# Generalised Inference and Inferential Modelling

R.E. Jennings
Dept of Philosophy
Simon Fraser University
Burnaby, B.C. V5A 1S6

C.W. Chan
Dept of Systems Science
Simon Fraser University
Burnaby, B.C. V5A 1S6

M.J. Dowad
No. 206
9061 Home Street
Burnaby, B.C. V3N 4L2

## Abstract

The authors argue for a generalisation of inference from the standard account in terms of truth preservation to one which countenances preservation of other desirable metalinguistic properties. The development is partly historical and partly analytic. A relational account of preservation is then presented and from this the two notions of inferential structure and inferential model are derived. To illustrate the generality of the relational conception of inference, we show how such structures and models can be realised in the development of a legal advisory system.

## 1. Inference as preservation

### 1.1 Historical illustrations

A conventional dogma takes the essential feature of inference to be the preservation of truth. Thus a system of inference is said to be sound if and only if it never permits the inference of a false conclusion from true premisses. The doctrine is sometimes expressed in terms of *inferential closure* by saying that a system of inference S is sound iff for any set of sentences $\Sigma$, if $\Sigma$ contains only true sentences, then the inferential closure of $\Sigma$ under $\vdash_S$ contains only true sentences.

Historically, this view was inherited from Boole, Shrocder and Frcgc and finds its quintessential refinement in the material conditional of classical logic, whose one virtue, as we tell our bewildered students, is that it preserves truth. The secret vice of material implication is revealed only when there is no truth to preserve, for a false sentence materially implies any sentence whatsoever. That vice has been the pulse of an important portion of twentieth century logical research. C.I. Lewis, whose most important work had tins interest as its inspiration, and to whom we owe the resurgence of interest in modal logic in this century, located the problem with material implication in its cxtensionality. Mis proposed solution lay in introducing a notion of intensional (later called **strict)** implication. Ultimately the focus of his research lay in the investigation of logics of necessity in which, if only the correct system could be identified, a suitable conditional would emerge as the necessary material conditional. In the strongest of his systems, strict implication coincides with provable material implication. Lewis mitigated the difficulty. A merely false sentence does not strictly imply every sentence whatsoever. His success, however, was imperfect, for every necessarily false sentence strictly implies every sentence whatsoever. But he had improved matters, and the improvement can be expressed in the language of preservation. It is this: Strict implication preserves both truth and possibility. Whereas inference which is based upon material implication becomes unprincipled if premisses are false, inference based upon strict implication, even from false premisses will be principled, provided that the premisses are possible. And the licensed conclusions will all be possible, even if some of them arc false.

Although Lewis argued vigorously that unprincipled inference from impossibilities is no vice, others have not been content. So Lewis' advance has spawned new research, this time into systems of implication, and of inference more generally, which permit controlled inferences in the presence of contradiction. Richard Sylvan and Graham Priest have introduced the term 'paraconsistem' as a label for systems of this sort. Our account will be schematic, because our interest in paraconsistent systems here is merely as illustrations of inference understood as preservation.

There have been numerous proposals for contradiction-tolerant inferential strategies. We will mention just three, conveniently partitioned by national boundaries. Although there has been some debate among their proposers, the three need not be seen as competitors, but rather should be looked upon as having different kinds of applications in different circumstances.

I. The American Programmme

There have in fact been many American proposals, more than one of them originating with Nuel Belnap. We mention just one of his. In [Belnap *11]* Belnap imagines a computer required to draw inferences from some stock of sentences to which truth value assignments arrive from more than one source. How in particular, he asks, should the computer exercise its inferential duties when the assignments are in conflict? In this circumstance, he argues, classical inference places no constraint upon what the computer may infer. Our interest lies in Belnap's diagnosis of the difficulty. It lies, he maintains, in the failure of classical inference to preserve truth when both truth and falsity are assigned. The resulting semantic is four-valued, the values being 00, 01, 10, 11.

## II. The Australian Programme

Again, there have been many Australian initiatives, and of them we mention only one, the strongly paraconsistent approach. The barebones of this approach, also dependent on the truth preservative paradigm, is this: II we do not accept that p follows from $q \& \sim q$, then there must be models which have $q \& \sim q$ true at indices where p is false. Again, the locus of the failure of classical inference, according to this programme, is its failure to preserve truth at a crucial point in inferential life. To accommodate the truth of contradictions, the truth condition for negations is altered to:

$$\models^{\mathcal{M}}_{x} \sim\alpha \Leftrightarrow \not\models^{\mathcal{M}}_{x} \alpha.$$

## III. The Canadian Programme

Once again we mention only one of many suggested approaches. In [Schotch and Jennings 80], [Jennings and Schotch 84] and elsewhere, the doctrine that truth is the sole preservation-worthy property is called into doubt. In particular, an automated reasoner receiving data from more than one source must be capable of drawing inferences in some disciplined way even when a datum from one source is inconsistent with a datum from another. The approach advocated there looks for additional metalinguistic properties which are worthy of preservation when not all of the sentences of a premiss set $\Sigma$ can be true. One such property is the *level (of* coherence of $\Sigma$.

Definition

> The coherence level $\ell(\Sigma)$ of a set $\Sigma$ of sentences, is the cardinality of the least partition of $\Sigma$ into coherent subsets. If there is no such partition, then $\ell(\Sigma) = \infty$.

It is a plausible constraint upon an automated reasoner that, in the case in which it receives contradictory data from different sources, it take account that difference of provenance in its reasoning. The requirement that the level of coherence of the inferential closure of $\Sigma$ be no greater than $\ell(\Sigma)$ may be thought of as a way of making that requirement formally precise. The requirement of itself does not define any system of inference, but particular schemes of level preservation do. In [Schotch and Jennings 80] we systematise an inference relation $\vdash$ which is defined:

Definition

$$\Sigma \vdash \alpha \text{ iff } \forall \pi \in \Pi_{\ell(\Sigma)}(\Sigma), \exists c \in \pi: c \vdash \alpha.$$

where $\Pi_{\ell(\Sigma)}(\Sigma)$ is the set of $\ell(\Sigma)$-fold partitions of $\Sigma$.

In [Jennings and Schotch 84] we introduce the inference relation $\Vdash$ which is defined:

Definition

$$\Sigma \Vdash \alpha \text{ iff } \Sigma \vdash \alpha \;\&\; \ell(\Sigma) = \ell(\Sigma \cup \{\alpha\}).$$

One may find a logic attractive without liking every model-theoretic account that happens to characterise it, and it is tempting to conjecture that a semantic account can be given of the Sylvan-Priest logics which characterises those metalinguistic properties in addition to truth that their implication preserves. Specifically, some of the central examples of allegedly true contradictions that Sylvan and Priest want to see dealt with arise out of paradox. On their account, a world in which the Russell Set exists is a world about one of whose inhabitants contradictory sentences will be true. At that world, the Russell set both contains itself as an element and does not. What such a view neglects is that the Russell set *necessarily* contains itself and *necessarily* does not. These properties hold of it in virtue of the definition of the set It is, therefore, the adopted language, and not some feature of a world which permits the Russell set to be defined which gives rise to the contradiction. Accordingly if one's only concern were that *paradoxical contradictions (i.e., those conjoining necessary truth with necessary falsehood* of such a language not give rise to unprincipled inference an implication would suffice that preserved in addition to truth, the language-derivativeness of truth values. Such an implication would license the inference, from paradoxical contradictions, only of sentences also dependent for their truth value upon the language. For a preliminary discussion of such a paradox-tolerant implication, see [Jennings and Johnston S3].[1] It is not claimed of that account that it proves the conjecture. It merely illustrates a point: truth is not the only inferentially preservable property. A system of inference essentially provides procedures by which a set £ of sentences (for example, the set of one's beliefs) having some complex of metalinguistic properties can be unfailingly extended to a larger set $\Sigma'$ having the same complex of properties. By all means, we may regard *truth* as one of the properties to be preserved, but what other properties are to be preserved can depend upon our interests. A practical corollary is that for particular applications, a system of inferential principles should be bespoke, not off the rack.

Some welcome relaxation in the grip of the conventional dogma became evident recently among Situation Semanticists, who have shifted their attention from truth preservation to the conveyance of information.

> ...the study of valid inference as a situated activity shifts attention from *truth preservation* to *information extraction* and *information processing.* [Barwise 84 p. xiv}

> Inference is an activity that attempts to use facts about the world to extract additional information, information implicit in the facts...a sound inference is

1. *Cf.* [Kripke 75] where ihe noli on of the *groundedness* of a sentence is introduced

> Given a sentence A of $\mathcal{L}$, let us define A to be *grounded* if it has a value in the smallest fixed point $\mathcal{L}_{\sigma}$; otherwise *ungrounded. .It A* is grounded, define the *level* of *A* to be the smallest ordinal a such that *A* has a truth value in $\mathcal{L}_{\sigma}$

It is a plausible restriction to place upon an implication relation that it preserve, in addition to *truth,* the *level* of sentences in Kripkc's sense, when that level is infinite.

one which has the structure necessary to serve as a link in an informational chain. [Barwise84]p.56.

Having said all of that, one must still say what distinguishes valid inference from invalid inference in this idiom. Here again, the informational requirement for validity is readily viewed as a preservational constraint. For if we consider what the inference is made from on the one hand (call it P), and what is inferred on the other (call it C), it is evident that the information conveyed by P is a subportion of the information conveyed by P and the property of conveying only subportions of that information in C must be one mark of sound inference on the information-extraction view of what inference is.

## 2. Generalising Preservation

### 2.1 The lesson

Inference as typically construed may be regarded as a sentence-formulating activity constrained by preservational requirements. The so-called paradoxes of material implication and of strict implication make those relations inferential liabilities when the premisses which are the basis for the formulation of new sentences lack the only metalinguistic properties which those implication relations arc capable of preserving. The usefulness of an inference relation depends upon its capacity to preserve significant metalinguistic properties. Which such properties are significant depends upon the particularities of an application. But in general, inference will remain principled provided that at each step it is constrained by the requirement that what is inferred preserves some property that the premisses happen to have, even if what that is may degrade from inferential step to inferential step. A thought experiment might ideally take us from premisses known to be true to conclusions whose truth is discoverable, but if the premisses are possible though suppositional, we will nevertheless expect the conclusions also to be possible, whether true or not. It is a desideratum of *a priori* design of inference procedures for particular fields of application, as for example, in diagnosis, that the preservative constraints of inference be fully specified for every stage of the investigation.

Two final points on the subject of truth preservation: first, a completely satisfactory account of the nature of truth has yet to be given. In logical practice as distinct from metaphysical chat, it is not truth, but the assigned truth value 1 which is preserved. Merely to call this truth is to give no more than a reading of the semantic value, certainly not an account of it. Formally, it is merely the designated value. By contrast, in many real-life cases, the question of the applicability of truth and falsity is at best dubious, and frequently irrelevant. We ask of physical theories, not that they be true, but only that they be consistent with experience and have expressive power and explanatory utility. Yet we still must draw conclusions from them in order to test their utility. In this context, the preservation of truth is a curious and puzzling requirement. For the falsification of a consequence of a theory prompts as often as not a modification of the language of the theory rather than the insertion of a negation into it. It is an oddity of the history of philosophical logic that it should so glutinously cling to the preservation of a commodity of which it professes to understand so little.

Finally, we should note that the preservation even of assigned truth is not strictly a feature of inferential practice, but only a necessary feature of *valid* inference on a particular account of validity. We may think of the principled formulation of hypotheses as an inferential practice, which may involve numerous intermediate steps. But the hypothesis is made before the performance of experiments intended to provide a negative check upon the preservation of truth. The preservation of truth provides a standard by which inferential practices are judged; it is not directly a guide for inferential practice itself. There is nothing new in this claim; it is no more than what we teach to introductory students in insisting that they keep their proof theory and their truth theory distinct.

### 2.2 Inference and Logic

It is, of course, possible to define a language as an element of a particular formal system, and specify rules for the introduction and elimination of the logical constants of the language. But it is also possible to discuss inference independently of any particular language, even to prove results whose effects we normally notice only when the language of premises and conclusions has been given. Thus for example, without knowing anything of the language of premisses or conclusion we may know of a particular inference relation that it satisfies certain *structural* conditions (say, Transitivity, Monotonicity, and Reflexivity), and we know that any inference relation satisfying those structural properties corresponds to a class of valuations. [Scott 72J Again, we may know quite generally that if a relation preserves a property, its converse preserves the absence of that property. *Modus Ponendo Ponens* and *Modus Tollendo Tollens* for material implication are just particular instances of this structural requirement. Another, at least as important for some inferential tasks, is that if causation preserves normality, then the converse of causation preserves abnormality. To take a concrete example from the simplest kind of diagnosis, consider the natural relation between a patient and a thermometer. It is this physical relation which enables the diagnostician to infer from an abnormal length of the column of liquid in the latter to a fever in the former. Thus, although ultimately we will want to say what role our ordinary logical vocabulary is to play in inference, it is nevertheless possible to discuss and commit ourselves to general inferential practices independently of any discussion of or commitment to detailed rules governing what sentences follow from what.

### 2.3 Preservation as Inclusion

It is a simple generalisation from a monadic to a binary account of the preservation of designated value in inference. Define a classical truth relation $\mathcal{T}_\mathcal{V}$ by the following:

$$\mathcal{T}_\mathcal{V} =_{\text{df.}} \{<P_i, P_j> \mid \mathcal{V}(P_i) \leq \mathcal{V}(P_j)\}$$

where v is a classical truth value assignment. An inference relation / between sentences preserves truth iff $I \subsetneq \mathcal{T}_\mathcal{V}$ for every valuation *V.*

Clearly any n-valued function will similarly generate a partial ordering of uniform depth n (that is, having uniform length n chains), but preservational sense can be made even when the restriction to partial orders of uniform depth is dropped. The utility of the approach becomes evident when we consider cases in which we hypothesise causal explanations of events in real physical settings. To be izuidcd in one's causal inferences by chains of mechanical, electrical, or electromagnetic connexions, or spatio-temporal adjacence is to make inferences which preserve, in the sense outlined, the transitive closures and asymmetric decompositions of relevant *natural* relations between states. Such inference as from smoke to fire or from symptoms to disease may plausibly be called natural inference.

The idiom of preservation allows us to represent at a usefully abstract level more particularistic requirements of inference that have already been discussed in the literature. Notable in this connexion is the work of Nau and Reggia |Nau and Reggia, 861:

> Part of the reason why the usual rule-based approach cannot produce diagnoses satisfying [certain of their criteria] is that the information contained in the production rules is simply incorrect. The underlying causal knowledge is not of the form
>
> IF *manifestations* THEN *disorder*
>
> typically found in rule-based expert systems but is instead of the form
>
> IF *disorder* THEN *manifestations.*

Their own approach looks for a principle which will permit an account of diagnostic reasoning expressible in terms of truth preservation.

> Suppose, for example, that a manifestation $m_1$ can be caused by any of the disorders $d_1$, $d_2$, and $d_3$. If $m_1$ is present, then we cannot deduce the presence of $d_1$ nor of $d_2$, nor of $d_3$. The correct action would instead be to postulate $d_1$ $d_2$, and $d_3$ as alternative possible hypotheses for what is causing $m_1$. However, if we further knew that $d_1$, $d_2$, and $d_3$ were the *only* disorders capable of causing $m_1$, then we could correctly deduce that at least one of $d_1$, $d_2$, and $d_3$ must be present. This is a special case of a diagnostic principle which we will call t*he principle of abduction.*

To be sure, if $d_1$, $d_2$, and d^ are not only severally sufficient but also exhaustive of the sufficient conditions for $m_{1t}$ then even in the ordinary way of things, there are grounds for inferring their disjunction from $m_1$. But whatever role such a closed-world-like assumption may play for a particular diagnostic algorithm, it is not necessary in a formal account of the kind of inference required here. And from a representational viewpoint, not altogether desirable. The IF-THEN rules from disorders to manifestations assert *causal* relations. The IF-THEN rules from manifestations to disorders, on the present view, represent possible *inferential* relations which preserve the transitive closure of the converse of the causal relation. That a certain set of possible causes *exhausts* the possible causes is, if it is so, a useful additional piece of information, but it is causal information, and belongs to the level of domain representation, not to inference.

The preservation paradigm of inference promotes a clear demarcation between representation and inference. Its value is best seen when relational distinctions are maintained rather than obscured. It is capable of exploiting the diversity of causal or more generally, productive, relationships (mechanical, electrical, electromagnetic, thermodynamic *etc.)* and non-causal relationships (topographical proximity, comparative probability, comparative reliability of sensors *etc.)* as well.

## 3. Inferential structures and models

An inferential structure consists of two elements:

1. A domain D of objects, where objects are thought of as internally composed of no more than a set of functions and their values. (The assumption is that for each of these a normal range of values can be primitively specified or computed.)

2. A set $\mathfrak{R}$ of intransitive and asymmetric relations defined over D. (These arc the relations which preserve the normality of the values of the functions that constitute the objects of D.)

An inferential model is an inferential structure together with:

3. A partial specification of normal ranges for the functions consisting the objects of D and

4. A preference ordering on $\wp(\mathfrak{R})$.. This specifies what is ideally to be preserved, and what given up in what order in the course of drawing inferences about the objects of D.

The distinction between structure and model is intended to suggest the parallel use of that vocabulary in the model theory of various propositional logics. To cite a familiar example, recall that, in modal correspondence theory, typically what counts as a model for a first order theory, counts as a *model structure* (or 'frame') upon which a model for a modal language is defined. What counts as a model for one language counts as the structure for another. The point of insisting upon the distinction here is to emphasise that the language of inferential structures is assumed to be the language of physics, while the language of normality and of inferential strategy is not assumed to be independent of the values, even aesthetic sensibilities of the user of it. Both what counts as normal values and what is taken to be preservable depends upon the interests that the model is intended to serve. If the structure represents a complex of devices in an industrial operation, and the functions represent transducers, then the normal values will typically be those which represent whatever is acceptable for the ordinary running of the plant, but they may vary depending upon what the plant is being used to produce. And the preference orderings may vary depending, for example, upon whether the inference is a projection of hypothetical values resulting from a contemplated adjustment or a diagnosis of a detected abnormality.

### 3.1 Some examples

Examples of such structures and models that come most readily to mind have to do with machinery wearing sensors and patients wearing thermometers or the electrodes of an ECG. Certainly these conform closely to the description implicit in our account of inferential models. All represent physical connexions which may be described as preserving normality of the values from certain magnitudes to other magnitudes. In the case of ECG use, we sec particularly the point of treating specifications of normality as of the model and not of the structure. Though lead placements do not vary, computerised diagnostic criteria differ from manufacturer to manufacturer and clinical criteria of normality as well as inferential strategies differ from cardiologist to cardiologist. But the conception of inference that we arc advancing goes beyond the standard examples of what is normally regarded as *diagnostics* to all cases which can be brought within the organisational framework of the general account of inferential structure and model outlined. All that is required is that within the proposed domain of application there be discernible objects in natural relationships, about which we make judgements of normality and within which we seek explanations of abnormalities according to the perceived natural relationships among the objects. To illustrate the generality of the conceptions of inferential structure and model, we consider how they might be realised within the realm of legal advisory systems.

Take for example an application of the inferential model to the remoteness domain |Chan, 1988). The issue addressed is that of remoteness in the law of negligence. Specifically, the question is: how far down the chain of causality could a person be liable for a negligent act so that a damage would be recoverable? Negligence usually involves the creation of risk of harm to someone. The injuries of the primary subject of the risk are generally not tcx) remote; the remoteness issue arises when other persons suffer a loss as a result of the negligence to the primary subject [Smith, 1984].

The objectives of the Remoteness Advisor are (1) to establish whether negligence and remoteness is involved and if so, to categorise the negligent act that has led to the damage; (2) once it is determined that remoteness is involved, to retrieve all prior court cases which could be used as precedents in court. The major objective of developing the knowledge based system is the second one. The extent to which the system is able to identify precedent cases that are as similar (or dissimilar) as possible to the one that is to be argued in court is one of the key measures of its values to the lawyer [Chan and Benbasat, 1990].

In applying the inferential structure, the domain D of objects is the class of all cases past and pending or some subclass of cases satisfying a condition of minimal similarity to the present case. The set of relations R among these objects include similarity relations (with respect to various properties or functions of each case such as physical circumstances or nature of the activity involving risk), jurisdictional relations (cases being adjacent if in the same jurisdiction and distant according to their capacity to create precedence), circumstantial relations (cases being distant according to their circumstantial relevance), temporal, citational (cases being in this relation if the one was cited as a precedent in the decision of the other), seniority (whether decided at county court, Queen's bench, superior or supreme court) and so on. For simplicity's sake, let us assume that the cases are all civil.

A binary function for an object may be the result of a case, which takes the value of 1 if the case was decided for the plaintiff, and 0 if it was decided for the defendant. Non-binary functions for an object may be activities in a case that involve risk and settlement details.

So much for the structure. For the model, we must provide two further items, we must say which value counts as the normal value for the function, and provide a preference ordering on subsets of the relations.

In the first particularly we see concretely that what we want inference to preserve depends upon our own interest. For the solicitor consulting such a system, whether 1 counts as the normal value or 0 depends upon whether he or she is representing the plaintiff or the defendant. The judgement in each precedent case may have favoured either the defendant or the plaintiff, which may or may not coincide with the orientation with which the present case is being prepared. Hence for each of the precedent cases also, either 1 or 0 can count as the normal value.

The preference orderings on $\wp(\Re)$ arc also affected by that orientation. But typically, one supposes, the consulting solicitor will hypothesise that there is risk to the client (the abnormal value) and want the system to provide an account of where such risk comes from; then she will hypothesise risk to the opponent and want a similar account. The system will find the source of the hypothesised risk first in the most similar cases already decided within the same jurisdiction and then in cases cited in those in whatever jurisdiction. Or, finding no sufficiendy similar cases in the same jurisdiction, it will look farther abroad, and finding cases there will move to cases cited in them.

## 4. Conclusion

John Austin is said to have rejected the notion that over-simplification is an occupational hazard of the philosopher in favour of the view that it is his occupation. The same may be said of the logician. It is, accordingly, a mistake to take the lessons of logic too literally. The logician speaks of the preservation of truth, and may not himself see, as it is likely that C.I. Lewis did not see, the more natural generality even in his own contributions. The canny consumer will discount a logician's rhapsodic attachment to truth as he would a poet's to beauty. He should consider all that there is to be preserved, and design with a view to preserving it. Insofar as what he is already doing can be described in such terms, even if they are not his own, he is engaged in constructing a system of inference; it is the logician who should be taking notes.

## References


[Barwise 1984] Jon Barwise, "Logic and Information" in *The Situation in Logic,* CSLI Lecture Notes No. 17

[Belnap 1977] N.D. Belnap Jr. "How a computer should think", *Contemporary Aspects of Philosophy. Proceedings of the Oxford International Symposium,* Boston: Oriel Press, 36-51.

[Chan 1988] C.W. Chan, *Knowledge Acquisition : Observations from Three Case Studies,* unpublished Master's thesis, Faculty of Commerce and Business Administration, University of British Columbia, November 1988.

|Chan and Bcnbasat 1990] C. Chan and I. Benbasat, "Case Research in Knowledge Acquisition : Observations and Lessons", forthcoming in *Knowledge Engineering Review.* October 1990.

[Jennings and Johnston 1983] R.E. Jennings and D.K. Johnston, "Paradox-tolerant Logic", *Logique el Analyse*, 1984.

[Jennings and Schotch 1984] R.E. Jennings and P.K. Schotch, "The Preservation of Coherence" *Studia Logica* 43:1/2 (1984).

[Kripke 1975] Saul Kripke, "Outline of a Theory of Truth" *The Journal of Philosophy* 72 (1975), 690 - 716.

[Nau and Reggia 1986] Dana S. Nau and James A. Reggia, "Relationships between Deductive and Abductive Inference in Knowledge-Based Diagnostic Problem Solving*' in Larry Kerschberg, Ed. *Expert Database Systems* (The Benjamin Cummings Publishing Company, Inc. 1986), 549-558.

[Priest 1979] G. Priest, "Logic of Paradox", *Journal of Philosophical Logic* 8, 219-241.

[Routley 1972] R. Routley, "A semantical analysis of implicational system I, and of the first degree of entailment", *Mathematische Annalen* 196, 58-84.

[Schotch and Jennings 1980] P.K. Schotch and R.E. Jennings, "Inference and Necessity" *Journal of Philosophical Logic* 9, August 1980, 327-340.

[Scott 1972] D. Scott, "Completeness and axiomatizability in many-valued logic", *Proceedings of the Tarski Symposium,* AMS (1972), 411-435.