrained by the assertions. Each assertion states that the conditional probability of its consequent given its antecedent is greater than or equal to $1 - \epsilon$ for any $\epsilon > 0$. Thus:

Definition 1 *The conditional assertion $\alpha \sim \beta$ denotes the fact that $\Pr(\beta|\alpha) \geq 1 - \epsilon$ for all $\epsilon > 0$.*

Probabilistic consistency is defined as the existence of at least one probability distribution which satisfies these constraints (Adams 1975); probabilistic entailment of a further conditional is defined as probabilistic inconsistency of its counterpart, that is:

Definition 2 *$\alpha \sim \beta$ is p-entailed by Δ iff $\Delta \cup \{\alpha \sim \neg\beta\}$ is not p-consistent.*

This implies that all probability distributions that satisfy Δ also satisfy $\alpha \sim \beta$. However this result may only be achieved by using infinitesimal analysis so that the derived conditional will be constrained to be greater than $1 - \delta$ for any $\delta > 0$ if the ϵ of the original conditionals is made small enough. This can be paraphrased as saying that System P allows us to make our conclusions as close to certainty as we like, provided the conditional probabilities associated with the initial assertions are sufficiently close to certainty.

Using this interpretation of the rules thus means assuming that we are able to give the conditional assertions arbitrarily high conditional probabilities. Often we have less reliable information, and so ϵ is not infinitesimal. When such a set of conditional assertions are used to derive new assertions and these new assertions are themselves used as the basis for new deductions, then some ϵ values will be far from infinitesimal and it is not possible to assess the impact of these non-infinitesimal values on the strength of the new assertions.

Because of this concern, we investigated (Bourne & Parsons 1998) the impact of non-infinitesimal ϵ values

$\alpha \vdash_{\epsilon} \alpha$	Reflexivity	$\frac{\models \alpha \leftrightarrow \beta, \alpha \vdash_{\epsilon_1} \gamma}{\beta \vdash_{\epsilon_1} \gamma}$	Left Logical Equivalence
$\frac{\models \beta \rightarrow \gamma, \alpha \vdash_{\epsilon_1} \beta}{\alpha \vdash_{\epsilon_1} \gamma}$	Right Weakening	$\frac{\alpha \vdash_{\epsilon_1} \beta, \alpha \vdash_{\epsilon_2} \gamma}{\alpha \vdash_{\epsilon_1 + \epsilon_2} \beta \wedge \gamma}$	And
$\frac{\alpha \vdash_{\epsilon_1} \beta, \alpha \vdash_{\epsilon_2} \gamma}{\alpha \wedge \beta \vdash_{\frac{\epsilon_1 + \epsilon_2}{1-\epsilon_1}} \gamma}$	Cautious Monotonicity	$\frac{\alpha \vdash_{\epsilon_1} \gamma, \beta \vdash_{\epsilon_2} \gamma}{\alpha \vee \beta \vdash_{\epsilon_1 + \epsilon_2} \gamma}$	Or

Figure 1: Rules with associated bounds

in the context of a set of rules of inference for System P given by Kraus *et al*, showing how ϵ values are propagated. The results of this investigation are the rules in Figure 1. The rules are written in the usual Gentzen style, with antecedents above the line and consequents below it. Thus the rule ‘And’ says that if it is possible to derive $\alpha \vdash \beta$ such that $\Pr(\beta | \alpha) \geq 1 - \epsilon_1$ and $\alpha \vdash \gamma$ such that $\Pr(\gamma | \alpha) \geq 1 - \epsilon_2$, then it is possible to derive $\alpha \vdash \beta \wedge \gamma$ such that $\Pr(\beta \wedge \gamma | \alpha) \geq 1 - (\epsilon_1 + \epsilon_2)$.

While this work solved the problem of determining the impact of the non-infinitesimal values, it fell short of providing a practical reasoning system. The problem is that although in System P we can tell whether or not $\alpha \vdash \beta$ follows from the initial set of defaults, the procedures for determining this do not permit the propagation of the ϵ values. Thus we can tell if $\alpha \vdash \beta$ follows, and so we can find out if a proof exists, but we can’t determine the associated ϵ value. What we need is a proof theory which allows the ϵ values to be propagated through the proof so that every inferred default has its ϵ value determined, and providing such a proof theory is the subject of this paper.

A proof theory for System P

Normally in generating a proof theory for some logical system the procedure (Gabbay 1996) is to establish two rules for each connective in the underlying language. One rule relates to introducing the connective into a formula, and one relates to eliminating the connective from a formula. The set of rules then define all the legal transformations between formulae, and thus define what may be proved from some initial set of formulae. The process of defining a proof theory thus proceeds from the underlying language to the proof rules.

The situation here is a little different. System P already has a set of proof rules defined. However, these rules do not include introduction and elimination rules for all the connectives in the underlying language, and so do not support a conventional proof theory. However, it is possible to use the existing rules to define a proof theory for a significant part of the underlying language of System P, and this is the approach we adopt.

The proof system

We start with a set of propositions \mathcal{S} , a set of connectives, $\{\neg, \wedge, \vee, \rightarrow, \leftrightarrow, \Rightarrow\}$, and the following rules for

building well-formed formulae in this language:

1. If $\alpha \in \mathcal{S}$, then α is a basic well-formed formula (*bwff*).
2. If α and β are *bwffs* then $\neg\alpha$, $\alpha \wedge \beta$, $\alpha \vee \beta$, $\alpha \rightarrow \beta$, $\alpha \leftrightarrow \beta$ are *bwffs*.
3. If γ and δ are *bwffs*, then $\gamma \Rightarrow_{\epsilon} \delta$ is a default well-formed formula (*dwoff*).
4. Nothing else is a *bwff* or a *dwoff*.

Together all these formulae constitute a language $\mathcal{L}_{\mathcal{S}}$. The denotation of basic well-formed formulae is as in propositional logic, while the meaning of *dwoffs* is the following:

Definition 3 The default $\gamma \Rightarrow_{\epsilon} \delta$ is taken to mean $\Pr(\delta | \gamma) \geq 1 - \epsilon$.

Comparing Definitions 1 and 3 it is clear that the defaults of $\mathcal{L}_{\mathcal{S}}$ are exactly the conditional assertions of System P for a particular finite ϵ , and any set of conditional assertions $\Delta = \bigcup_i \{\alpha_i \vdash \beta_i\}$ will have a corresponding set of *dwoffs* $\Delta' = \bigcup_i \{\alpha_i \Rightarrow_{\epsilon_i} \beta_i\}$. We say that Δ' is the *default dual* of Δ , and Δ is the *assertion dual* of Δ' . We apply the same terminology to single *dwoffs* and conditional assertions. The reason for writing the defaults in this way is to distinguish between the conditional assertions themselves, and the consequence relation which defines what may be inferred from them—a distinction which is not always clear in work on System P. Assuming that we have a knowledge base Δ which consists of a set of *dwoffs*, we can then define the valid set of conclusions which may be drawn from Δ as those sanctioned by the consequence relation $\vdash_{\mathcal{P}}$ defined in Figure 2¹.

The proof rules that define $\vdash_{\mathcal{P}}$ may need a little explanation. The rule Ax is a form of ‘bootstrap’ rule which says that if some default $\alpha \Rightarrow_{\epsilon} \beta$ is in Δ , then were α added to Δ , it would be possible to infer β with probability not less than $1 - \epsilon$. The rule And says that if adding α to Δ makes it possible to infer β with probability no less than $1 - \epsilon_1$ and γ with probability no less than $1 - \epsilon_2$, then adding α to Δ makes it possible to infer $\beta \wedge \gamma$ with probability no less than $1 - (\epsilon_1 + \epsilon_2)$. Thus the denotation of the consequence:

$$\Delta, \alpha \vdash_{\mathcal{P}} (\beta, \epsilon)$$

¹Note that this includes the two rules Cut and S which can be derived from the basic set of rules.

$\frac{\alpha \Rightarrow_{\epsilon} \beta \in \Delta}{\Delta, \alpha \vdash_P (\beta, \epsilon)}$	Ax	$\Delta, \alpha \vdash_P (\alpha, 0)$	Ref
$\frac{\Delta, \alpha \vdash_P (\beta, \epsilon_1) \quad \Delta, \alpha \vdash_P (\gamma, \epsilon_2)}{\Delta, \alpha \vdash_P (\beta \wedge \gamma, \epsilon_1 + \epsilon_2)}$	And	$\frac{\Delta, \alpha \vdash_P (\beta, \epsilon_1) \quad \Delta, \alpha \vdash_P (\gamma, \epsilon_2)}{\Delta, \alpha \wedge \beta \vdash_P (\gamma, \frac{\epsilon_2}{1-\epsilon_1})}$	CM
$\frac{\Delta, \alpha \vdash_P (\beta, \epsilon_1) \quad \beta \vdash \gamma}{\Delta, \alpha \vdash_P (\gamma, \epsilon_1)}$	RW	$\frac{\Delta, \alpha \vdash_P (\gamma, \epsilon_1) \quad \vdash \alpha \leftrightarrow \beta}{\Delta, \beta \vdash_P (\gamma, \epsilon_1)}$	LLE
$\frac{\Delta, \alpha \vdash_P (\gamma, \epsilon_1) \quad \Delta, \beta \vdash_P (\gamma, \epsilon_2)}{\Delta, \alpha \vee \beta \vdash_P (\gamma, \epsilon_1 + \epsilon_2)}$	Or	$\frac{\Delta, \alpha \wedge \beta \vdash_P (\gamma, \epsilon_1)}{\Delta, \alpha \vdash_P (\beta \rightarrow \gamma, \epsilon_1)}$	S
$\frac{\Delta, \alpha \wedge \beta \vdash_P (\gamma, \epsilon_1) \quad \Delta, \alpha \vdash_P (\beta, \epsilon_2)}{\Delta, \alpha \vdash_P (\gamma, \epsilon_1 + \epsilon_2)}$	Cut		

Figure 2: The consequence relation \vdash_P .

is that on the basis of what is given in Δ , we can infer $\Pr(\beta \mid \alpha) \geq 1 - \epsilon$.

The rules RW and LLE are a little unusual in that both have antecedents which involve \vdash , which stands for the consequence relation of standard propositional calculus. Thus RW says that you can replace any inference made by \vdash_P with any logical consequence, and LLE says that you can replace anything on the lefthand side of \vdash_P with something that is logically equivalent to it.

This proof system we will call \mathcal{SP} . As with any proof system we are interested in the soundness and completeness of the conclusions which may be drawn using \mathcal{SP} . We define:

Definition 4 A default base is a set of default well-formed formulae.

Definition 5 A basic well-formed formula β is a p-consequence of a default base Δ , conditional on α , iff:

$$\Delta, \alpha \vdash_P (\beta, \epsilon)$$

The value ϵ is known as the *strength* of the consequence. With these definitions, suitable soundness results are easy to obtain. The first relates what can be inferred using \vdash_P to System P:

Theorem 6 For every p-consequence β , conditional on α of a default base Δ , $\alpha \vdash \beta$ is p-entailed by the set of assertions Δ' which is the assertion dual of Δ .

Proof: \mathcal{SP} has a set of proof rules which mirror those of System P, and anything that may be proved using these rules is a p-consequence. Since in (Kraus, Lehmann, & Magidor 1990) it is shown that anything proved using the rules of System P from a given set of conditional assertions Δ' is p-entailed by that set, it follows that any p-consequence of Δ , the default dual of Δ' , is p-entailed by Δ' .

Thus \mathcal{SP} allows us to infer exactly the same things as System P. We also need to show the soundness of the mechanism for propagating the strength of the consequences. This is given by the following:

Theorem 7 The strengths of the p-consequences of a default base are those justified by probability theory.

Proof: The soundness of the propagation of ϵ values with respect to probability theory is proved in (Bourne & Parsons 1998).

Together these two results guarantee that \mathcal{SP} is sound—it will generate conclusions sanctioned by System P with probabilistically correct strengths. Since Kraus *et al.* show that the rules of System P are sufficient to infer all the consequences of System P, the following completeness result is immediate:

Theorem 8 For every $\alpha \vdash \beta$ which is p-entailed by a set of conditional assertions Δ , β is a p-consequence of the default dual of Δ conditional on α .

What this theorem guarantees is any conditional assertion which is p-entailed by a given set of defaults will, when those defaults are translated into the language of \mathcal{SP} , be a consequence of the corresponding set of *dwoffs*. However, this result gives no clue as to the kinds of conclusions we can draw from a given set of *dwoffs*. It does not tell us if a particular p-consequence will be found, it just says that it will be found if its assertion dual is p-entailed.

What we would also like are results which say exactly what kind of conclusions we can infer from some initial set of defaults, which therefore give us some idea of what it is reasonable to establish from some initial set of values, and that is what we consider in the remainder of the paper.

Defining the scope of \mathcal{SP}

We start by considering that we have a set of *simple defaults* of the form $\alpha \Rightarrow_{\epsilon_i} \gamma_i$ which all have the same antecedent. These form a simple default base:

Definition 9 A simple default base for a language \mathcal{L}_S is a default base:

$$\Delta = \bigcup_{i=1, \dots, n} \{\alpha \Rightarrow_{\epsilon_i} \gamma_i\}$$

where α and the γ_i are bwffs in \mathcal{L}_S .

We can think of the consequents of this set of defaults forming a set \mathbf{G} . In general, we have:

Definition 10 *The consequent set of a simple default base Δ is the set \mathbf{G} such that:*

$$\mathbf{G} = \{\gamma_i \mid \{\alpha \Rightarrow_{\epsilon_i} \gamma_i\} \in \Delta\}$$

and this set \mathbf{G} has what we will call an *associated conjunction* Γ which is a conjunction of all the γ_i in \mathbf{G} . Now, applying Ax and CM to $\alpha \Rightarrow_{\epsilon_1} \gamma_1$ and $\alpha \Rightarrow_{\epsilon_2} \gamma_2$, we obtain:

$$\Delta, \alpha \wedge \gamma_1 \vdash_P \left(\gamma_2, \frac{\epsilon_2}{1 - \epsilon_1} \right) \quad (1)$$

Using the same rules on $\alpha \Rightarrow_{\epsilon_1} \gamma_1$ and $\alpha \Rightarrow_{\epsilon_3} \gamma_3$ gives:

$$\Delta, \alpha \wedge \gamma_1 \vdash_P \left(\gamma_3, \frac{\epsilon_3}{1 - \epsilon_1} \right) \quad (2)$$

and combining the latter with (1) will give:

$$\Delta, \alpha \wedge \gamma_1 \wedge \gamma_2 \vdash_P \left(\gamma_3, \frac{\epsilon_3}{1 - \epsilon_1 - \epsilon_2} \right)$$

If we imagine repeating this process it is clear that given Δ we can recursively apply CM to obtain, for some ϵ :

$$\Delta, \alpha \wedge \mathbf{B}' \vdash_P (\gamma_i, \epsilon)$$

for any $\gamma_i \in \mathbf{G}$, and for any \mathbf{B}' which is the associated conjunction of a set \mathbf{B}' such that $\mathbf{B}' \subseteq \mathbf{G}$. Since And makes it possible to build up conjunctions on the consequent side, similar reasoning makes it obvious that recursively applying the And rule to the same initial set of defaults will give:

$$\Delta, \alpha \vdash_P (\Gamma', \epsilon')$$

where Γ' is the associated conjunction of a set \mathbf{G}' and $\mathbf{G}' \subseteq \mathbf{G}$. Clearly, then, if we use both rules together, we can derive conclusions of the form:

$$\Delta, \alpha \wedge \mathbf{B}' \vdash_P (\Gamma', \epsilon'')$$

where \mathbf{B}' and Γ' are the associated conjunctions of sets \mathbf{B}' and \mathbf{G}' such that $\mathbf{B}' \subseteq \mathbf{G}$ and $\mathbf{G}' \subseteq \mathbf{G}$. Note it is possible that $\mathbf{B}' \cap \mathbf{G}' \neq \emptyset$. Thus we have:

Theorem 11 *Given a simple default base with antecedent α and consequent set \mathbf{G} , the consequence relation \vdash_P will generate all consequences:*

$$\Delta, \alpha \wedge \mathbf{B}' \vdash_P (\Gamma', \epsilon)$$

for some ϵ , where \mathbf{B}' and Γ' are the associated conjunctions of sets \mathbf{B}' and \mathbf{G}' such that $\mathbf{B}' \subseteq \mathbf{G}$ and $\mathbf{G}' \subseteq \mathbf{G}$.

Proof: This follows directly from the above discussion.

This result characterises the kind of consequences we can prove using CM and And on a set of simple defaults. It is possible to generalise these results to wider sets

of defaults. Consider that instead of a set of simple defaults, we have, instead, a general set of *conjunctive defaults* of the form $\alpha \wedge \mathbf{B}_i \Rightarrow_{\epsilon_i} \Gamma_i$ where α , as before, is a single proposition and \mathbf{B}_i and Γ_i are conjunctions of propositions. This set of defaults is a conjunctive default base:

Definition 12 *A conjunctive default base for a language \mathcal{L}_S is a default base:*

$$\Delta = \bigcup_{i=1, \dots, n} \{\alpha \wedge \mathbf{B}_i \Rightarrow_{\epsilon_i} \Gamma_i\}$$

where α is a bwff in \mathcal{L}_S , and the \mathbf{B}_i and Γ_i are conjunctions of such bwffs.

The formula α is known as the *base antecedent* of the conjunctive default base, and for each default, \mathbf{B}_i is the *conjunctive antecedent* and Γ_i is the *consequent*. We distinguish two subsets of Δ , the *simple subset*, which is the set of all simple default rules in Δ . Applying the rules Ax and S to any conjunctive default in Δ will give:

$$\Delta, \alpha \vdash_P (\mathbf{B}_i \rightarrow \Gamma_i, \epsilon_i)$$

Now, if we can obtain:

$$\Delta, \alpha \vdash_P (\mathbf{B}_i, \epsilon_j)$$

Applying And will give us

$$\Delta, \alpha \vdash_P (\mathbf{B}_i \wedge (\mathbf{B}_i \rightarrow \Gamma_i), \epsilon_i + \epsilon_j)$$

Now, RW makes it possible to replace any p-consequence with any of its logical consequences. This makes it possible to obtain:

$$\Delta, \alpha \vdash_P (\mathbf{B}_i \wedge \Gamma_i, \epsilon_i + \epsilon_j)$$

and hence:

$$\Delta, \alpha \vdash_P (\varphi_{ij}, \epsilon_i + \epsilon_j)$$

for any $\varphi_{ij} \in \mathbf{B}_i \cup \mathbf{G}_i$ where \mathbf{B}_i and Γ_i are the associated conjunctions of \mathbf{B}_i and \mathbf{G}_i . This gives us:

Theorem 13 *Given a conjunctive default base Δ with base antecedent α , whose simple subset has consequent set \mathbf{G} , the consequence relation \vdash_P will generate all the consequences of the form:*

$$\Delta, \alpha \vdash_P (\gamma^*, \epsilon)$$

for some ϵ , where Δ contains a default $\alpha \wedge \mathbf{B}_i \Rightarrow_{\epsilon_i} \Gamma_i$, Γ_i and \mathbf{B}_i are the associated conjunctions of \mathbf{G}_i and \mathbf{B}_i respectively, $\gamma^* \in \mathbf{B}_i \cup \mathbf{G}_i$ and $\mathbf{B}_i \subseteq \mathbf{G}$.

Proof: From the previous discussion the theorem follows provided that we can infer $\Delta, \alpha \vdash_P (\mathbf{B}_i, \epsilon_j)$. By Theorem 11, this is possible if $\mathbf{B}_i \subseteq \mathbf{G}$.

What this tells us is that we can use the simple subset of a set of defaults to effectively² break down more complex defaults into simple defaults. We can then use these to build up more complex p-consequences just as in Theorem 11:

²We use the term "effectively" since we don't obtain new simple defaults, but p-consequences which we would get from simple defaults.

