Report Number: WUCS-91-27

1991-05-01

Reasoning for Inconsistency-- A Taxonomy and a Connectionist Approach

Gadi Pinkas and Ron P. Loui

A consequence relation (CR) relates sets of beliefs to the appropriate conclusions that might be deduced. Of special interest to Artificial Intelligence are CRs that cope with inconsistency within the set of beliefs. Default reasoning, belief revision, social choice and reasoning from conflicting knowledge sources are just a few examples of mechanisms that need to handle inconsistency, In this paper we show a taxonomy in which many existing mechanisms are mapped, and new interesting ones are revealed. We identify simple relations among the CRs and give a language for their specification. We then show that a large portion of... Read complete abstract on page 2.

Reasoning for Inconsistency– A Taxonomy and a Connectionist Approach

Gadi Pinkas and Ron P. Loui

Complete Abstract:

A consequence relation (CR) relates sets of beliefs to the appropriate conclusions that might be deduced. Of special interest to Artificial Intelligence are CRs that cope with inconsistency within the set of beliefs. Default reasoning, belief revision, social choice and reasoning from conflicting knowledge sources are just a few examples of mechanisms that need to handle inconsistency, In this paper we show a taxonomy in which many existing mechanisms are mapped, and new interesting ones are revealed. We identify simple relations among the CRs and give a language for their specification. We then show that a large portion of the CRs described by the language can be implementable in neural networks like Boltzman machines and Hopfield nets. The result demonstrates the flexibility of these connectionist models for the approximation of a variety of knowledge level theories.

Reasoning from Inconsistency - A Taxonomy and a Connectionist Approach

Gadi Pinkas and Ron P. Loui

WUCS-91-27

May 1991

Department of Computer Science
Washington University
Campus Box 1045
One Brookings Drive
Saint Louis, MO 63130-4899

Reasoning from Inconsistency-
A Taxonomy and a
Connectionist Approach

Gadi Pinkas
Ron P. Loui
May 13, 1991
WUCS-91-27

Department of Computer Science
Washington University
509 Bryan, Campus Box 1045
One Brookings Drive.
St. Louis, Missouri 63130
pinkas@cics.wustl.edu
Tel:(314) 726-7526

Abstract

A consequence relation (CR) relates sets of beliefs to the appropriate conclusions that might be deduced. Of special interest to Artificial Intelligence are CRs that cope with inconsistency within the set of beliefs. Default reasoning, belief revision, social choice and reasoning from conflicting knowledge sources are just a few examples of mechanisms that need to handle inconsistency. In this paper we show a taxonomy in which many existing mechanisms are mapped, and new interesting ones are revealed. We identify simple relations among the CRs and give a language for their specification. We then show that a large portion of the CRs described by the language can be implemented using a connectionist constraint satisfaction approach and therefore implementable in neural networks like Boltzman machines and Hopfield nets. The result demonstrates the flexibility of these connectionist models for the approximation of a variety of knowledge level theories.

*Supported by NSF grant number: 22-1321-57136

1. Introduction

A variety of inference mechanisms has been proposed for reasoning from a set of beliefs expressed as logic well formed formulas (WFFs). Of special interest to artificial intelligence are those mechanisms that infer intelligently from possibly inconsistent set of beliefs. This is the core of what Nicholas Rescher calls a theory of plausible reasoning whose main task is "to deal with cognitive dissonance" [Rescher 76]. Default reasoning, belief revision, and reasoning from unreliable and conflicting source(s) of knowledge, to name just few, must all deal with inconsistent knowledge. These are tasks that humans so excel in doing, and yet formal systems find hard to solve.

Few inference engines exist to solve these problems and those that do try to solve it suffer from inherently slow performance and from inflexibility to learn and adjust. Connectionist networks *may* provide an answer to these problems; however, currently only few connectionist systems implement a strong knowledge level theory (for examples of such connectionist systems see: [Shastri 85], [Derthick 88]).

In the first part of the paper we define *consequence relations* (CRs) that link together sets of beliefs and their conclusions, and describe a taxonomy of such CRs. A language to describe consequence relations is given, and operators are defined that relate and tie together different such CRs. The result is an elegant taxonomy of new and existing inference mechanisms (for example, non-monotonic (NM) reasoning, voting systems, belief revision, logic programming etc.) that sheds light on the relations among those mechanisms.

In the second part of the paper we demonstrate that many such CRs can be implemented on a connectionist constraint satisfaction architecture. Hopfield networks, [Hopfield 84], Boltzman machines [Hinton, Sejnowski 86], mean field theory [Hinton89] or the special energy functions of [Derthick 88] can be used as the target networks.

2. A taxonomy of cautiousness

2.1. Consequence relations

Let L be a finite set of well formed formulas (WFFs from either propositional or predicate logic). A consequence relation (CR) is a binary relation defined over $\mathcal{P}(L) \times L$. We use the notation $\psi R \varphi$ to denote a pair that belongs to the CR R, where ψ is a finite subset of L and φ is a WFF in L. ψ is the set of beliefs used as a knowledge base while φ is a conclusion allowed by the relation R. The relation $\models$ (or $\vdash$) used in standard logic is an example of a CR which in case of inconsistent set of beliefs concludes trivially all WFFs. We'll show that using a simple language we can define many more consequence relations, among them is the standard $\models$ relation, as well as other interesting ones.

Before we formally define our notation, let's examine two examples: Given a possibly inconsistent set of WFFs, classical logic trivially entails everything if indeed there is a contradiction in the set. However, a more plausible mechanism may use the following inference procedure: 1) Find all maximal consistent (MC) subsets (maximal in the sense of set inclusion). 2) Take the closure of each such subset under classical entailment and 3) Conclude only the WFFs that are in the intersection of all the closures. Several systems use this kind of skeptical principle (which we call R_0, see figure 1). This includes some of the first NM systems [Reiter 1987], dealing with inconsistency [Rescher, Manor 70] and some social choice (voting systems) with the Pareto relation [Campbell 79].

Now, assume we have a mechanism that uses the negation as failure principle on the consequences of R_0; i.e, we conclude φ from a set ψ iff we can not conclude the negation of φ from ψ in R_0. If our set of beliefs ψ is complete [1] (e.g. when using a closed world assumption) then the negation as failure

[1] ψ is complete in R iff for all φ either $\psi R \varphi$ or $\psi R \neg\varphi$.

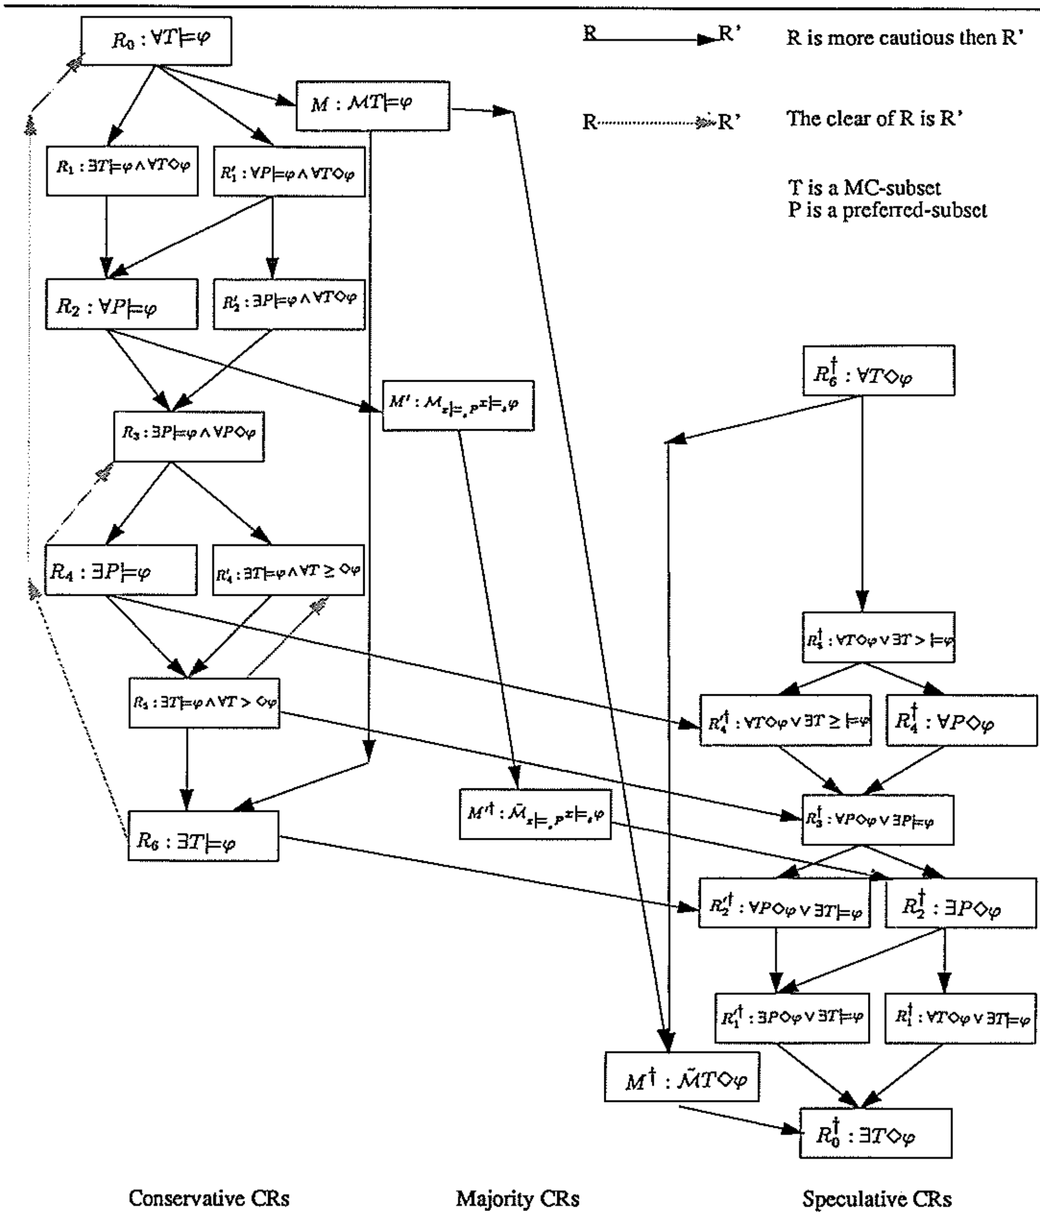

Figure 1: a small portion of the taxonomy that includes the examples given in this paper.

principle produces the same set of conclusion as the original CR (i.e, R_0). However, if ψ is incomplete, we obtain a very bold CR. It concludes φ from a set ψ iff there exists a model that satisfy both a MC-subset of ψ and the WFF φ. We call this CR, the dual of R_0 and use the notation $R_0^\dagger$. Some voting systems use similar mechanism; i.e, we agree to prefer candidate x to candidate y if not all voters prefer y to x (the dual of the Pareto relation). Also, Derthick's Mundane reasoning ([Derthick 88]) can be seen as a variation on the dual of R_0, since it is based on finding a (most likely) model that satisfies both a (most likely) subset of ψ and φ (see $R_2^\dagger$ in figure 1). We can observe also that R_0 is much more cautious than $R_0^\dagger$; any conclusion made by R_0 is also concluded by $R_0^\dagger$ and clearly they are not equal. The dual of R_0 is also not safe in the sense that both φ and $\neg\varphi$ may be concluded from some ψ. However, if we "clear" it from all these ambiguous conclusions we obtain again the original R_0 .

2.2. A partial order and some operations on consequence relations

We are now ready to formally define our taxonomy. We define the cautiousness relation, duality and the clear operation. These operations tie and relate different CRs and help to identify new interesting CRs.

A partial order of cautiousness is defined over CRs. We say that $R_1 \succ R_2$ (R_1 is more cautious than R_2) iff $R_1 \subset R_2$. If R_1 is more cautious than R_2, then every conclusion available by R_1, is also available by R_2. However, R_2 may contain some conclusions that R_1 is too cautious to conclude.

The *dual* of a CR R written as $R^{\dagger}$, is the set of of all pairs $< \psi, \varphi >$ such that $< \psi, \neg\varphi > \notin R$; i.e, $\psi R^{\dagger} \varphi$ iff $\neg(\psi R \neg\varphi)$. The dual of $\models$ corresponds to negation as failure; i.e, we conclude the negation of a proposition if we can't prove the proposition itself.

The *clear* operation clears the ambivalent conclusions from the relation; that is, it deletes from R all pairs $\psi R \varphi$ when both $\psi R \varphi$ and $\psi R \neg\varphi$ hold. We say that R is *safe* if it includes no ambivalent conclusions.

The CRs that are generated using the clear and dual operations possess the following properties:

1) The dual of a dual is the relation itself; i.e $R^{\dagger\dagger} = R$;
2) clear(clear(R)) = $clear(R)$;
3) The clear of a relation is no less cautious than the relation itself; i.e, clear(R) $\succeq R$.
4) The clear of a relation is equal to the clear of the dual of the relation; i.e, $clear(R) = clear(R^{\dagger})$;
5) If R_1 is more cautious than R_2, then the dual of R_2 is more cautious than the dual of R_1; i.e, $R_1 \succ R_2$ then $R_2^{\dagger} \succ R_1^{\dagger}$.
6) If R is safe then $clear(R) = R$; if R is also complete[2] then $R = R^{\dagger}$.

2.3. A language to describe consequence relations

To specify CRs, we'll use the language of predicate logic (as a meta language) with the standard predicates, functions and axioms for set theory augmented by the following predicates:
- Satisfiability: $x \models_s \varphi$ holds iff x is a model that satisfies φ in standard model theory.
- Preference order: $>$ is a partial order defined on consistent subsets of L. We call T "better" then T' iff $T > T'$. T and T' are non-comparable (equally good) if no order between them exists.

Using these basic predicates we can define new useful notations:
- Entailment: $\psi \models \varphi$ holds iff *all* the models that satisfy ψ also satisfy φ (classical logic entailment).
- Entailment by consistency: $\psi \Diamond \varphi$ holds iff $\neg(\psi \models \neg\varphi)$; i.e, ψ and φ are consistent (the dual of $\models$).
- Maximal consistent subset (MC-subset) predicate: $MC(\psi, T)$ holds iff T is a consistent subset of ψ that is not strictly contained in any other consistent subset of ψ.
- The preferred maximal consistent subset (P-subset) predicate: $P(\psi, T)$ holds iff T is a MC-subset of ψ and there is no "better"($>$) MC-subset of ψ.
- The majority quantifier $((\mathcal{M}_{C(x)} x) Q(x))$ denotes that $Q(x)$ holds for the majority of the x's for which $C(x)$ holds[3]. The dual $(\tilde{\mathcal{M}})$ denotes not-minority.

For example: R_0 is specified by $(\forall T) MC(\psi, T) \rightarrow (T \models \varphi)$, while $R_0^{\dagger}$ is specified by $(\exists T) MC(\psi, T) \wedge (T \Diamond \varphi)$. In figure 1, we use shorthand: T represents a MC-subset and P represents a P-subset. We therefore can use $\forall T \models \varphi$ to specify R_0. Observation: The language is closed under the "clear" and "dual" operations.

[2] R is complete iff for every ψ and φ either $\psi R \varphi$ or $\psi R \neg\varphi$.

[3] Majority is expressed in terms of cardinality: $(\mathcal{M}_{C(c)} x) P(x)$ iff $| \{x \mid C(x) \wedge Q(x)\} | > | \{x \mid C(x) \wedge \neg Q(x)\} |$.

2.4. More examples

EXAMPLE **2.1** R_2 is a safe CR defined as follows: $\psi R_2 \varphi$ iff all the P-subsets (preferred subsets) of ψ entails φ; i.e, $(\forall T)P(\psi, T) \rightarrow (T \models \varphi)$. It is a cautious CR, but many researchers find it attractive: [Rescher, Manor 70] suggested indexing methods to determine the preference relation among the MC-theories. Also, many skeptical NM systems can be looked as implementing R_2 ([Touretzky, Horty, Thomason][4]
The dual of R_2 can be described by: $\psi R_2^\dagger \varphi$ iff there exists a preferred subset P that is consistent with φ; i.e, $(\exists P)P(\psi, P) \wedge (P \Diamond \varphi)$. Note, that this is the CR that Derthick uses for his Mundane reasoning. $R_2^\dagger$ is still a very bold CR and of-course $R_2 \succ R_2^\dagger$ holds. The clear of both R_2 and $R_2^\dagger$ is R_2 itself since R_2 is a safe relation. We can also observe that $R_0 \succ R_2 \succ R_2^\dagger \succ R_0^\dagger$.

EXAMPLE **2.2** The relation R_3 is defined as follows: $\psi R_3 \varphi$ holds iff there exists a preferred subset P of ψ that entails φ and the rest of the preferred subsets are consistent with φ. Although bolder than R_2, it is still a safe CR and therefore has an advantage over mechanisms (like R_4) where a conclusion is justified based on a single preferred subset. The dual of R_3 is: $\psi R_3^\dagger \varphi$ iff all the preferred subsets are consistent with φ or at least one preferred subset entails φ.

EXAMPLE **2.3** $\psi R_4' \varphi$ holds iff any MC-subset of ψ entails φ and all better or equal MC-subsets (in the sense of the preference order $\geq$), are consistent with φ. Argument systems like [Loui 87] may be seen as using R_4': an argument (which is a proof based on some consistent subset of beliefs) wins, unless there is a better or equally good ($\geq$) subset that proves the negation. A variation of this CR is described in [Packard,Heiner 83] where a conclusion is made if a consistent subset entails φ and all better or equal consistent subsets are consistent with φ. It is still a safe CR although bolder than R_3.

EXAMPLE **2.4** $\psi R_4 \varphi$ holds iff *a* preferred subset of ψ entails φ.
It is not a safe CR, however certain credulous NM systems ([Touretzky, Horty, Thomason]), correspond to it. The dual of R_4 is: $\psi R_4^\dagger \varphi$ iff all preferred subsets are consistent with φ. The inference is done by searching for at least one model consistent with φ for any of the most likely MC-subsets. Although quite bold, $R_4^\dagger$ is more cautious than Derthick's $R_2^\dagger$. Note, that $R_0 \succ R_1 \succ R_2 \succ R_3 \succ R_4 \succ R_4^\dagger \succ R_3^\dagger \succ R_2^\dagger \succ R_1^\dagger \succ R_0^\dagger$.

Some other interesting examples are (see the graph of figure 1):

- R_1: a MC-subset entails φ and all MC-subsets are consistent with φ (mentioned in [Packard,Heiner 83]).
- R_1': All P-subsets entail φ and all the MC-subsets are consistent with φ.
- R_2': A P-subset entails φ and all the MC-subsets are consistent with φ.
- R_5: A MC-subset entails φ and all better ($>$) MC-subsets are consistent with φ.
- R_6: A MC-subset entails φ.
- M: The majority of the MC-subsets entail φ. This CR is useful for social choice systems based on democratic voting (see for example in [Campbell 79]). The dual of M is based on the quantifier "not-minority" ($\tilde{\mathcal{M}}$); i.e, at least half of the MC-subsets are consistent with φ.
- M': The majority of the models that satisfy any P-subset also satisfy φ. Although, not even a single subset has to entail φ, this CR is safe and makes sense from a probabilistic view. The dual insists that this property holds for at least half of the models (not-minority).

Finally, from a look at figure 1, we can identify three columns: The CRs from the left are called "conservative" CRs, since we can construct at least one subset that entails any conclusion of the CR (an argument supporting the conclusion). The CRs from the right are called "speculative" since a

[4] Different non-monotonic systems may be seen as if they are implementing different preference relation.

conclusion is inferred based on our ability to construct at least one plausible model that justifies the conclusion. The third column includes CRs that are based on majority (or not-minority) of models (possibly in conjunction with other types of entailment).
Note, that the standard entailment relation is a special case of any conservative CR: When we have a consistent set of beliefs, all conservative CRs conclude the same set of conclusions as the standard $\models$ relation.

3. A connectionist constraint satisfaction approach

Although all of the CR's describe in this paper can be implemented using symbolic techniques, connectionist approach has the following advantages: 1) It can be implemented on a massively parallel architecture and be used for real time applications; 2) Although the problem is NP-hard and the worse case convergence is exponential, we can trade time with correctness[5]; i.e, the network may give us an answer after arbitrary time limit and improves its probability of giving a correct answer as more time is given; 3) The ability of these architectures to learn gives our inference engines the capability of dynamically changing the knowledge base as well as the inference mechanism itself. A network may learn a new set of beliefs and develop its own intuition by acquiring a new preference relation.

In the following sections we show that most of the consequence relations discussed earlier can be implemented naturally using a constraint satisfaction approach on a connectionist network. Our main objective is to show that we can develop energy functions implementable using standard connectionist models. We want to express the constraints needed for a wide range of reasoning mechanisms, so that the global minima of the energy function developed are exactly the solutions to the reasoning problem.

3.1. Specifying weighted logical constraints in Symmetric Connectionist networks

Weighted logical constraints (both hard and soft) can be translated into energy functions that can be then implemented on a connectionist network ([Derthick 88], [Pinkas 91], [Pinkas 90b][6]. The global minima of these energy functions correspond exactly to the models that minimize the violation of the constraints. Using these algorithms, it is possible to construct basic connectionist modules that we further use as building blocks in order to express CRs as energy functions.

3.1.1. Preference order. Given a set (ψ) of propositional logic formulas augmented by real positive numbers called penalties, we define a partial order over subsets of ψ using the function $penalty_{\psi}(S) = \sum_{\varphi_i \in \psi - S} \rho_i$; i.e, we penalize S for all the φ_is in ψ that are not in S.
We say that $S_1 > S_2$ (S_1 is better than S_2), if $penalty_{\psi}(S_1) < penalty_{\psi}(S_2)$, and that two subsets are non-comparable (equally good) if the have equal penalty.
A preferred subset P of ψ, is a consistent subset such that no better consistent subset exist. The preferred-subsets (P-subsets) are therefore those MC-subsets that minimize the penalty function.

3.2. Basic connectionist modules

We describe now several useful network modules. Each module we specify has the form of
$name\text{-}of\text{-}module_{<\text{logical constraints}>}(<\text{visible units}>)$, meaning that the network module is accessible

[5] Using architectures like Boltzman machines that are able of escaping from a local minimum as more time is given.
[6] In [Derthick 88] special energy functions are generated that need special architecture. In [Pinkas 90b] quadratic energy functions are generated that can be implemented on symmetric connectionist networks like Boltzman machines or Hopfield nets.

via the <visible units> and is wired using the <logical constraints>. The visible units for each network module are used as inputs, by external clamping or by biasing, and as output to obtain an answer from the system[7].

The following notation will be used in describing the network modules: $\psi = \{< \rho_i, \varphi_i >\}$ is a set of logical beliefs (propositional formulas augmented by penalties); φ is a propositional formula (the conclusion we want to check); X is the set of visible units $X_1 \ldots X_n$, representing the atomic propositions accessible in ψ or φ; T is a set of visible units representing the φ_i's that are in ψ (the constraints $T_i \leftrightarrow \varphi_i$ must be satisfied); A is the visible unit (output) representing the answer given by the network.

We assume that the following network modules can be constructed; the implementation details appear in [Pinkas 90b] and [Pinkas 90c]:

A network that searches for a preferred model:
$NET_\psi(X)$ is a network module that allows us to combine hard and soft (weighted) logical constraints. The network then searches for a truth assignment to the atomic propositions of ψ (the visible variables in X), such that the weighted sum of the constraints that are violated ($\sum_{\vec{x} \not\models \varphi_i} \rho_i$), is minimized. In [Pinkas 90b] we show that the truth assignments are exactly the models of the preferred subsets of ψ.

A network that searches for a proof:
$PROOF_{\psi,\varphi}(T, A)$ is a network that searches for a subset of ψ that form an argument (a proof) for the query φ. If such argument in support of φ is found, the network sets A to one, and the units T_i in T represent the φ_i's that participate in the proof. The T_i's may be clamped (or biased) externally and thus force the network to answer whether a proof for φ can be constructed using this subset (T) of ψ. Note, that we can implement R_6 directly using $PROOF_{\psi,\varphi}(T, A)$.

Assuming we have these two last network modules, we can build the following ones:
A network that searches for a preferred subset:
$PREFER_\psi(X,T)$ searches for a preferred subset of ψ, and returns the subset of beliefs in the units of T; i.e, T_i is set iff φ_i is in the preferred subset that was found. The implementation is by using the module NET with the hard constraints $\{T_i \leftrightarrow \varphi_i\}$ and with the soft constraints $< \rho_i, T_i >$.

A network that searches for a proof based on a preferred subset:
The network $PREFPROOF_{\psi,\varphi}(X, T, A)$ searches for an answer to whether there exists a preferred subset of ψ that contains an argument (a proof) that supports φ. If such proof exists then A is set to one. The implementation is by combining $PROOF_{\psi,\varphi}(T, A)$ with $PREFER_\psi(X, T)$. By giving higher penalties to the constraints in $PREFER_\psi$, we bias the T_is for $PROOF_{\psi,\varphi}(T, A)$, so that a proof may be found only from a preferred subset. We can implement R_4 directly with $PREFPROOF_{\psi,\varphi}(X, T, A)$.

A network that checks for consistency with a preferred subset:
The network $PREFCONS_{\psi,\varphi}(X, A)$ checks whether there exist a preferred model (one that satisfies a P-subset of ψ) that also satisfies φ. If such model exists then A is set to one. Implementation is by combining the hard constraint ($A \leftrightarrow \varphi$) and the soft constraint $< \epsilon, A >$ with $NET_\psi(X)$. The network then tries to satisfy both a preferred-subset of ψ and φ. If it is possible to satisfy both of the constraints, then φ is consistent with a preferred subset and A is set; otherwise, the soft constraint $< \epsilon, A >$ cannot be satisfied and A becomes zero[8].

3.3. Some consequence relations that can be stated as connectionist constraint satisfaction.

All preferred subsets entail (R_2): We want to build the network $R_{2_{\psi,\varphi}}(A)$ such that $A = 1$ holds iff $\psi R_2 \varphi$. We start using the network module $PREFER_\psi(T)$ adding the hard constraints $A_1 \leftrightarrow \varphi$ and

[7] Visible units have a binary value (zero or one) and the global minima of the energy functions associated with the network projected into the visible units are the solutions to the problem specified.

[8] ϵ is a bias that is weaker then the rest of the constraints. The network should satisfy it only if the other constraints are maximally satisfied.

$\neg A_1 \leftrightarrow A$. We also add the soft bias $< \epsilon, \neg A_1 >$. The new network searches for a model that satisfies both a preferred subset of ψ and $\neg\varphi$. If it succeeds in finding such a model, we conclude that there is a preferred subset consistent with $\neg\varphi$ and therefore $< \psi, \varphi >$ is not in R_2, else we conclude that $\psi R_2 \varphi$. The ϵ bias causes the network to prefer a P-model that does not satisfy φ, causing $\neg A_1$ to hold and therefore set A to one. However, if such a preferred model cannot be found, then A_1 becomes true and A is set to zero.

A preferred subset entails and all preferred subsets are consistent (R_3):
We want to build the network $R_{3_{\psi,\varphi}}(A)$ that sets $A = 1$ iff $\psi R_3 \varphi$. We combine the network modules $PREFPROOF_{\psi,\varphi}(X, T, A_1)$ with $PREFPROOF_{\psi,\neg\varphi}(X, T, A_2)$ and add the hard constraint $(A_1 \wedge \neg A_2) \leftrightarrow A$. The new network searches for a preferred subset with an argument supporting φ and a preferred subset with an argument supporting $\neg\varphi$. Only if there exists such an argument for φ and no such argument exists for $\neg\varphi$, are we ready to set $A = 1$ meaning that $\psi R_3 \varphi$. Using the module $PROOF_{\psi,\varphi}$ instead of $PREFPROOF_{\psi,\varphi}$ we can build also the network for R_1.

A subset entails and all better or equal subsets are consistent (R'_4): We want to generate $R'_{4_{\psi,\varphi}}(A)$ that sets $A = 1$ iff there exists a subset of ψ that supports φ and there is no better or equally good subset of ψ that supports its negation. Implementation is by combining $PROOF_{\psi,\varphi}(T, A)$ with $PROOF_{\psi,\neg\varphi}(T', A')$ and adding the hard constraint $\neg A \leftrightarrow A'$ and the soft constraint $< \epsilon, A' >$. The network searches for proofs for either φ or $\neg\varphi$. Because of $\neg A \leftrightarrow A'$ only one proof will win; the one that is "better". If both proofs are equally good then the bias (ϵ) for A' causes the network to prefer the proof for $\neg\varphi$ and A becomes zero (false).

In a similar way we can implement also the rest of the conservative CRs that are in figure 1.

Implementing the dual: If R is implementable on a constraint satisfaction connectionist network, then $R^\dagger$ is also implementable on such network using the following construction:
Let $R_{\psi,\varphi}(A)$ be the network that searches for the answer to whether $\psi R \varphi$. The net $R_{\psi,\neg\varphi}(A')$ plus the hard constraint $A' \leftrightarrow \neg A$, searches for an answer to whether $\neg(\psi R \neg\varphi)$ holds. It is exactly the answer to whether $\psi R^\dagger \varphi$ holds.

Conclusion: We can implement $R_0, R_1, R'_1, R_2, R'_2, R_3, R_4, R'_4, R_5, R_6$ and all their duals. In fact we can implement any CR that can be described using a subset of the language of section 2.3. This subset includes any boolean combination[9] of basic expressions of either the form

$(\exists T)\left\{\begin{array}{c} MC(\psi,T) \\ P(\psi,T) \end{array}\right\} \wedge (\left\{\begin{array}{c} T\models \\ T\diamond \end{array}\right\}\left\{\begin{array}{c} \varphi \\ \neg\varphi \end{array}\right\})$ or the form

$(\exists T)(\left\{\begin{array}{c} MC(\psi,T) \\ P(\psi,T) \end{array}\right\}) \wedge (T\models\left\{\begin{array}{c} \varphi \\ \neg\varphi \end{array}\right\}) \wedge ((\forall T')(T'\left\{\begin{array}{c} > \\ \geq \end{array}\right\}T) \rightarrow (T'\diamond\left\{\begin{array}{c} \varphi \\ \neg\varphi \end{array}\right\}))$.

4. Summary

We have shown a taxonomy of consequence relations from possibly inconsistent set of beliefs. Many existing inference mechanisms can be mapped into this taxonomy. A partial order and several operations defined on CRs reveal an elegant structure that underlies the taxonomy and shed light on the relationships among them. Some interesting new CRs appear that deserve further research (e.g., $R'_1, R'_2, R_3, R_4^\dagger, R_6^\dagger, M'$).

Most of the CRs discussed can be stated as constraint satisfaction problems and be implemented[10] on connectionist architectures like Boltzman machines, Hopfield networks or using the special energy

[9] A boolean combination of such expressions uses the connective: $\vee, \wedge, \rightarrow$ and most important $\neg$, that enables us to use also the $\forall$, $<$ and $\leq$ inside the basic expressions.

[10] We have not figured out how to implement the majority quantifier.

functions of [Derthick 88]. In fact we show that a significant portion of the CRs that our language can express may be implemented. The capability of expressing such variety of mechanisms demonstrates the flexibility of these architectures as well as their potential to implement a wide range of knowledge level theories.

Acknowledgments: We thank Jon Doyle and Fritz Lehmann for helpful discussions.

References

[Campbell 79] D.E. Campbell, "Manipulation of social choice rules by strategic nomination of candidates", *Theory and decision* 10, 247-263, 1979.

[Derthick 88] M. Derthick, "Mundane reasoning by parallel constraint satisfaction", PhD Thesis, TR. CMU-CS-88-182, Carnegie Mellon 1988.

[Hinton89] G.E Hinton, "Learning in mean field theory", *Neural Computation* vol 1-1, 1989.

[Hinton, Sejnowski 86] G.E Hinton and T.J. Sejnowski "Learning and Re-learning in Boltzman Machines" in J. L. McClelland, D. E. Rumelhart *"Parallel Distributed Processing:Explorations in the Microstructure of Cognition"* Vol I pp. 282 - 317 MIT Press 1986

[Hopfield 84] J.J. Hopfield "Neurons with graded response have collective computational properties like those of two-state neurons". *Proceedings of the National Academy of Sciences* USA, 1984,Vol 81, pp. 3088-3092.

[Loui 87] R.P. Loui, "Defeat among arguments", Computational Intelligence, 1987.

[Packard,Heiner 83] D.J Packard, R.A. Heiner "Inconsistency resolution and collective chose", *Theory and Decision* 14, pp. 225-236, 1982.

[Pinkas 91] G. Pinkas, "Symmetric neural networks and propositional logic satisfiability", to appear in *Neural Computation* vol 3-2, 1991.

[Pinkas 90b] G. Pinkas, " NM reasoning and inconsistency using symmetric connectionist networks", to appear in the *proceedings of IJCAI-91*, Sydney, Australia, August 1991.

[Pinkas 90c] G. Pinkas, " Searching for proofs using connectionist constraint satisfaction networks", In preparation, Washington University, Dept. of Computer Science Technical report, 1991

[Reiter 1987] R. Reiter, "A logic of default reasoning", in M.L. Ginsberg (ed.) *Readings in Nonmonotonic Reasoning*, Morgan Kaufmann, Los Altos California, 1987

[Rescher, Manor 70] N. Rescher, R. Manor " On inference from inconsistent premises" *Theory and Decision* 1, pp. 179-217, 1970.

[Rescher 76] N. Rescher, " Plausible Reasoning", Van Gorcum, 1976.

[Shastri 85] L. Shastri, "Evidential reasoning in semantic networks:A formal theory and its parallel implementation", *PhD thesis, TR 166, University of Rochester*, Sept. 1985.

[Touretzky, Horty, Thomason 87] "A clash of intuitions:The current state of nonmonotonic multiple inheritance systems", in *Proceedings IJCAI 87*, pp. 476-482, Milan 1987.