Relevance Logic in Computer Science*

Stuart C. Shapiro
Department of Computer Science
State University of New York at Buffalo
Buffalo, NY 14260
shapiro@cs.buffalo.edu
during 1987–88 on sabbatical at
USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292-6695
shapiro@vaxa.isi.edu

October 12, 1987

§83. Relevance logic in computer science (by Stuart C. Shapiro). Artificial Intelligence (AI) is the branch of Computer Science that uses computational methods to study the kinds of processing that make up human intelligence. One means of pursuing this study is by building computer models (i.e., writing computer programs) that perform intellectual tasks, but recently more and more AI researchers have become concerned with the logical foundations of such processes. It is not surprising, then, that a group of AI researchers have been attracted to relevance logic as an appropriate foundation for human and computer reasoning systems.

We can categorize the uses of relevance logic that have been suggested in the AI literature in two groups: those that have made use of, or modified, **R**'s proof theory to design AI reasoning systems; those that have stressed the four-valued semantics of **R**.

§83.1. Use of the proof theory. One of the first suggestions that **R** would be useful for Artificial Intelligence reasoning systems was by Shapiro and Wand 1976. Their first point is that, "In a question-answering system, an implication has imperative as well as declarative content: an implication ought to be a useful inference rule" (Shapiro and Wand 1976 p.8, see also Hewitt 1972). In this view, an implication, such as $A \rightarrow B$, is also treated as a rule that says, "if

*To appear as Chapter XII §83 of ENTAILMENT, VOLUME II by Alan R. Anderson, Nuel D. Belnap, Jr. and J. Michael Dunn, Princeton University Press you want to know the truth of B, check the truth of A." If A is irrelevant to B (the worst case being that A is a contradiction), this is not a reasonable rule.

Shapiro and Wand modify the notation of $\mathbf{FR}_{\rightarrow\&}$(§27.2) to eliminate the sub-proof structure. They suggest that the knowledge base (KB) of a reasoning system be considered to contain "assertions of the form $\langle A, \varphi, \alpha\rangle$, where A is some formula, $\varphi \in \{0, 1\}$, and α is a set." (The angle brackets were not in the original.) The rules of inference they present are:

hyp: $\langle A, 0, \{k\}\rangle$ may be added to the KB as long as $\{k\}$ is a singleton set such that no assertion of the form $\langle B, 0, \{k\}\rangle$ is already in KB.

add: $\langle A, 0, \{k\}\rangle$ may be removed from KB and replaced by $\langle A\&B, 0, \{k\}\rangle$.

$\rightarrow$**E:** If $\langle A, \varphi, \alpha\rangle \in$ KB and $\langle A\rightarrow B, \xi, \beta >\rangle \in$ KB, then $\langle B, 1, \alpha \cup \beta\rangle$ may be added to KB.

$\rightarrow$**I:** If $\langle A, 0, \{k\}\rangle \in$ KB and $\langle B, \varphi, \beta\rangle \in$ KB and $k \in \beta$, then $\langle A\rightarrow B, 1, \beta - \{k\}\rangle$ may be added to KB.

&**E:** If $\langle A\&B, \varphi, \alpha\rangle \in$ KB, then $\langle A, 1, \alpha\rangle$ may be added to KB and $\langle B, 1, \alpha\rangle$ may be added to KB.

&**I:** If $\langle A, \varphi, \alpha\rangle \in$ KB and $\langle B, \xi, \alpha\rangle \in$ KB, then $\langle A\&B, 1, \alpha\rangle$ may be added to KB.

Using the later terminology of Martins and Shapiro (see §83.1.1 below), we may refer to φ as the *origin tag* and α as the *origin set* of the assertion. All assertions whose origin tags are 0 are hypotheses entered into the KB by some user. All assertions of the form $\langle A, 1, \alpha\rangle$ are derived assertions which have been derived under the set of assumptions $\{\langle B, 0, \{k\}\rangle \mid k \in \alpha\}$.

Shapiro and Wand discuss the use of their system for using hypothetical reasoning to derive new rules:

> Consider a universe of discourse, α, and the new, hypothetical world produced by assuming $\langle P, 0, \{p\}\rangle$. If, in this hypothetical world, we can derive $\langle Q, 1, \alpha \cup \{p\}\rangle$, we can then derive the new deduction rule $\langle P\rightarrow Q, 1, \alpha\rangle$ in the original universe by use of $\rightarrow$I. This is a productive rule in the sense that if we later learn that $\langle P, \varphi, \beta\rangle$ is true, we can derive $\langle Q, 1, \alpha \cup \beta\rangle$. ... The rules of $\mathbf{FR}_{\rightarrow\&}$are precisely the right ones to ensure that any derived [rules] are in fact relevant to the hypothetical situation. (pp. 16–17)

Shapiro and Wand also use origin sets to define the notion of a *context*: "A context is a set γ and is said to contain the set of assertions $\{\langle A, \varphi, \alpha\rangle \mid \alpha \subseteq \gamma\}$" (p. 15). They point out, in the light of suggestions made by Shapiro 1971 (pp. 107–109), that the rules of $\mathbf{FR}_{\rightarrow\&}$can be used to keep contradictory contexts separate and that origin sets can be used to discover and remove the source of contradictions if any arise during reasoning. These ideas were subsequently key ideas in Belief Revision systems and Assumption-based Truth Maintenance systems (see below).

§83.1.1. SWM. The work of Shapiro and Wand 1976 was continued by Martins and Shapiro, whose work is described in a series of papers (Martins 1983; Martins and Shapiro 1981, 1983, 1984, 1986a, 1986b, 1986c, 198+a, 198+b; see also Martins 1987). The logic developed by Martins and Shapiro, called **SWM**, operates on *supported wffs*, which are expanded versions of the assertion triples of Shapiro and Wand, and which we shall here refer to as *assertions*. An **SWM** assertion $\mathcal{A}$ is a quadruple, $\langle A, \tau, \alpha, \rho \rangle$, where A is a wff called the wff of $\mathcal{A}$, τ is a member of the set {hyp, der, ext}, and is called the *origin tag* (OT) of $\mathcal{A}$, α is a set of wffs and is called the *origin set* (OS) of $\mathcal{A}$, and ρ is a set of sets of wffs and is called the *restriction set* (RS) of $\mathcal{A}$. If $\mathcal{A}$ is the assertion $\langle A, \tau, \alpha, \rho \rangle$, the functions wff, ot, os, and rs are defined so that $\text{wff}(\mathcal{A}) = A$, $\text{ot}(\mathcal{A}) = \tau$, $\text{os}(\mathcal{A}) = \alpha$, and $\text{rs}(\mathcal{A}) = \rho$. (The notation in most of Martins and Shapiro's papers differs slightly from that given here.)

A set of hypotheses, α, is known to be inconsistent as soon as an assertion is derived whose wff is a contradiction and whose os is α, or as soon as two assertions, $\mathcal{A}_1$and $\mathcal{A}_2$, are derived for which $\text{wff}(\mathcal{A}_1) = \sim\text{wff}(\mathcal{A}_2)$ and $\alpha = \text{os}(\mathcal{A}_1) \cup \text{os}(\mathcal{A}_2)$. The rules of inference of **SWM** guarantee that for every derived assertion $\mathcal{A}$, $\text{os}(\mathcal{A})$ contains every hypothesis wff that was used in $\mathcal{A}$'s derivation and only those hypothesis wffs and that $\text{rs}(\mathcal{A})$ contains every set of hypothesis wffs known to be inconsistent with $\text{os}(\mathcal{A})$. The rules of inference do not allow the derivation of any assertion $\mathcal{A}$ for which $\text{os}(\mathcal{A})$ would be a set of hypothesis wffs already known to be inconsistent.

§83.1.1.1. Rules of inference of SWM. To make the rules of inference of **SWM** easier to state, several functions are defined.

First, to prevent any use of a context already known to be inconsistent, the rules require all parent assertions to be combinable, as defined by:

$$\text{Combine}(\mathcal{A}_1, \mathcal{A}_2) = \forall r \in \text{rs}(\mathcal{A}_1) : r \not\subseteq \text{os}(\mathcal{A}_2) \;\&\; \forall r \in \text{rs}(\mathcal{A}_2) : r \not\subseteq \text{os}(\mathcal{A}_1)$$

The OT 'hyp' tags assertions that are hypotheses; 'der' tags assertions that are normal derived assertions; 'ext' tags derived assertions whose later use is restricted. To prevent irrelevancies from arising, the rule of And Introduction must be restricted to parent assertions with the same OS. However, if $\mathcal{A}_1 = \langle A, t_1, o_1, r_1 \rangle$ and $\mathcal{A}_2 = \langle B, t_2, o_2, r_2 \rangle$ are two assertions, it intuitively seems unobjectionable for a reasoner to assert $\mathcal{A}_3 = \langle A\&B, t_3, o_1 \cup o_2, r_3 \rangle$. There is, in fact nothing wrong with this as long as certain rules are prevented from acting on $\mathcal{A}_3$ or any of its descendants. For this reason, $\mathcal{A}_3$ and all its descendants are given an OT of 'ext'. The function Λ correctly computes OTs:

$$\Lambda(a,b) = \begin{cases} \text{ext} & \text{if } a = \text{ext or } b = \text{ext} \\ \text{der} & \text{otherwise} \end{cases}$$

The final four functions are used in the computation of RSs to insure that no two sets in an RS overlap, and that all are disjoint with the OS.

$$\begin{aligned} \psi(R,O) &= \{\alpha \mid (\alpha \in R \,\&\, \alpha \cap O = \phi) \vee (\exists \beta \in R)[\beta \cap O \neq \phi \,\&\, \alpha = \beta - O]\} \\ \sigma(R) &= \{\alpha \in R \mid \sim(\exists \beta)(\beta \neq \alpha \,\&\, \beta \in R \,\&\, \beta \subset \alpha)\} \\ \mu(\{r_1,\ldots,r_m\},\{o_1,\ldots,o_n\}) &= \sigma(\psi(r_1 \cup \cdots \cup r_m, o_1 \cup \cdots \cup o_n)) \\ \int(O) &= \mu(\{r \mid \exists H \in O : r = \text{rs}(H)\}, \{o \mid \exists H \in O : o = \text{os}(H)\}) \end{aligned}$$

Given these functions, the rules of inference of **SWM** are:

Hypothesis (Hyp): For any wff A and sets of wffs $R_1,\ldots,R_n (n \geq 0)$, such that $\forall r \in \{R_1,\ldots,R_n\} : r \cap \{A\} = \phi$ and $\forall r,s \in \{R_1,\ldots,R_n\} : r \not\subset s$, we may add the assertion $\langle A, \text{hyp}, \{A\}, \{R_1,\ldots,R_n\}\rangle$ to the knowledge base, provided that A has not already been introduced as a hypothesis.

Negation Introduction ($\sim$I):

From $\langle A, t_1, o, r\rangle$,
$\langle \sim A, t_2, o, r\rangle$,
and $\{H_1,\ldots,H_n\} \subset o$,
infer $\langle \sim(H_1 \& \cdots \& H_n), \Lambda(t_1,t_2), o - \{H_1,\ldots,H_n\}, \int(o - \{H_1,\ldots,H_n\})\rangle$.

From $\mathcal{A}_1 = \langle A, t_1, o_1, r_1\rangle$,
$\mathcal{A}_2 = \langle \sim A, t_2, o_2, r_2\rangle$,
$o_1 \neq o_2$,
Combine($\mathcal{A}_1$, $\mathcal{A}_2$),
and $\{H_1,\ldots,H_n\} \subset (o_1 \cup o_2)$,
infer $\langle \sim(H_1 \& \cdots \& H_n), \text{ext}, (o_1 \cup o_2) - \{H_1,\ldots,H_n\}, \int((o_1 \cup o_2) - \{H_1,\ldots,H_n\})\rangle$.

This rule may be applied before URS (see below).

Negation Elimination ($\sim$E): From $\langle \sim\sim A, t, o, r\rangle$, infer $\langle A, \Lambda(t,t), o, r\rangle$.

And Introduction (&I):

From $\langle A, t_1, o, r\rangle$
and $\langle B, t_2, o, r\rangle$,
infer $\langle A \& B, \Lambda(t_1,t_2), o, r\rangle$.

From $\mathcal{A}_1 = \langle A, t_1, o_1, r_1\rangle$,
$\mathcal{A}_2 = \langle B, t_2, o_2, r_2\rangle$,
$o_1 \neq o_2$,
and Combine($\mathcal{A}_1$, $\mathcal{A}_2$),
infer $\langle A \& B, \text{ext}, o_1 \cup o_2, \mu(\{r_1, r_2\}, \{o_1, o_2\})\rangle$.

And Elimination (&E):

From $\langle A\&B, t, o, r\rangle$,
and $t \neq \text{ext}$,
infer either $\langle A, \text{der}, o, r\rangle$
or $\langle B, \text{der}, o, r\rangle$ or both.

Or Introduction (truth functional) (∨I):

From $\langle A, t, o, r\rangle$,
infer either $\langle A \vee B, \Lambda(t,t), o, r\rangle$
or $\langle B \vee A, \Lambda(t,t), o, r\rangle$, for any wff B.

Or Introduction (intensional) (⊕I):

From $\langle {\sim}A \rightarrow B, t_1, o, r\rangle$
and $\langle {\sim}B \rightarrow A, t_2, o, r\rangle$,
infer $\langle A \oplus B, \Lambda(t_1, t_2), o, r\rangle$.

Or Elimination (⊕E):

a. From $\mathcal{A}_1 = \langle A \oplus B, t_1, o_1, r_1\rangle$,
$\mathcal{A}_2 = \langle {\sim}A, t_2, o_2, r_2\rangle$,
and Combine($\mathcal{A}_1$, $\mathcal{A}_2$),
infer $\langle B, \Lambda(t_1, t_2), o_1 \cup o_2, \mu(\{r_1, r_2\}, \{o_1, o_2\})\rangle$.
From $\mathcal{A}_1 = \langle A \oplus B, t_1, o_1, r_1\rangle$,
$\mathcal{A}_2 = \langle {\sim}B, t_2, o_2, r_2\rangle$,
and Combine($\mathcal{A}_1$, $\mathcal{A}_2$),
infer $\langle A, \Lambda(t_1, t_2), o_1 \cup o_2, \mu(\{r_1, r_2\}, \{o_1, o_2\})\rangle$.

b. From $\mathcal{A}_1 = \langle A \oplus B, t_1, o_1, r_1\rangle$,
$\mathcal{A}_2 = \langle A \rightarrow C, t_2, o_2, r_2\rangle$,
$\mathcal{A}_3 = \langle B \rightarrow C, t_3, o_2, r_2\rangle$,
and Combine($\mathcal{A}_1$, $\mathcal{A}_2$),
infer $\langle C, \Lambda(t_1, \Lambda(t_2, t_3)), o_1 \cup o_2, \mu(\{r_1, r_2\}, \{o_1, o_2\})\rangle$.

Implication Introduction (→I):

From $\langle B, \text{der}, o, r\rangle$
and any hypothesis $H \in o$,
infer $\langle H \rightarrow B, \text{der}, o - \{H\}, \int(o - \{H\})\rangle$.

Modus Ponens—Implication Elimination, Part 1 (MP):

From $\mathcal{A}_1 = \langle A, t_1, o_1, r_1\rangle$,
$\mathcal{A}_2 = \langle A \rightarrow B, t_2, o_2, r_2\rangle$,
and Combine($\mathcal{A}_1$, $\mathcal{A}_2$),
infer $\langle B, \Lambda(t_1, t_2), o_1 \cup o_2, \mu(\{r_1, r_2\}, \{o_1, o_2\})\rangle$.

Modus Tollens—Implication Elimination, Part 2 (MT):

From $\mathcal{A}_1 = \langle A \rightarrow B, t_1, o_1, r_1\rangle$,
$\mathcal{A}_2 = \langle {\sim}B, t_2, o_2, r_2\rangle$,
and Combine($\mathcal{A}_1$, $\mathcal{A}_2$)
infer $\langle {\sim}A, \Lambda(t_1, t_2), o_1 \cup o_2, \mu(\{r_1, r_2\}, \{o_1, o_2\})\rangle$.

Updating of Restriction Sets (URS): From $\langle A, t_1, o_1, r_1\rangle$, and $\langle \sim A, t_2, o_2, r_2\rangle$, we *must* replace each hypothesis $\langle H, \text{hyp}, \{H\}, R\rangle$ such that $H \in (o_1 \cup o_2)$ by $\langle H, \text{hyp}, \{H\}, \sigma(R \cup ((o_1 \cup o_2) - H))\rangle$. Furthermore, we *must* also replace every assertion $\langle F, t, o, r\rangle$ ($t = \text{der}$ or $t = \text{ext}$) such that $o \cap (o_1 \cup o_2) \neq \phi$ by $\langle F, t, o, \sigma(r \cup \{(o_1 \cup o_2) - o\})\rangle$. However, the rule of $\sim I$ may be applied before the restriction sets are updated.

$\forall$ Introduction ($\forall$ I): From $\langle B(t), \text{der}, o \cup \{A(t)\}, r\rangle$, in which $A(t)$ is a hypothesis that uses a term (t) never used in the system prior to A's introduction, and t is not in o or r, infer $\langle \forall(x)[A(x) \rightarrow B(x)], \text{der}, o, \int(o)\rangle$. (According to this rule of inference, the universal quantifier can only be introduced in the context of an implication. This is not a drawback, as it may seem at first, since the role of the antecedent of the implication ($A(x)$) is to define the type of objects that are being quantified.)

$\forall$ Elimination—Universal Instantiation ($\forall$ E):

From	$\mathcal{A}_1 = \langle \forall(x)[A(x) \rightarrow B(x)], t_1, o_1, r_1\rangle$,
	$\mathcal{A}_2 = \langle A(c), t_2, o_2, r_2\rangle$,
and	Combine($\mathcal{A}_1$, $\mathcal{A}_2$),
where	c is any individual symbol,
infer	$\langle A(c) \rightarrow B(c), \Lambda(t_1, t_2), o_1 \cup o_2, \mu(\{r_1, r_2\}, \{o_1, o_2\})\rangle$.

$\exists$ Introduction ($\exists$ I):

From	$\langle A(c), t, o, r\rangle$, where c is an individual constant,
infer	$\langle \exists(x)[A(x)], \Lambda(t, t), o, r\rangle$.

$\exists$ Elimination ($\exists$ E):

From	$\langle \exists(x)[A(x)], t, o, r\rangle$
infer	$\langle A(c), \Lambda(t, t), o, r\rangle$
	where c is any individual constant that was never used before.

The rules of ~I (part 1), &I (part 1), and $\oplus$I are only applicable to assertions that have the same OS and the same RS. This condition is not as constraining as it may seem at first glance, since Martins and Shapiro prove that if two assertions have the same OS, then they also have the same RS. In fact, this justifies a different view of the database of assertions. One may think of the KB as containing a set of wffs. For every wff A and every assertion $\mathcal{A}$ in which A = wff($\mathcal{A}$), A is a wff of type ot($\mathcal{A}$) in the context os($\mathcal{A}$) and in every context γ such that os($\mathcal{A}$) $\subseteq \gamma$. Two contexts α and β are known to be inconsistent if, in the previous way of thinking, there is an assertion $\mathcal{A}$ such that $\alpha = \text{os}(\mathcal{A})$ & $\beta \in \text{rs}(\mathcal{A})$ or $\beta = \text{os}(\mathcal{A})$ & $\alpha \in \text{rs}(\mathcal{A})$. The rules of inference of **SWM** apply with the obvious modifications. However, Martins and Shapiro show that if one restricts the reasoner to a consideration of only wffs in a single context, not known to be inconsistent, the Combine test need never be made, and if a new contradiction is uncovered within the context, the removal of any wff in the OS of the contradictory assertion will restore the context to the status of not being known to be inconsistent. This is the logical basis for assumption-based truth maintenance, or belief revision (Martins 1987; Martins and Shapiro 1981, 1984, 1986a, 198+a, 198+b).

§83.1.1.2 Example. The main advantages of SWM are that the OSs show precisely the hypotheses required to derive each assertion, so that when a contradiction is found, no innocent hypothesis will be blamed, and that once a set of hypotheses is found to be contradictory, reasoning will no longer occur in the context formed by that set of hypotheses. In actual computer reasoning systems based on SWM, the user may explicitly decide to reason in a context known to be inconsistent.

As an example of SWM, we show the derivation that the existence of the Russell set is self-inconsistent.

1. $\langle \exists(s)[Set(s) \& \forall(x)[Set(x) \rightarrow ((x \in s \rightarrow \sim(x \in x)) \& (\sim(x \in x) \rightarrow x \in s))]], hyp, \{1\}, \{\}\rangle$ Hyp

2. $\langle Set(R) \& \forall(x)[Set(x) \rightarrow ((x \in R \rightarrow \sim(x \in x)) \& (\sim(x \in x) \rightarrow x \in R))], der, \{1\}, \{\}\rangle$ ∃E 2

3. $\langle Set(R), der, \{1\}, \{\}\rangle$ &E 2

4. $\langle \forall(x)[Set(x) \rightarrow ((x \in R \rightarrow \sim(x \in x)) \& (\sim(x \in x) \rightarrow x \in R))], der, \{1\}, \{\}\rangle$ &E 2

5. $\langle ((R \in R \rightarrow \sim(R \in R)) \& (\sim(R \in R) \rightarrow R \in R)), der, \{1\}, \{\}\rangle$ ∀E 4,3

6. $\langle (R \in R) \rightarrow \sim(R \in R), der, \{1\}, \{\}\rangle$ &E 5

7. $\langle R \in R, hyp, \{7\}, \{\}\rangle$ Hyp

8. $\langle \sim(R \in R), der, \{1, 7\}, \{\}\rangle$ MP 7, 6

9. $\langle \sim(R \in R), ext, \{1\}, \{\{7\}\}\rangle$ ~I 7, 8

URS is now required by the presence of 7 and 8. Every assertion with an OS of $\{1\}$ now has $\{7\}$ added to its RS, and every assertion with an OS of $\{7\}$ now has $\{1\}$ added to its RS. The two hypotheses are now:

1′. $\langle \exists(s)[Set(s) \& \forall(x)[Set(x) \rightarrow ((x \in s \rightarrow \sim(x \in x)) \& (\sim(x \in x) \rightarrow x \in s))]], hyp, \{1\}, \{\{7\}\}\rangle$ Hyp; URS 7, 8

7′. $\langle R \in R, hyp, \{7\}, \{\{1\}\}\rangle$ Hyp; URS 7, 8

Other revised assertions will be shown when and only when they are about to be used.

9′. $\langle \sim(R \in R), ext, \{1\}, \{\{7\}\}\rangle$ ~I 7, 8; URS 7, 8

5′. $\langle(((R \in R \rightarrow \sim(R \in R))\&(\sim(R \in R) \rightarrow R \in R)), der, \{1\}, \{\{7\}\}\rangle$ ∀E 4,3; URS 7, 8

10. $\langle \sim(R \in R) \rightarrow R \in R, der, \{1\}, \{\{7\}\}\rangle$ $\& E5'$

11. $\langle R \in R, ext, \{1\}, \{\{7\}\}\rangle$ MP 9′, 10

URS is now required by the presence of 11 and 9′. In this case, $o_1 = o_2 = o_1 \cup o_2 = \{1\}$, so hypothesis 1 becomes:

1″. $\langle \exists(s)[Set(s)\&\forall(x)[Set(x) \rightarrow ((x \in s \rightarrow \sim(x \in x))\&(\sim(x \in x) \rightarrow x \in s))]], hyp, \{1\}, \{\{\}\}\rangle$ Hyp; URS 7, 8; URS 11, 9′

The existence of the empty set in the RS of 1″ means that 1″ is self-inconsistent and not combinable with any other assertion. Within the context of the hypothesis {1} we may reason about the Russell set, but that hypothesis may not be combined with any other, so the contradiction has been isolated.

§83.1.2. Implementations. Martins and Shapiro implemented a computer reasoning system, SNeBR (Martins 1983; Martins and Shapiro 1983, 1984, 1986c, 198+b), based on a version of **SWM** for the non-standard propositional connectives of SNePS, the Semantic Network Processing System (Shapiro 1979, Shapiro and Rapaport 1987).

Ohlbach and Wrightson 1984 used the Markgraf Karl Refutation Procedure (Raph 1983), a resolution based theorem prover, to show that

$$(A \rightarrow (B \rightarrow B)) \rightarrow (A \rightarrow (A \rightarrow (B \rightarrow B)))$$

follows from the axioms of $\mathbf{T}_{\rightarrow}$ (see §8.13).

Thistlewaite and McRobbie have implemented KRIPKE, an **R** based automatic theorem prover (see Malkin 1987 and Thistlewaite, McRobbie and Meyer 198+).

Brachman, Gilbert, and Levesque 1985 mention their intention to implement an inference mechanism based on a relevance logic as part of the KRYPTON knowledge representation/reasoning system.

§83.2. Use of the four-valued semantics of R. Belnap 1975, 1977 was the first to suggest that the four-valued semantics of **R** make it a useful model for computer reasoning systems. A revised version of these papers appears as §81 of this volume, so the discussion will not be repeated here beyond noting the meaning, in a computer reasoning context, of the four values. Most database management systems assume what in Artificial Intelligence has been called the Closed World Assumption (Reiter 1978). This is that the database contains all relevant true information, so whatever information is not in the database is false. The Closed World Assumption is unreasonable for any reasoning system that might learn new facts. For such a system, false assertions as well as true assertions may be explicitly stored in the database. An assertion that is not stored in the database as either true or false must only be assumed to be unknown. True, false, and unknown are three of the four truth values. The fourth, both, is used if more than one informant put information into the database and one informant said that an assertion was true while another said that it was false. Perhaps a single informant at one time said that the assertion was true, and at another time that it was false. Perhaps the actual situation changed, so that an assertion that was true at one time later became false, or maybe a simple error was made in entering information, and this led to a contradiction. Of course, an assertion's having a truth value of both indicates some problem to be resolved in the database, unless it is true in one context and false in another. However, until the problem is resolved, the use of **R** can prevent the contradiction from polluting the database with every possible conclusion (derivable from a contradiction in standard logics).

The Closed World Assumption is also unreasonable for a database management system or reasoning system that, for reasons of speed, must produce information before it can develop all the implications of its stored data. Such a system might not find some information, not because it was not in or implied by its database, but because it was not given enough time (or other resources). Call the information retrievable by such a system within its resource limits its *explicit beliefs* and all the information it could retrieve given an arbitrary amount of resources its *implicit beliefs*. Semantics for relevance logics appropriate for the set of explicit beliefs of such systems have been discussed by Levesque 1984a, 1984b; Fagin and Halpern 1985, 198+; Frisch 1985, 1986; and Lakemeyer 1986 (see also Levesque 1986).

Lakemeyer 1987 extends the model of Levesque 1984b to one that an agent can use to hold meta-beliefs (beliefs about its own beliefs) and reason about them efficiently.

Mitchell and O'Donnell 1986 (see also O'Donnell 1985) are particularly interested in the use of **R** for database systems that may have errors in the data. They present two versions of realizability semantics for relevance logic, show soundness for the first, and soundness and completeness over a nonstandard set of models for the second.

Patel-Schneider 1985a, 1985b presents a decidable variant of relevance logic including quantifiers as an appropriate logic for reasoning systems.

Allowing unknown as a truth value invites one to consider inferences based on lack of knowledge; e.g., if P is unknown conclude Q. The Closed World Assumption then amounts to $\forall P,$ if P is unknown then $\sim P$, but less overriding rules are useful for the sort of default reasoning people seem to engage in. (The favorite example in Artificial Intelligence is if x is a bird and it is not known that x doesn't fly, then x does fly.) If a previously unknown datum, used for one of these lack-of-knowledge inferences, is later learned to be false, the earlier conclusion may no longer be justified. This phenomenon, of once valid conclusions becoming invalid due to the gaining of knowledge, has been termed non-monotonicity, and several non-monotonic logics have been proposed as the foundation of such reasoning (see Perlis 1987). Sandewall 1985a, 1985b discusses a functional approach to non-monotonic logic with the four-valued semantics of **R**.

A particular kind of database used in Artificial Intelligence is the inheritance net (see Touretzky 1987). Thomason, Horty, and Touretzky 1986 discuss inheritance nets in which nodes represents either individuals or kinds, and in which there are two kinds of links. The link $p \rightarrow q$ means that p is a q (or all p's without exception are q's), and the link $p \not\rightarrow q$ means that p is not a q (or p's are not q's, again without exception). They give a proof theory and a model theory for inference in these nets, show the soundness and completeness of the proof theory relative the the model theory, and show that the four-valued semantics of **R** is an appropriate interpretation of this logic.

ADDITIONAL TWO PARAGRAPHS PROVIDED BY BELNAP & DUNN HERE

Bibliography

BELNAP, NUEL D., JR.

1975. *How a computer should think.* **Contemporary aspects of philosophy: proceedings of the oxford international symposium.** Oriel Press, Stocksfield, England, pp. 30–56.

1977. *A useful four-valued logic.* **Modern Uses of Multiple-Valued Logic,** ed. G. Epstein and J. M. Dunn, Dordecht (Reidel) pp. 8–37.

BRACHMAN, RONALD J., GILBERT, VICTORIA P. AND LEVESQUE, HECTOR J.

1985. *An essential hybrid reasoning system: knowledge and symbol level accounts of KRYPTON.* **Proceedings of the ninth international joint conference on artificial intelligence.** Morgan Kaufmann, Los Altos, CA, pp. 532–539.

FAGIN, RONALD AND HALPERN, JOSEPH Y.

1985. *Belief, awareness, and limited reasoning: preliminary report.* **Proceedings of the ninth international joint conference on artificial intelligence.** Morgan Kaufmann, Los Altos, CA, pp. 491–501.

198+. *Belief, awareness, and limited reasoning.* **Artificial intelligence.**

FRISCH, ALAN M.

1985. *Using model theory to specify AI programs.* **Proceedings of the ninth international joint conference on artificial intelligence.** Morgan Kaufmann, Los Altos, CA, pp. 148–154.

1986. **Knowledge retrieval as specialized inference.** University of Rochester doctoral dissertation.

HEWITT, CARL E.

1972. **Description and theoretical analysis (using schemata) of PLANNER.** A.I. Report TR-258, MIT, Cambridge, MA.

LAKEMEYER, GERHARD

1986. *Steps towards a first-order logic of explicit and implicit belief.* **Theoretical aspects of reasoning about knowledge,** ed. J. Y. Halpern, Los Altos, CA (Morgan Kaufmann) pp. 325–340.

1987. *Tractable meta-reasoning in propositional logics of belief.* **Proceedings of the tenth international joint conference on artificial intelligence.** Morgan Kaufmann, Los Altos, CA, pp. 402–408.

LEVESQUE, HECTOR J.

1984a. *A logic of implicit and explicit belief.* **Proceedings of the national conference on artificial intelligence.** Morgan Kaufmann, Los Altos, CA, pp. 198–202.

1984b. **A logic of implicit and explicit belief.** Technical Report No. 32, Fairchild Laboratory for Artificial Intelligence.

1986. *Knowledge representation and reasoning.* **Annual review of computer science,** ed. J. F. Traub, B. J. Grosch, B. W. Lampson and N. J. Nilsson, Palo Alto, CA (Annual Reviews Inc.) pp. 255–287.

See Brachman, Gilbert and Levesque.

MALKIN, PETER

1987. *Automated reasoning project: activities and accomplishments.* **The computers and philosophy newsletter,** vol. 2, pp. 25–30.

MARTINS, JOÃO P.

1983. **Reasoning in multiple belief spaces.** doctoral dissertation, Technical Report No. 203, Department of Computer Science, SUNY at Buffalo, Buffalo, NY.

1987. *Belief revision.* **Encyclopedia of artificial intelligence,** ed. S. C. Shapiro, New York (John Wiley & Sons) pp. 58–62.

See Martins and Shapiro.

MARTINS, JOÃO P. AND SHAPIRO, STUART C.

1981. **Belief revision system based on relevance logic and heterarchical contexts.** Technical Report No. 175, Department of Computer Science, SUNY at Buffalo, Buffalo, NY.

1983. *Reasoning in multiple belief spaces.* **Proceedings of the eighth international joint conference on artificial intelligence.** Los Altos, CA (Morgan Kaufmann), pp. 370–373.

1984. *A model for belief revision.* **Non-monotonic reasoning workshop.** Menlo Park, CA (The American Association for Artificial Intelligence), pp. 241–294.

1986a. *Theoretical foundations for belief revision.* **Theoretical aspects of reasoning about knowledge,** ed. J. Y. Halpern, Los Altos, CA (Morgan Kaufmann) pp. 383–398.

1986b. *Hypothetical reasoning.* **Applications of artificial intelligence to engineering problems: proceedings of the 1st international conference.** Berlin (Springer-Verlag), pp. 1029–1042.

1986c. *Belief revision in SNePS.* **Proceedings of the sixth Canadian conference on artificial intelligence.** Montréal (Presses de l'Université du Québec), pp. 230–234.

198+a. *A logic for belief revision.* **The journal of philosophical logic.**

198+b. *A model for belief revision.* **Artificial intelligence.**

MITCHELL, JOHN C. AND O'DONNELL, MICHAEL J.

1986. *Realizability semantics for error-tolerant logics (preliminary version).* **Theoretical aspects of reasoning about knowledge,** ed. J. Y. Halpern, Los Altos, CA (Morgan Kaufmann) pp. 363–381.

O'DONNELL, MICHAEL J.

1985. **A type-theoretic foundation for epistemic and relevance logic.** JHU/EECS-85/02, Department of Electrical Engineering and Computer Science, Johns Hopkins University, Baltimore, MD.

See Mitchell and O'Donnell.

OHLBACH, HANS-JÜRGEN AND WRIGHTSON, GRAHAM

1984. *Solving a problem in relevance logic with an automated theorem prover.* **7th International conference on automated deduction; Lecture notes in computer science 170,** ed. R. E. Shostack, Berlin (Springer-Verlag) pp. 496–508.

PATEL-SCHNEIDER, PETER F.

1985a. *A decidable first-order logic for knowledge representation.* **Proceedings of the ninth international joint conference on artificial intelligence.** Morgan Kaufmann, Los Altos, CA, pp. 455–458.

1985b. **A decidable logic for knowledge representation.** AI Technical Report No. 45, Schlumberger Palo Alto Research Center.

PERLIS, DONALD

1987. *Nonmonotonic reasoning.* **Encyclopedia of artificial intelligence,** ed. S. C. Shapiro, New York (John Wiley & Sons) pp. 849–853.

RAPH, KARL MARK G.

1983. **The Markgraf Karl refutation procedure.** Interner Bericht, University of Karlsruhe.

REITER, RAYMOND

1978. *On closed world data bases.* **Logic and databases,** ed. H. Gallaire and J. Minker, New York (Plenum) pp. 55–76.

SANDEWALL, ERIK

1985a. *A functional approach to non-monotonic logic.* **Proceedings of the ninth international joint conference on artificial intelligence.** Morgan Kaufmann, Los Altos, CA, pp. 100–106.

1985b. *A functional approach to non-monotonic logic.* **Computational intelligence,** vol. 1, pp. 80–87.

SHAPIRO, STUART C.

1971. **The MIND system: a data structure for semantic information processing.** R-837-PR, The Rand Corp., Santa Monica, CA.

1979. *The SNePS semantic network processing system.* **Associative Networks: The Representation and Use of Knowledge by Computers,** ed. N. V. Findler, New York (Academic Press) pp. 179–203.

See Martins and Shapiro, Shapiro and Rapaport, Shapiro and Wand.

SHAPIRO, STUART C. AND RAPAPORT, WILLIAM J.

1987. *SNePS considered as a fully intensional propositional semantic network.* **The Knowledge Frontier: Essays in the Representation of Knowledge,** ed. G. McCalla and N. Cercone, New York (Springer-Verlag) pp. 262–315.

Shapiro, Stuart C. and Wand, Mitchell

1976. **The relevance of relevance.** Technical Report No. 46, Computer Science Department, Indiana University, Bloomington, IN.

Thistlewaite, P. B., McRobbie, M. A. and Meyer, Robert K.

198+. **Automated theorem-proving in non-classical logics; Research notes in theoretical computer science,** New York (John Wiley & Sons).

Thomason, Richmond H., Horty, John F. and Touretzky, David S.

1986. **A calculus for inheritance in monotonic semantic nets.** CMU-CS-86-138, Department of Computer Science, Carnegie-Mellon University Pittsburgh, PA.

Touretzky, David S.

1987. *Inheritance hierarchy.* **Encyclopedia of artificial intelligence,** ed. S. C. Shapiro, New York (John Wiley & Sons) pp. 422–431.

See Thomason, Horty and Touretzky.