UNIVERSITÄT DES SAARLANDES
FACHBEREICH INFORMATIK
D-66041 SAARBRÜCKEN
GERMANY
WWW: http://js-sfbsun.cs.uni-sb.de/pub/www/

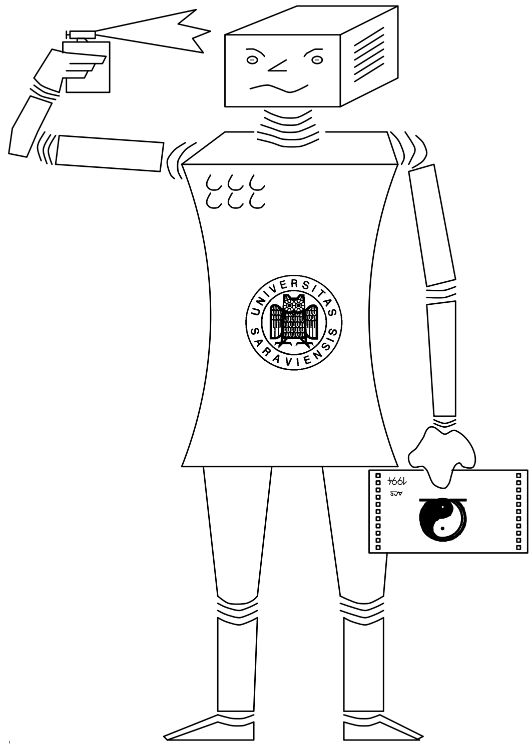

Reasoning with Assertions and Examples

Manfred Kerber Erica Melis
Jörg H. Siekmann

Published as: In Terry Dartnell, Steven Kim, Fay Sudweeks, eds., *Working Notes of the AAAI Spring Symposium on AI and Creativity*, Stanford, California, USA, p.61–66, 1993.

Reasoning with Assertions *And* Examples

Manfred Kerber[1] and **Erica Melis**[2] and **Jörg H. Siekmann**[3]
Universität des Saarlandes
Im Stadtwald
W-6600 Saarbrücken
Germany

Abstract. The hallmark of traditional Artificial Intelligence (AI) research is the symbolic representation and processing of knowledge. This is in sharp contrast to many forms of human reasoning, which to an extraordinary extent, rely on cases and (typical) examples. Although these examples could themselves be encoded into logic, this raises the problem of restricting the corresponding model classes to include only the intended models.

There are, however, more compelling reasons to argue for a hybrid representation based on assertions *as well as* examples. The problems of adequacy, availability of information, compactness of representation, processing complexity, and last but not least, results from the psychology of human reasoning, all point to the same conclusion: Common sense reasoning requires different knowledge sources and hybrid reasoning principles that combine symbolic as well as semantic-based inference.

In this paper we address the problem of integrating semantic representations of examples into automated deduction systems. The main contribution is a formal framework for combining sentential with direct representations. The framework consists of a hybrid knowledge base, made up of logical formulae on the one hand and direct representations of examples on the other, and of a hybrid reasoning method based on the resolution calculus. The resulting hybrid resolution calculus is shown to be sound and complete.

1. Introduction

The traditional paradigm of knowledge representation and processing, dating at least to the time of McCarthy's Advice Taker [McCarthy, 1968], is to represent knowledge by a collection of sentences of a formal language that are viewed as the symbolic knowledge base of the reasoner. Keeping with this view, the reasoning process itself is then seen as the application of explicit inference rules (or procedures) to sentences of the formal language. The best evidence of this point of view is the dominance of predicate, modal and nonmonotonic logics in AI research.

The advantage of such a logic-based approach is the formal logical framework with its precise semantics. Furthermore, the specific logic under consideration provides a clean inference mechanism for deriving results from the theory, where applying inference rules to a knowledge base takes care of the procedural aspects of an AI-system. This view of traditional AI, now quite commonly accepted, arose from the procedural versus declarative representation debate of the seventies.

There are, however, several problems with this traditional approach: In the first place, it has become clear that logics are not always the best or easiest way to represent complex forms of information. Consider, for instance, a variation of a well-known example from nonmonotonic logic: Suppose Birdy is a bird and Tweety is a typical bird. Assume also that in addition to the object language (say first-order predicate logic) there is a means of representing a typical bird, e.g. in a neural net, so that the query `can_fly(Tweety)` evaluates to `true` if the representation of the typical bird Tweety contains the information that Tweety can, in fact, fly. In other words, we have two different levels of information: the syntactic information, `bird(Birdy)`, which is stated in the object language, where (deductive) inferences are drawn, and the semantic information, where for example a typical case of a bird, namely `Tweety`, is represented. If `can_fly(Tweety)` evaluates to true, we tentatively conclude that Birdy can fly, too. Apparently, this procedure is similar to the way humans reason under these circumstances: the default knowledge is stored in the form of an example [Motte-Haber, 1988], from which the conclusion is drawn *by analogy* rather than by an explicit rule of deductive inference — be it monotonic or not. If we want to know something about an arbitrary bird, for instance, whether or not it has teeth, we have no rule in mind such as: "typical birds have no teeth" (or rules to reflect the myriad of other facts that are not the case), but we think of a typical representative for the concept bird and reason by analogy: Tweety has no teeth, hence Birdy has no teeth. A purely logical approach would require some complicated rule of inference

[1] Fachbereich Informatik, e-mail: `kerber@cs.uni-sb.de`
[2] Fachbereich Informatik, e-mail: `melis@cs.uni-sb.de`
[3] DFKI, German Research Center for Artificial Intelligence, e-mail: `siekmann@dfki.uni-sb.de`

such as circumscription in order to minimize the intended models.

Thus the first problem refers to the *adequacy* of representing and processing knowledge only in terms of formulae in some formal language. It is a well-known observation that human reasoning relies to an extra-ordinary extent on cases and typical example (see next paragraph). The criticism of classical AI by Dreyfus and Dreyfus [Dreyfus and Dreyfus, 1986], for example, claims that only the lower stages of human knowledge processing are based on an explicit use of rules whereas for the three higher stages — competence, proficiency, and expertise — large amounts of well-chosen examples are imperative. Many experiments in cognitive psychology have provided ample evidence for an explicit rule application, as well as for processing information that is somehow more directly represented.

A second problem concerns the availability of logically formalized knowledge and explicit rules of inference. Often, knowledge about examples, diagrams and the like is encoded implicitly, as for example, in the analogical representation of a map [Sloman, 1975].

Another problem becomes obvious by comparing the compactness of syntactic and semantic encodings of information, e.g., information about a road map. The direct, "semantic" representation is far more concise than a corresponding explicit symbolic representation (if one can be obtained at all).

Finally, the complexity of knowledge processing, e.g. theorem proving, within the traditional paradigm depends on the expressive power of the formal language (cf. [Halpern and Vardi, 1991]).

In this paper we present an approach to knowledge representation and processing that combines a direct "semantic" representation of examples with the usual sentential representation. The main idea of this approach is to replace some information traditionally inferred using syntactic rules by appropriate semantic information which we assume to be somehow directly represented, perhaps in a neural net, some data structures, or a case base.

Hybrid approaches to knowledge representation and reasoning have also been suggested, inter alia, by Halpern and Vardi [Halpern and Vardi, 1991], Johnson-Laird and Byrne [Johnson-Laird and Byrne, 1990], and Myers and Konolige [Myers and Konolige, 1992]. In [Kerber *et al.*, 1992] we demonstrated how analogical reasoning with typical examples can be done with hybrid knowledge bases. In this paper we are interested in the combination of *deductive* reasoning with reasoning by examples that is based on the same hybrid knowledge base.

First, we motivate our approach by giving psychological evidence for the use of examples in human reasoning. Then we introduce a formal framework consisting of a semantics for hybrid reasoning based on three-valued logic together with a resolution calculus.

2. Psychological Evidence

Do people reason by applying explicit rules to logically formalized knowledge or do they reason by example? This problem has been widely investigated by cognitive psychologists, and psychological experiments have provided evidence for both modes of reasoning. For instance, Cherniak [Cherniak, 1984] and Medin and Ross [Medin and Ross, 1989] found support for their thesis that people reason using information directly extracted from examples/models in an experiment clearly exhibiting retrieval of examples and the subsequent use of these examples for analogical reasoning. In addition, the importance of *typical examples* in reasoning has been shown by several researchers [Rips, 1975], [Cherniak, 1984]. Kaiser, Jonides, and Alexander [Kaiser *et al.*, 1986], for example, demonstrated that people draw on their formal models of physics only after they are unable to find an acceptable solution by analogy, and Ross [Ross, 1987] has shown that novices make use of analogy with earlier solution instances even when a principle or a rule has been presented explicitly. On the other hand, some psychological findings point toward an explicit rule application of facts (see [Smith *et al.*, 1992] for an overview). Braine, Reiser, and Rumain [Braine *et al.*, 1984] have shown that the more rules are required in order to determine the validity of an argument, the longer the reaction time and the lower the accuracy of the final response to questions on the argument. Yet they found that people generally agree on the correctness of modus ponens applications. Such evidence strongly suggests that people do pay attention to the structure of an argument, that is, its logical form.

In other words, both views of reasoning are legitimate and can be useful depending on the context, the conditions, and the aim of the respective reasoning process. Since in many real life situations, neither one nor the other position is sufficient to cover the whole case, we argue for a hybrid approach to knowledge representation and reasoning in AI, rather than for a purely symbolic or a purely representation-based approach. Further experimental support for this position is given in [Smith *et al.*, 1992] and by Galotti, Baron, and Sabini [Galotti *et al.*, 1986]. The latter have performed tests with syllogisms and concluded that there exist deduction rules in human reasoning as well as non-rule entities which we shall sometimes refer to as noncompositional, following Myers and Konolige [Myers and Konolige, 1992]. Smith, Langston, and Nisbett [Smith *et al.*, 1992] also discussed hybrid reasoning mechanisms.

3. A Framework for Hybrid Reasoning

In the following paragraph we shall introduce a general semantics for knowledge bases represented by formulae *and* examples. After that we will present a logic for hybrid reasoning that is based on the resolution calculus, and extend this to a hybrid calculus. Finally, we shall show soundness and completeness of this extended calculus with respect to the given semantics.

3.1 Model Theory for Hybrid Reasoning

As described above, we want to base our reasoning mechanism upon logical formulae as well as examples. For our purpose, a knowledge base will consist of a set of formulae Γ and a set of example sets $\mathcal{E}$ (for each concept one set of examples). Of course, the formulae and the example sets should be connected in some way, and we will express the relationship between them by a partial interpretation function $\partial\mathcal{I}$.

Definition (Knowledge Base): A *knowledge base* is a triple $\Delta = \langle \Gamma, \mathcal{E}, \partial\mathcal{I} \rangle$, where Γ is a formula set expressed in a logic $\mathcal{L}$. $\mathcal{E}$ is a set of example sets, and $\partial\mathcal{I}$ is a partial interpretation of $\mathcal{L}$-formulae. ■

More precisely, we assume a sorted (first-order) logic $\mathcal{L}$, where each sort can be viewed as a concept like bird, human, or female. We denote the sorts by lowercase greek letters such as κ or μ. $\mathcal{E}$ is a set of sets $\{\mathcal{E}_\kappa\}_\kappa$, where each $\mathcal{E}_\kappa$ is called the *set of examples of sort* κ. The $\mathcal{E}_\kappa$ are such that their structure corresponds to the sort structure of $\mathcal{L}$, that is, if $\mu \sqsubseteq \kappa$ (i.e. μ is *subsort* of κ) then $\mathcal{E}_\mu \subseteq \mathcal{E}_\kappa$. $\mathcal{E}$ forms the *frame* (the collection of universes) for the partial interpretation of the terms. $\partial\mathcal{I}$ is a fixed partial interpretation function (corresponding to three-valued strong Kleene logic $\mathcal{L}^{\mathbf{K}}$ [Kleene, 1952, Urquhart, 1986]) in the frame $\mathcal{E}$. Each term t of sort κ is either interpreted by an example in $\mathcal{E}_\kappa$ or by the bottom element $\perp$. Formulae may be evaluated by $\partial\mathcal{I}$ to $\mathtt{true}$, to $\mathtt{false}$, or to $\mathtt{undef}$. Furthermore, we assume that for every element $e^i_\kappa \in \mathcal{E}_\kappa$ there exists a constant c^i_κ of sort κ with $\partial\mathcal{I}(c^i_\kappa) = e^i_\kappa \in \mathcal{E}_\kappa$ and that there are only finitely many examples in $\mathcal{E}$.

The semantics of composed formulae is defined as usual, based on the propositional connectives as defined by the following truth tables:

$\neg$	
false	true
undef	undef
true	false

$\vee$	false	undef	true
false	false	undef	true
undef	undef	undef	true
true	true	true	true

In order to fix the semantics of the universal quantifier, assignments ξ for the interpretation of the variables into the frame are necessary. If ξ is an arbitrary assignment, $\xi[x \leftarrow a]$ denotes the assignment equal to ξ for all variables except for x and $\xi[x \leftarrow a](x) = a$. $\partial\mathcal{I}_\xi(\forall x_\kappa \varphi) :=$

$$\begin{cases} \mathtt{true} & \text{if } \partial\mathcal{I}_{\xi[x\leftarrow a]}(\varphi) = \mathtt{true} \text{ for all } a \in \mathcal{E}_\kappa \\ \mathtt{false} & \text{if } \partial\mathcal{I}_{\xi[x\leftarrow a]}(\varphi) = \mathtt{false} \text{ for one } a \in \mathcal{E}_\kappa \\ \mathtt{undef} & \text{else} \end{cases}$$

The semantics of $\wedge$, $\Rightarrow$, $\Leftrightarrow$, and $\exists$ can then be defined in the usual manner.

Note that these definitions do not assume a concrete representation of the examples. The examples may, in fact, be represented by a data structure like a semantic net or by some frame structure, or they may be represented within a neural net – the only requirement is that we get an answer to certain questions, thus fixing the interpretation function. In other words, $\partial\mathcal{I}$ has to be effectively computable for all ground formulae (i.e., variable free formulae) and, consequently, for all formulae, since we assume the number of examples in $\mathcal{E}$ to be finite.

As usual, we give an (extended) set theoretic semantics for a formula set. Our semantics is such that it is compatible with the examples. An interpretation of a knowledge base Δ is defined as an extension of the partial interpretation, given by $\langle \mathcal{E}, \partial\mathcal{I} \rangle$.

Definition ($\langle \mathcal{E}, \partial\mathcal{I} \rangle$-Interpretation): Let $\mathcal{E} = \{\mathcal{E}_\kappa\}_\kappa$ be a given set of example sets and let $\partial\mathcal{I}$ be a partial interpretation function in $\mathcal{E}$. An interpretation $\langle \{\mathcal{D}_\kappa\}_\kappa, \mathcal{I} \rangle$ is called an $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-*interpretation* iff

- there are mappings $\natural_\kappa : \mathcal{E}_\kappa \to \mathcal{D}_\kappa$ with $\mathcal{I}(c_\kappa) = \natural_\kappa(\partial\mathcal{I}(c_\kappa))$ for all constant symbols c_κ with $\partial\mathcal{I}(c_\kappa) \neq \perp$. (When the sort is not important, we omit the index κ and simply write $\natural$.)
- for all terms t and arbitrary ground instances $\sigma(t)$ with $\partial\mathcal{I}(\sigma(t)) \neq \perp$ holds $\mathcal{I}(\sigma(t)) = \natural(\partial\mathcal{I}(\sigma(t)))$.
- for all formulae φ and all ground instances σ of φ holds, if $\partial\mathcal{I}(\sigma(\varphi)) \neq \mathtt{undef}$ then $\mathcal{I}(\sigma(\varphi)) = \partial\mathcal{I}(\sigma(\varphi))$

If $\mathcal{I}_\xi(\varphi) = \mathtt{true}$ for all assignments ξ then $\mathcal{I}$ is called an $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-*model* of φ. If φ has no $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-model, it is said to be $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-*unsatisfiable*. Γ $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-*entails* the formula φ iff each $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-model of Γ is an $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-model of φ, too (i.e. $\Gamma \models_{\langle \mathcal{E}, \partial\mathcal{I} \rangle} \varphi$). ■

Example: Let $\mathcal{L}$ be the sorted logic with just one sort bird and two constants tweety and birdy. Furthermore, let the knowledge base Δ consist of the formula set

$$\begin{aligned} \Gamma = \ & \{\mathsf{red}(\mathsf{tweety}), \\ & \neg\mathsf{canfly}(\mathsf{birdy}), \\ & \forall x_{\mathsf{bird}} \mathsf{canfly}(x) \Rightarrow \mathsf{has\_feather}(x)\}, \end{aligned}$$

the example set be $\mathcal{E}_{\mathsf{bird}}$ which contains one example for tweety and another for birdy, and the partial interpretation be $\partial\mathcal{I}$. By looking up the examples in the knowledge base, $\partial\mathcal{I}$ evaluates $\mathsf{canfly}(\mathsf{tweety})$ to $\mathtt{true}$ and $\mathsf{canfly}(\mathsf{birdy})$ to $\mathtt{undef}$. Then in all $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-models of Γ, $\mathsf{red}(\mathsf{tweety})$, $\mathsf{canfly}(\mathsf{tweety})$, and $\mathsf{has\_feather}(\mathsf{tweety})$ hold, whereas $\mathsf{canfly}(\mathsf{birdy})$ does not hold. The assertion $\mathsf{has\_feather}(\mathsf{birdy})$ may hold or not. ■

A knowledge base $\langle \Gamma, \mathcal{E}, \partial\mathcal{I} \rangle$ is called *consistent* if an $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-model of Γ exists, else *inconsistent*. As in classical logic, a knowledge base is rather useless if it contains incompatible knowledge. If Γ contains, for instance, a ground formula φ such that $\partial\mathcal{I}(\varphi) = \mathtt{false}$, then the knowledge base is not consistent.

3.2 Proof Theory for Hybrid Reasoning

Assume the two inference rules **RES** and **FACT** for resolution and factoring, respectively [Robinson, 1965], and let $\mathcal{C}$ be the initial clause set generated from a knowledge base Γ and the negation of a hypothesis φ by normalization. φ can be *derived classically* from Γ, denoted $\Gamma \vdash \varphi$, if the empty clause $\Box$ is in the transitive closure of $\mathcal{C}$ under the rules **RES** and **FACT**.

Often there are simplification rules **SIMP**, for the sake of efficiency, which simplify the clause set, for instance, by deleting redundant clauses.

In order to prove that a hypothesis φ follows hybridly from a knowledge base $\Delta = \langle \Gamma, \mathcal{E}, \partial\mathcal{I} \rangle$, first Γ and $\neg\varphi$ are transformed into the set $\mathcal{C}$ of clauses in normal form. Then the empty clause $\Box$ is to be derived from $\mathcal{C}$ by application of **RES**, **FACT** and an additional rule, **EXEMP** (for exemplification), which looks up the knowledge base of examples. (It is not too hard to see that the transformation to clausal normal form does not change $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-unsatifiability. Newly introduced Skolem-terms are not known to $\partial\mathcal{I}$ and, consequently, they are interpreted to $\perp$.)

Let $L_1 \vee \ldots \vee L_m \vee K_1 \vee \ldots \vee K_n$ be a clause, and σ be a substitution such that all $\sigma(L_i)$ are ground and $\partial\mathcal{I}(\sigma(L_i)) = \mathtt{false}$. Then **EXEMP** is defined as:

$$\boxed{\textbf{EXEMP} \quad \frac{L_1 \vee \ldots \vee L_m \vee K_1 \vee \ldots \vee K_n}{\sigma(K_1) \vee \ldots \vee \sigma(K_n)}}$$

This rule is in a sense the *semantic* equivalent to the usual resolution rule (consequently we call the calculus "exemplification-based resolution"). In order to see this, consider the modus ponens mode of reasoning: if A implies B and A holds then B. Now exemplification-based reasoning rests on the following form: if A implies B and A holds in the database of cases then B. Now as usual, if $\neg A$ or B and A holds in the database of cases, then $\neg A$ is false, hence B. In other words the above rule **EXEMP** generalizes "semantic modus ponens" just as resolution is a generalization of modus ponens in the sense that there may be more than just the two predicates A and B and that possible substitutions of the variables must be considered.

We may assume that for all K_j in the result of the **EXEMP**-rule, $\partial\mathcal{I}(\sigma(K_j)) \neq \mathtt{false}$, otherwise, such literals K_j could be eliminated by repeated exemplification. For an arbitrary clause C, different applications of **EXEMP** may be possible (as a result, for example, of different ground substitutions σ). But, since there are only finitely many examples in $\mathcal{E}$, there are only finitely many possible applications of **EXEMP** for each clause.

EXAMPLE: If we want to derive $\mathsf{has\_feather}(\mathsf{tweety})$ in the example given above, we obtain as initial clause set

$$\mathcal{C} = \{\mathsf{red}(\mathsf{tweety}),\ \neg\mathsf{canfly}(\mathsf{birdy}),\ \neg\mathsf{canfly}(x) \vee \mathsf{has\_feather}(x),\ \neg\mathsf{has\_feather}(\mathsf{tweety})\}.$$

Because $\partial\mathcal{I}(\neg\mathsf{canfly}(\mathsf{tweety})) = \mathtt{false}$, we can apply **EXEMP** to the second last clause and get the clause $\mathsf{has\_feather}(\mathsf{tweety})$. An application of **RES**, together with the last of the initial clauses, then yields the empty clause as the resolvent. ■

In addition to the traditional simplification rules for resolution there is a simplification rule based on $\partial\mathcal{I}$, namely, a clause can be deleted if one of its literals, say L, is ground and $\partial\mathcal{I}(L) = \mathtt{true}$.

Exemplification-based resolution, that is, reasoning that is based on the inference rules **RES**, **FACT** and **EXEMP**, can be considered as a semantic analogue to syntactic T-resolution [Stickel, 1985].

The next two theorems show that this hybrid resolution calculus is adequate for the notion of hybrid entailment.

THEOREM: *The calculus consisting of the rules* **RES**, **FACT**, *and* **EXEMP** *is refutation sound.*

Proof: We first show the soundness of **RES** (for **FACT** the proof is analogous). Let C_1 and C_2 be arbitrary clauses and let $\mathcal{I}$ be an $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-model of C_1 and C_2. Let R be a resolvent of C_1 and C_2; we show that $\mathcal{I}$ is an $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-model of R, as well. Let C_1 and C_2 be

$$C_1 = L \vee K_1 \vee \ldots \vee K_l \text{ and}$$
$$C_2 = \neg L' \vee M_1 \vee \ldots \vee M_n$$

with most general unifier σ of L and L'. Their resolvent is

$$R = \sigma(K_1) \vee \ldots \vee \sigma(K_l) \vee \sigma(M_1) \vee \ldots \vee \sigma(M_n).$$

Since $\mathcal{I}$ is an $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-model of C_1 and C_2, it is a standard model of C_1 and C_2. Hence, because of the soundness of standard resolution, $\mathcal{I}$ is a standard model of R. Assume $\mathcal{I}$ is *not* an $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-model of R, then $\partial\mathcal{I}(R) = \mathtt{false}$. Consequently, for all literals K_i and M_j, $\partial\mathcal{I}(\sigma(K_i)) = \partial\mathcal{I}(\sigma(M_j)) = \mathtt{false}$. That is why $\partial\mathcal{I}(K_i) = \mathtt{false}$ and $\partial\mathcal{I}(M_j) = \mathtt{false}$, and by the definition of $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-interpretations, $\mathcal{I}(K_i) = \mathtt{false}$ and $\mathcal{I}(M_j) = \mathtt{false}$. Finally, because of $\mathcal{I}(C_1) = \mathcal{I}(C_2) = \mathtt{true}$, $\mathcal{I}(L) = \mathcal{I}(\neg L') = \mathtt{true}$ also holds, contradicting the unifiability of L and L'.

To prove the soundness of **EXEMP**, we assume $\mathcal{I}$ to be an $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-model for a clause $L_1 \vee \ldots \vee L_m \vee K_1 \vee \ldots \vee K_n$ and assume that the L_i can be eliminated by **EXEMP** and the K_j can not. Then $\mathcal{I}$ is an $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-model of $\sigma(L_1) \vee \ldots \vee \sigma(L_m) \vee \sigma(K_1) \vee \ldots \vee \sigma(K_n)$ for any substitution σ. For substitutions σ occuring within the applications of **EXEMP**, the literals $\sigma(L_i)$ are ground. This implies that $\mathcal{I}(\sigma(L_i)) = \partial\mathcal{I}(\sigma(L_i)) = \mathtt{false}$ for all i and thus guarantees $\mathcal{I}(\sigma(K_1) \vee \ldots \vee \sigma(K_n)) = \mathtt{true}$. Since $\mathcal{I}$ is an $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-interpretation, $\mathcal{I}$ is an $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-model of the resulting exemplified clause, $\sigma(K_1) \vee \ldots \vee \sigma(K_n)$.

The soundness of the calculus follows since all three rules are sound and the empty clause is trivially $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-unsatisfiable. ■

THEOREM: *The calculus consisting of the rules* **RES**, **FACT**, *and* **EXEMP** *is refutation complete.*

Proof: Let $\mathcal{C}$ be a finite, $\langle \mathcal{E}, \partial\mathcal{I} \rangle$-unsatisfiable clause set. In order to show completeness, the derivability of $\Box$

from $\mathcal{C}$ by **RES**, **FACT**, and **EXEMP** must be proved. We therefore form the transitive closure $\bar{\mathcal{C}}$ of $\mathcal{C}$ under **EXEMP** which can be done in finitely many steps and prove that $\bar{\mathcal{C}}$ is unsatisfiable in the standard sense (and that $\Box$ is derivable from $\bar{\mathcal{C}}$ by **RES** and **FACT** according to the completeness result for standard resolution).

For a proof by refutation, assume $\bar{\mathcal{C}}$ has a standard model $\langle\{\mathcal{D}_\kappa\}_\kappa, \mathcal{I}\rangle$. We show that there is an $\langle\mathcal{E}, \partial\mathcal{I}\rangle$-model $\langle\{\mathcal{D}_\kappa\}_\kappa, \mathcal{I}^\star\rangle$ of $\mathcal{C}$.

Construction of $\mathcal{I}^\star$: Let $\natural_\kappa$ be partial mappings from $\mathcal{E}_\kappa$ to $\mathcal{D}_\kappa$ such that for all constants c_κ denoting examples of sort κ, $\natural(\partial\mathcal{I}(c_\kappa)) := \mathcal{I}(c_\kappa)$ holds. For all constants c in the signature of the underlaying language $\mathcal{L}$ we fix $\mathcal{I}^\star(c) := \mathcal{I}(c)$.

Let P be a predicate symbol of sort $\kappa_1 \times \cdots \times \kappa_m \rightarrow \kappa$ in $\mathcal{L}$ and let d_{κ_i} be arbitrary elements of $\mathcal{D}_{\kappa_i}$. Then

$$\mathcal{I}^\star(P(d_{\kappa_1}, \ldots, d_{\kappa_m})) := \begin{cases} \partial\mathcal{I}(P)(\natural_{\kappa_1}(d_{\kappa_1}), \ldots, \natural_{\kappa_m}(d_{\kappa_m})) & \text{if not } \mathtt{undef} \\ \mathcal{I}(P)(d_{\kappa_1}, \ldots, d_{\kappa_m}) & \text{else.} \end{cases}$$

By defining the interpretation of function symbols analogously, $\mathcal{I}^\star$ is fixed as an $\langle\mathcal{E}, \partial\mathcal{I}\rangle$-interpretation for all formulae of the language $\mathcal{L}$. It remains to show that $\mathcal{I}^\star$ is indeed a model of $\mathcal{C}$.

Let $C \in \mathcal{C}$ and $C = K_1 \vee \ldots \vee K_m$ and assume $\mathcal{I}^\star(C) = \mathtt{false}$. By Herbrand's theorem there is a ground substitution σ such that $\mathcal{I}^\star(\sigma(C)) = \mathtt{false}$. Since $\bar{\mathcal{C}} \supseteq \mathcal{C}$ and $\mathcal{I} \models \mathcal{C}$, we must have $\mathcal{I}(\sigma(C)) = \mathtt{true}$, and since $\mathcal{I}^\star(\sigma(C)) = \mathtt{false}$, $\partial\mathcal{I}(\sigma(C))$ must evaluate to $\mathtt{false}$ (definition of $\mathcal{I}^\star$), i.e. , $\partial\mathcal{I}(\sigma(K_i))$ must be $\mathtt{false}$ for all i. Thus, the empty clause $\Box$ can be derived from C by **EXEMP**, and so $\Box$ is an element of $\bar{\mathcal{C}}$. This contradicts the assumption that $\bar{\mathcal{C}}$ has a model. ■

Remark: Of course, it is *not* the idea of the exemplification-based resolution calculus to perform all possible applications of **EXEMP** in advance as in the completeness proof. The application of the **EXEMP**-rule, the selection of a clause, and the choice of σ, are tasks for an efficient control strategy. In particular, it is useful to have a partial order on the examples and to apply the exemplification rule preferentially, checking only typical examples. ■

4. A General Trend to Hybrid Reasoning

While the classical view of AI stresses the assertional part of knowledge representation and rule-based reasoning, in pure case-based reasoning another extreme position is taken: only examples are stored and processed.

Since neither the mere logical nor the mere semantic-based approaches are adequate in all cases, several subareas of AI independently witnessed the development of techniques for reasoning based on various semantically oriented representations mixed with the traditional deductive approach. To wit:

- Model-based methods in theorem proving rely on hybrid representations. Typical approaches include model construction methods [Slaney and Lusk, 1992], model checking [Halpern and Vardi, 1991], "vivid reasoning" [Etherington *et al.*, 1989], and the various methods of representing a model in the form of a diagram in geometry theorem proving [Gelernter, 1959].
- Recently, diagrams have become important again as a representational medium, and the manipulating of diagrams [Myers and Konolige, 1992] or depictions [Latecki and Pribbenow, 1992] as a means of inferencing.

  The investigation of analogical reasoning based on the semantic representation of typical examples as opposed to rule-based analogical inference, is becomming an active area of research in hybrid reasoning (see, e.g., [Kerber *et al.*, 1992] and the additional references there).

5. Conclusion: Exemplification-Based Reasoning

Hybrid reasoning as proposed in this paper, attempts to combine the strong tools of formalization taken from logic *and* the more semantic oriented, noncompositional representation of examples. A formal framework for combining sets of formulae with the information provided by an explicit representation of examples was presented. While the resulting hybrid reasoning presented in this paper is resolution-based, we are not particularly committed to the use of resolution. Indeed, our representation is suitable for combination with any deductive rule of inference. Moreover, this approach can also be employed in conjunction with non-deductive inference methods, e.g., for analogical reasoning based on typical examples, induction, abduction, etc. In a forthcoming paper [Kerber *et al.*, 1993], we present a uniform framwork for hybrid reasoning, called *exemplification*. The idea behind exemplification-based reasoning is that the different traditional modes of logical inference (deduction, abduction, and induction as, for example, proposed by Charles S. Peirce) can be augmented by an additional "exemplification" rule instructing the reasoning system to look up an (typical) example directly.

We have outlined here how our framework supports hybrid reasoning in the case of for *reliable* (deductive) reasoning. In conjunction with *tentative* reasoning, such as hybrid analogy or hybrid abduction, it is important to use the structure of the given sets of examples as in [Kerber *et al.*, 1992].

Acknowledgements. We would like to thank Patricia Johann for her comments on a draft of this paper.

References

Braine, M.D., Reiser, B.J. and Rumain, B.: 1984, Some empirical justification for a theory of natural porpositional logic. *Psychological Learning Motivation*, 18:313–371.

Cherniak, Ch.: 1984, Prototypicality and deductive reasoning. *Journal of Verbal Learning and Verbal Behavior*, 23:625–642.

Dreyfus, H.L. and Dreyfus, S.E.: 1986, *Mind over Machine.* The Free Press. The Power of Human Intuition and Expertise in the Era of the Computer.

Etherington, D., Borgida, A., Brachman, R.J., and Kautz, H.: 1989, Vivid knowledge and tractable reasoning. In *Proceedings of eleventh IJCAI-89*, pages 1146–1152.

Galotti, K.M., Baron, J., and Sabini, J.: 1986, Individual differences in syllogistic reasoning: Deduction rules or mental models? *Journal of Experimental Psychology: General*, 115:16–25.

Gelernter, H.: 1959, Realization of a geometry theorem-proving machine. In *Proceedings of the International Conference on Information Processing, UNESCO.*

Halpern, J.Y. and Vardi, M.Y.: 1991, Model checking vs. theorem proving. In *Proceedings of the Second International Conference on Knowledge Representation KR'91*, pages 325–334, Cambridge MA.

Johnson-Laird, P.N. and Byrne, R.M.: 1990, *Deduction.* Ablex Publishing Corp..

Kaiser, M.K., Jonides, J., and Alexander, J., 1986, Intuitive reasoning on abstract and familiar physics problems. *Memory & Cognition*, 14:308–312.

Kerber, M., Melis, E., and Siekmann, J.: 1992, Analogical reasoning with typical examples. SEKI-Report SR-92-13, Universität des Saarlandes, Saarbrücken.

Kerber, M., Melis, E., and Siekmann, J.: 1993, Exemplification-based reasoning. SEKI-Report, Universität des Saarlandes, Saarbrücken. forthcoming.

Kleene, S.C.: 1952, *Introduction to Metamathematics.* Van Nostrand, Amsterdam, The Netherlands.

Latecki, L. and Pribbenow, S.: 1992, On hybrid reasoning for spatial expressions. In B. Neumann, editor, *Proceedings of the 10th European Conference on Artificial Intelligence*, pages 289–393. John Wiley.

McCarthy, J.: 1968, Programs with common sense. In Marvin Minsky, editor, *Semantic Information Processing*, pages 403–418. MIT Press, Cambridge, Massachusetts.

Medin, D.L. and Ross, B.H.: 1989, The specific character of abstract thought: Categorization, problem solving, and induction. In R.J. Sternberg, editor, *Advances in the Psychology of human intelligence*, volume 5. Erlbaum.

de la Motte-Haber, H.: 1988, Regelmaß und Einmaligkeit als ästhetische Prinzipien. In R. Wille, editor, *Symmetrie in Geistes- und Naturwissenschaften*, pages 17–29. Springer Verlag, Berlin, Germany.

Myers, K.L. and Konolige, K.: 1992, Reasoning with analogical representations. In *Proceedings of KR'92*, pages 189–200.

Rips, L.: 1975, Inductive judgements about natural categories. *Journal of Verbal Learning and Verbal Behavior*, 14:665–681.

Robinson, J.A.: 1965, A machine oriented logic based on the resolution principle. *Journal of the ACM*, 12:23–41.

Ross, B.H.: 1987, This is like that. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 13:629–639.

Slaney, J.K. and Lusk, E.L.: 1992, Finding models. techniques and applications. Technical report, Tutorial at 11th CADE.

Sloman, A.: 1975, Afterthoughts on analogical representation. In *Proceedings of Theoretical Issues in Natural Language Processing.*

Smith, E.E., Langston, C. and Nisbett, R.E.: 1992, The case for rules in reasoning. *Cognitive Science*, 16(1):1–40.

Stickel, M.E.: 1985, Automated deduction by theory resolution. *Journal of Automated Reasoning*, 1:333–356.

Urquhart, A.: 1986, Many-valued logic. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, chapter III.2, pages 71–116. D.Reidel Publishing Company, Dordrecht, Netherlands. Volume III: Alternatives to Classical Logic.