Outline

Motivation

- Conventional NM logics are based on (ad hoc) modifications of the logical machinery (proofs/models).
- *Nonmonotonicity* is only a negative characterization: If we have $\Theta \mathrel{|\!\sim} \varphi$, we do not necessarily have $\Theta \cup \{\psi\} \mathrel{|\!\sim} \varphi$.
- Could we have a constructive positive characterization of default reasoning?

Plausible Consequences

- In conventional logic, we have the logical consequence relation $\alpha \models \beta$: If α is true, then also β is true.
- Instead, we will study the relation of plausible consequence $\alpha \mathrel{|\!\sim} \beta$: if α is all we know, can we conclude β?
- $\alpha \mathrel{|\!\sim} \beta$ does not imply $\alpha \wedge \alpha' \mathrel{|\!\sim} \beta$!!!
  Compare to conditional probability: $P(\beta|\alpha) \neq P(\beta|\alpha, \alpha')$!
- Find rules characterizing $\mathrel{|\!\sim}$: for example, if $\alpha \mathrel{|\!\sim} \beta$ and $\alpha \mathrel{|\!\sim} \gamma$, then $\alpha \mathrel{|\!\sim} \beta \wedge \gamma$.
- Write down all such rules!
- Perhaps we find a semantic characterization of $\mathrel{|\!\sim}$.

Desirable Properties 1: Reflexivity

- Reflexivity:

$$\frac{}{\alpha \mathrel{|\!\sim} \alpha}$$

  - Rationale: If α holds, this *normally implies* α.
  - Example: `Tom goes to a party` *normally implies* that `Tom goes to a party`.

Reflexivity in Default Logic

Plausible consequence as Reasoning in Default Logic

Let us consider relations $\mid\!\sim_\Delta$ that are defined in terms of Default Logic.
$\alpha \mid\!\sim_{\langle D,W\rangle} \beta$ means that β is a skeptical conclusion of $\langle D, W \cup \{\alpha\}\rangle \mid\!\sim \beta$.

Proposition

Default Logic satisfies Reflexivity.

Proof.

The question is: does α skeptically follow from $\Delta = \langle D, W \cup \{\alpha\}\rangle$?
For all extensions E of Δ, $W \cup \{\alpha\} \subseteq E$ by definition. Hence $\alpha \in E$ and α belongs to all extensions of Δ. □

Desirable Properties 2: Left Logical Equivalence

- Left Logical Equivalence:

$$\frac{\models \alpha \leftrightarrow \beta,\ \alpha \mid\!\sim \gamma}{\beta \mid\!\sim \gamma}$$

  - Rationale: It is not the syntactic form, but the logical content that is responsible for what we conclude normally.
  - Example: Assume that
    `Tom goes` *or* `Peter goes` *normally implies* `Mary goes`.
    Then we would expect that
    `Peter goes` *or* `Tom goes` *normally implies* `Mary goes`.

Left Logical Equivalence in Default Logic

Proposition

Default Logic satisfies Left Logical Equivalence.

Proof.

Assume that $\models \alpha \leftrightarrow \beta$ and γ is in all extensions of $\langle D, W \cup \{\alpha\}\rangle$. The definition of extensions is invariant under replacing any formula by an equivalent formula. Hence $\langle D, W \cup \{\beta\}\rangle$ has exactly the same extensions, and γ is in every one of them. □

Desirable Properties 3: Right Weakening

- Right Weakening:

$$\frac{\models \alpha \rightarrow \beta,\ \gamma \mid\!\sim \alpha}{\gamma \mid\!\sim \beta}$$

  - Rationale: If something can be concluded normally, then everything classically implied should also be concluded normally.
  - Example: Assume that
    `Mary goes` *normally implies* `Clive goes` *and* `John goes`.
    Then we would expect that
    `Mary goes` *normally implies* `Clive goes`.
  - From 1 & 3 supraclassicality follows:

$$\alpha \mid\!\sim \alpha \quad + \quad \frac{\models\alpha\rightarrow\beta,\ \alpha\mid\!\sim\alpha}{\alpha\mid\!\sim\beta} \quad \Rightarrow \quad \frac{\alpha\models\beta}{\alpha\mid\!\sim\beta}$$

Right Weakening in Default Logic

Proposition

Default Logic satisfies Right Weakening.

Proof.

Assume α is in all extensions of a default theory $\langle D, W \cup \{\gamma\}\rangle$ and $\models \alpha \rightarrow \beta$. Extensions are closed under logical consequence. Hence also β is in all extensions. □

Desirable Properties 4: Cut

- Cut:

$$\frac{\alpha \mid\!\sim \beta,\ \alpha \wedge \beta \mid\!\sim \gamma}{\alpha \mid\!\sim \gamma}$$

  - Rationale: If part of the premise is plausibly implied by another part of the premise, then the latter is enough for the plausible conclusion.
  - Example: Assume that
    `John goes` *normally implies* `Mary goes`.
    Assume further that
    `John goes` *and* `Mary goes` *normally implies* `Clive goes`.
    Then we would expect that
    `John goes` *normally implies* `Clive goes`.

Cut in Default Logic

Proposition

Default Logic satisfies Cut.

Proof idea.

Show that every extension E of $\Delta = \langle D, W \cup \{\alpha\}\rangle$ is also an extension of $\Delta' = \langle D, W \cup \{\alpha \wedge \beta\}\rangle$.
The consistency of the justifications of defaults are tested against E both in the $W \cup \{\alpha\}$ case and in the $W \cup \{\alpha \wedge \beta\}$ case.
The preconditions that are derivable when starting from $W \cup \{\alpha\}$ are also derivable when starting from $W \cup \{\alpha \wedge \beta\}$. $W \cup \{\alpha \wedge \beta\}$ does not allow deriving further preconditions because also with $W \cup \{\alpha\}$ at some point β is derived. Hence E is also an extension of Δ'.
Hence, because γ belongs to all extensions of Δ', it also belongs to all extensions of Δ. □

Desirable Properties 5: Cautious Monotonicity

- Cautious Monotonicity:

$$\frac{\alpha \mid\!\sim \beta,\ \alpha \mid\!\sim \gamma}{\alpha \wedge \beta \mid\!\sim \gamma}$$

  - Rationale: In general, adding new premises may cancel some conclusions. However, existing conclusions may be added to the premises without canceling any conclusions!
  - Example: Assume that
    `Mary goes` *normally implies* `Clive goes` and
    `Mary goes` *normally implies* `John goes`.
    `Mary goes` *and* `Jack goes` might not *normally imply* that `John goes`.
    However, `Mary goes` and `Clive goes` should *normally imply* that `John goes`.

Cautious Monotonicity in Default Logic

Proposition

Default Logic **does not** *satisfy Cautious Monotonicity.*

Proof.

Consider the default theory $\langle D, W\rangle$ with

$$D = \left\{ \frac{a:g}{g}, \frac{g:b}{b}, \frac{b:\neg g}{\neg g} \right\} \text{ and } W = \{a\}.$$

$E = \text{Th}(\{a,b,g\})$ is the only extension of $\langle D, W\rangle$ and g follows skeptically.
For $\langle D, W \cup \{b\}\rangle$ also $\text{Th}(\{a,b,\neg g\})$ is an extension, and g does not follow skeptically. □

Cumulativity

Lemma

Rules 4 & 5 can be equivalently stated as follows.

If $\alpha \mathrel{|\!\sim} \beta$, then the sets of plausible conclusions from α and $\alpha \wedge \beta$ are identical.

The above property is also called cumulativity.

Proof.

$\Rightarrow$: Assume that 4 & 5 hold and $\alpha \mathrel{|\!\sim} \beta$. Assume further that $\alpha \mathrel{|\!\sim} \gamma$. With rule 5 (CM), we have $\alpha \wedge \beta \mathrel{|\!\sim} \gamma$. Similarly, from $\alpha \wedge \beta \mathrel{|\!\sim} \gamma$ by rule 4 (Cut) we get $\alpha \mathrel{|\!\sim} \gamma$.
Hence the plausible conclusions from α and $\alpha \wedge \beta$ are the same.
$\Leftarrow$. Assume Cumulativity and $\alpha \mathrel{|\!\sim} \beta$. Now we can derive *rules 4 and 5*. □

The System **C**

1. Reflexivity
$$\frac{}{\alpha \mathrel{|\!\sim} \alpha}$$
2. Left Logical Equivalence
$$\frac{\models \alpha \leftrightarrow \beta,\ \alpha \mathrel{|\!\sim} \gamma}{\beta \mathrel{|\!\sim} \gamma}$$
3. Right Weakening
$$\frac{\models \alpha \rightarrow \beta,\ \gamma \mathrel{|\!\sim} \alpha}{\gamma \mathrel{|\!\sim} \beta}$$
4. Cut
$$\frac{\alpha \mathrel{|\!\sim} \beta,\ \alpha \wedge \beta \mathrel{|\!\sim} \gamma}{\alpha \mathrel{|\!\sim} \gamma}$$
5. Cautious Monotonicity
$$\frac{\alpha \mathrel{|\!\sim} \beta,\ \alpha \mathrel{|\!\sim} \gamma}{\alpha \wedge \beta \mathrel{|\!\sim} \gamma}$$

Derived Rules in **C**

- Equivalence:
$$\frac{\alpha \mathrel{|\!\sim} \beta,\ \beta \mathrel{|\!\sim} \alpha,\ \alpha \mathrel{|\!\sim} \gamma}{\beta \mathrel{|\!\sim} \gamma}$$
- And:
$$\frac{\alpha \mathrel{|\!\sim} \beta,\ \alpha \mathrel{|\!\sim} \gamma}{\alpha \mathrel{|\!\sim} \beta \wedge \gamma}$$
- MPC:
$$\frac{\alpha \mathrel{|\!\sim} \beta \rightarrow \gamma,\ \alpha \mathrel{|\!\sim} \beta}{\alpha \mathrel{|\!\sim} \gamma}$$

Proofs

Equivalence

Assumption:	$\alpha \mid\!\sim \beta, \quad \beta \mid\!\sim \alpha, \quad \alpha \mid\!\sim \gamma$
Cautious Monotonicity:	$\alpha \wedge \beta \mid\!\sim \gamma$
Left L Equivalence:	$\underline{\beta \wedge \alpha \mid\!\sim \gamma}$
Cut:	$\beta \mid\!\sim \gamma$

And

Assumption:	$\alpha \mid\!\sim \beta, \quad \alpha \mid\!\sim \gamma$
Cautious Monotonicity:	$\alpha \wedge \beta \mid\!\sim \gamma$
propositiona logic:	$\alpha \wedge \beta \wedge \gamma \models \beta \wedge \gamma$
Supraclassicality:	$\alpha \wedge \beta \wedge \gamma \mid\!\sim \beta \wedge \gamma$
Cut:	$\underline{\alpha \wedge \beta \mid\!\sim \beta \wedge \gamma}$
Cut:	$\alpha \mid\!\sim \beta \wedge \gamma$

MPC is an Exercise.

Undesirable Properties 1: Monotonicity and Contraposition

- Monotonicity:

$$\frac{\models \alpha \rightarrow \beta, \ \beta \mid\!\sim \gamma}{\alpha \mid\!\sim \gamma}$$

  - Example: Let us assume that
    `John goes` *normally implies* `Mary goes`.
    Now we will probably not expect that
    `John goes` *and* `Joan` (who is not in talking terms with Mary) `goes` *normally implies* `Mary goes`.

- Contraposition:

$$\frac{\alpha \mid\!\sim \beta}{\neg\beta \mid\!\sim \neg\alpha}$$

  - Example: Let us assume that
    `John goes` *normally implies* `Mary goes`.
    Would we expect that
    `Mary does not go` *normally implies* `John does not go`?
    What if `John goes` always?

Undesirable Properties 1: Monotonicity

$\alpha \models \beta, \beta \mid\!\sim \gamma$ but not $\alpha \mid\!\sim \gamma$ pictorially:

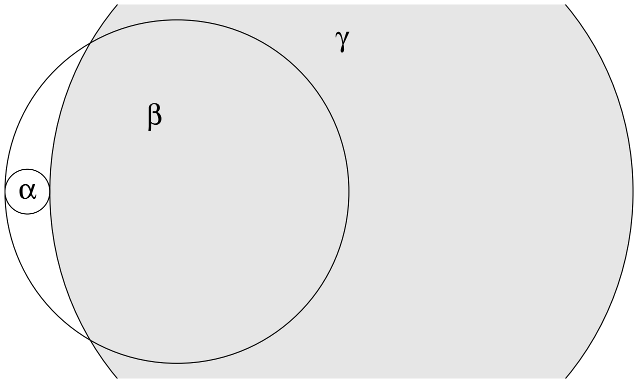

Undesirable Properties 1: Contraposition

$\alpha \mid\!\sim \beta$ but not $\neg\beta \mid\!\sim \neg\alpha$ pictorially:

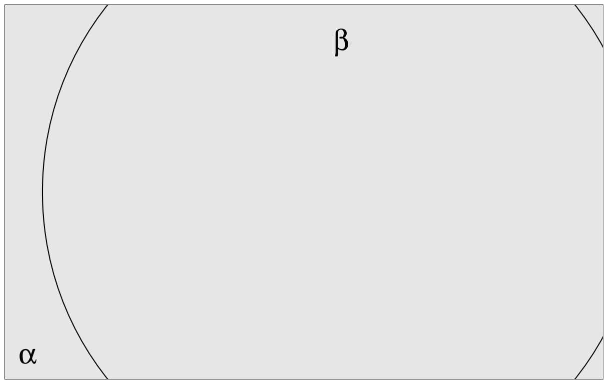

Undesirable Properties 2: Transitivity & EHD

- Transitivity:

$$\frac{\alpha \mid\!\sim \beta,\ \beta \mid\!\sim \gamma}{\alpha \mid\!\sim \gamma}$$

  - Example: Let us assume that `John goes` *normally implies* `Mary goes` and `Mary goes` *normally implies* `Jack goes`.
    Now, should `John goes` *normally imply* that `Jack goes`? If `John goes` very seldom?

- Easy Half of Deduction Theorem (*EHD*):

$$\frac{\alpha \mid\!\sim \beta \rightarrow \gamma}{\alpha \wedge \beta \mid\!\sim \gamma}$$

Undesirable Properties 2: Transitivity

$\alpha \mid\!\sim \beta,\ \beta \mid\!\sim \gamma$ but not $\alpha \mid\!\sim \gamma$ pictorially:

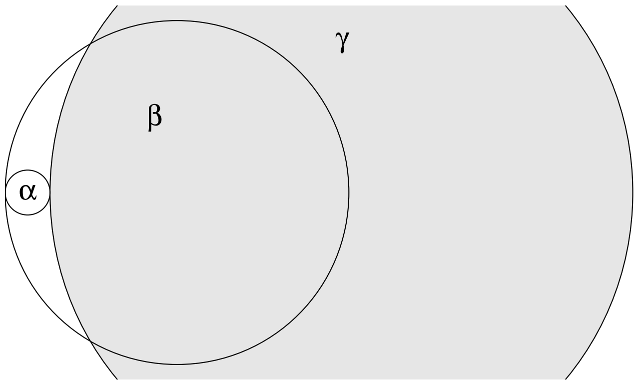

Undesirable Properties 2: EHD

$\alpha \mid\!\sim \beta \rightarrow \gamma$ but not $\alpha \wedge \beta \mid\!\sim \gamma$ pictorially:

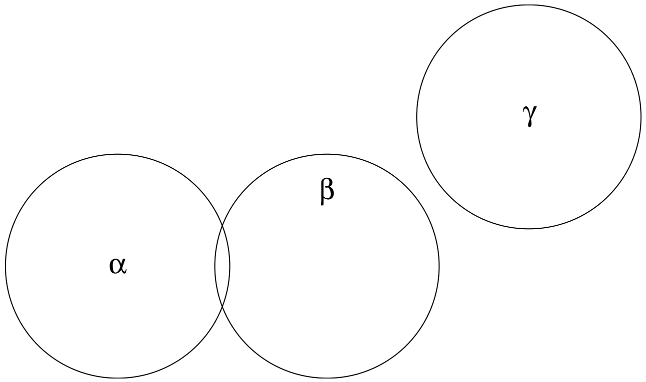

Undesirable Properties 3

Theorem

In the presence of the rules in system C, monotonicity and EHD are equivalent.

Proof.

Monotonicity $\Rightarrow$ *EHD*:

- $\alpha \mid\!\sim \beta \rightarrow \gamma$ (assumption)
- $\alpha \wedge \beta \mid\!\sim \beta \rightarrow \gamma$ (monotonicity)
- $\alpha \wedge \beta \mid\!\sim \alpha \wedge \beta$ (reflexivity)
- $\alpha \wedge \beta \mid\!\sim \beta$ (right weakening)
- $\alpha \wedge \beta \mid\!\sim \gamma$ (MPC)

Monotonicity $\Leftarrow$ *EHD*:

- $\models \alpha \rightarrow \beta, \beta \mid\!\sim \gamma$ (assumption)
- $\beta \mid\!\sim \alpha \rightarrow \gamma$ (right weakening)
- $\beta \wedge \alpha \mid\!\sim \gamma$ (EHD)
- $\alpha \mid\!\sim \gamma$ (left logical equivalence)

□

Undesirable Properties 4

Theorem
In the presence of the rules in system C, monotonicity and transitivity are equivalent.

Proof.

Monotonicity ⇒ transitivity:

- $\alpha \mathrel{|\!\sim} \beta, \beta \mathrel{|\!\sim} \gamma$ (assumption)
- $\alpha \wedge \beta \mathrel{|\!\sim} \gamma$ (monotonicity)
- $\alpha \mathrel{|\!\sim} \gamma$ (cut)

Monotonicity ⇐ transitivity:

- $\models \alpha \rightarrow \beta, \beta \mathrel{|\!\sim} \gamma$ (assumption)
- $\alpha \models \beta$ (deduction theorem)
- $\alpha \mathrel{|\!\sim} \beta$ (supraclassicality)
- $\alpha \mathrel{|\!\sim} \gamma$ (transitivity)

□

Undesirable Properties 5

Theorem
In the presence of right weakening, contraposition implies monotonicity.

Proof.

1. $\models \alpha \rightarrow \beta, \beta \mathrel{|\!\sim} \gamma$ (assumption)
2. $\neg\gamma \mathrel{|\!\sim} \neg\beta$ (contraposition)
3. $\models \neg\beta \rightarrow \neg\alpha$ (classical contraposition)
4. $\neg\gamma \mathrel{|\!\sim} \neg\alpha$ (right weakening)
5. $\alpha \mathrel{|\!\sim} \gamma$ (contraposition)

□

Note: *Monotonicity* does not imply *contraposition*, even in the presence of all rules of system C!

Cumulative Closure 1

- How do we *reason* with $\mathrel{|\!\sim}$ from φ to ψ?
- Assumption: We have a set K of conditional statements of the form $\alpha \mathrel{|\!\sim} \beta$.
  The question is: Assuming the statements in $\mathbf{K}$, is it plausible to conclude ψ given φ?
- Idea: We consider all cumulative consequence relations that contain $\mathbf{K}$.
- Remark: It suffices to consider only the minimal cumulative consequence relations containing $\mathbf{K}$.

Cumulative Closure 2

Lemma
The set of cumulative consequence relations is closed under intersection.

Proof.
Let $\mathrel{|\!\sim}_1$ and $\mathrel{|\!\sim}_2$ be cumulative consequence relations. We have to show that $\mathrel{|\!\sim}_1 \cap \mathrel{|\!\sim}_2$ is a cumulative consequence relation, that is, it satisfies the rules 1-5.
Take any instance of the any of the rules. If the preconditions are satisfied by $\mathrel{|\!\sim}_1$ and $\mathrel{|\!\sim}_2$, then the consequence is trivially also satisfied by both. □

Cumulative Closure 3

Theorem

For each finite set of conditional statements $\mathbf{K}$, there exists a unique smallest cumulative consequence relation containing $\mathbf{K}$.

Proof.

Assume the contrary, i.e., there are incomparable minimal sets $\mathbf{K}_1, \ldots, \mathbf{K}_m$. Then $\mathbf{K} = \mathbf{K}_1 \cap \ldots \cap \mathbf{K}_m$ is a unique smallest cumulative consequence relation containing $\mathbf{K}$: contradiction.

This relation is the cumulative closure $\mathbf{K}^C$ of $\mathbf{K}$. □

Cumulative Models – informally

- We will now try to characterize cumulative reasoning model-theoretically.
- Idea: Cumulative models consist of states ordered by a preference relation.
- *States* characterize beliefs.
- The *preference relation* expresses the normality of the beliefs.
- We say: $\alpha \mid\!\sim \beta$ is accepted in a model if in all most preferred states in which α is true, also β is true.

Preference Relation

- Let $\prec$ be a binary relation on a set U.
  $\prec$ is asymmetric iff

$$s \prec t \text{ implies } t \not\prec s \quad \text{for all } s,t \in U.$$

- Let $V \subseteq U$ and $\prec$ be a binary relation on U.
  - $t \in V$ is minimal in V iff $s \not\prec t$ for all $s \in V$.
  - $t \in V$ is a minimum of V (a smallest element in V) iff $t \prec s$ for all $s \in V$ such that $s \neq t$.
- Let $P \subseteq U$ and $\prec$ be a binary relation on U.
  P is smooth iff for all $t \in P$, either t is minimal in P or there is $s \in P$ such that s is minimal in P and $s \prec t$.
- Note: $\prec$ is not a partial order but an arbitrary relation!

Cumulative Models – formally

- Let $\mathcal{U}$ be the set of all *possible worlds (propositional interpretations)*.
- A cumulative model W is a triple $\langle S, l, \prec \rangle$ such that
  1. S is a set of states,
  2. l is a mapping $l: S \to 2^{\mathcal{U}}$,and
  3. $\prec$ is an arbitrary *binary relation*

  such that the smoothness condition is satisfied (see below).
- A state $s \in S$ satisfies a formula α ($s \models\!\equiv \alpha$) iff $m \models \alpha$ for all propositional interpretations $m \in l(s)$.
  The set of states satisfying α is denoted by $\widehat{\alpha}$.
- Smoothness condition: A cumulative model satisfies this condition iff for all formulae α, $\widehat{\alpha}$ is *smooth*.

Consequence Relation Induced by a Cumulative Model

A cumulative model W *induces a consequence* relation $\mid\!\sim_W$ as follows:

$$\alpha \mid\!\sim_W \beta \text{ iff } s \models\!\!\!\equiv \beta \text{ for every minimal } s \text{ in } \widehat{\alpha}.$$

Example

Model $W = \langle \{s_1, s_2, s_3\}, l, \prec \rangle$ with $s_1 \prec s_2, s_2 \prec s_3, s_1 \prec s_3$

$$\begin{aligned} l(s_1) &= \{\{\neg p, b, f\}\} \\ l(s_2) &= \{\{p, b, \neg f\}\} \\ l(s_3) &= \{\{\neg p, \neg b, f\}, \{\neg p, \neg b, \neg f\}\} \end{aligned}$$

Does W satisfy the smoothness condition?

$\neg p \wedge \neg b \mid\!\sim f$? N Also: $\neg p \wedge \neg b \not\mid\!\sim \neg f$!

$p \mid\!\sim \neg f$? Y

$\neg p \mid\!\sim f$? Y

Soundness 1

Theorem

If W is a cumulative model, then $\mid\!\sim_W$ is a cumulative consequence relation.

Proof.

- *Reflexivity:* satisfied $\surd$.
- *Left logical equivalence:* satisfied$\surd$.
- *Right weakening:* satisfied $\surd$.
- *Cut:* $\alpha \mid\!\sim \beta, \alpha \wedge \beta \mid\!\sim \gamma \Rightarrow \alpha \mid\!\sim \gamma$. Assume that all minimal elements of $\widehat{\alpha}$ satisfy β, and all minimal elements of $\widehat{\alpha \wedge \beta}$ satisfy γ. Every minimal element of $\widehat{\alpha}$ satisfies $\alpha \wedge \beta$.Since $\widehat{\alpha \wedge \beta} \subseteq \widehat{\alpha}$, all minimal elements of $\widehat{\alpha}$ are also minimal elements of $\widehat{\alpha \wedge \beta}$. Hence $\alpha \mid\!\sim_W \gamma$.

$\square$

Soundness 2

Proof continues...

- *Cautious Monotonicity:* To show: $\alpha \mid\!\sim \beta, \alpha \mid\!\sim \gamma \Rightarrow \alpha \wedge \beta \mid\!\sim \gamma$.

  Assume $\alpha \mid\!\sim_W \beta$ and $\alpha \mid\!\sim_W \gamma$. We have to show: $\alpha \wedge \beta \mid\!\sim_W \gamma$, i.e., $s \models\!\!\!\equiv \gamma$ for all minimal $s \in \widehat{\alpha \wedge \beta}$.

  Clearly, every minimal $s \in \widehat{\alpha \wedge \beta}$ is in $\widehat{\alpha}$.

  We show that every minimal $s \in \widehat{\alpha \wedge \beta}$ is *minimal* in $\widehat{\alpha}$.

  Assumption: There is s that is minimal in $\widehat{\alpha \wedge \beta}$ but not minimal in $\widehat{\alpha}$. Because of *smoothness* there is minimal $s' \in \widehat{\alpha}$ such that $s' \prec s$. We know, however, that $s' \models\!\!\!\equiv \beta$, which means that $s' \in \widehat{\alpha \wedge \beta}$. Hence s is not minimal in $\widehat{\alpha \wedge \beta}$. Contradiction! Hence s must be minimal in $\widehat{\alpha}$, and therefore $s \models\!\!\!\equiv \gamma$. Because this is true for all minimal elements in $\widehat{\alpha \wedge \beta}$, we get $\alpha \wedge \beta \mid\!\sim_W \gamma$.

$\square$

Consequence: Counterexamples

Now we have a method for showing that a principle does not hold for cumulative consequence relations. $\rightsquigarrow$ Simply construct a cumulative model that falsifies the principle.

Contraposition: $\alpha \mid\!\sim \beta \Rightarrow \neg\beta \mid\!\sim \neg\alpha$

$$\begin{aligned} W &= \langle S, l, \prec \rangle \\ S &= \{s_1, s_2, s_3, s_4\}, s_i \not\prec s_j \ \forall s_i, s_j \in S \\ l(s_1) &= \{\{a, b\}\} \\ l(s_2) &= \{\{\neg a, b\}, \{\neg a, \neg b\}\} \\ l(s_3) &= \{\{b, a\}, \{b, \neg a\}\} \\ l(s_4) &= \{\{\neg b, a\}, \{\neg b, \neg a\}\} \end{aligned}$$

W is a cumulative model with $a \mid\!\sim_W b$ but $\neg b \not\mid\!\sim_W \neg a$.

Completeness?

- Each cumulative model W *induces* a cumulative consequence relation $\mid\!\sim_W$.
- Problem: Can we generate all cumulative consequence relations in this way?
- We can! There is a representation theorem: For each cumulative consequence relation, there is a cumulative model and *vice versa*.
- Advantage: We have a characterization of the cumulative consequence independently from the set of inference rules.

Transitivity of the Preference Relation?

- Could we strengthen the preference relation to transitive relations without sacrificing anything?
  No!
- In such models, the following additional principle called Loop is valid:

$$\frac{\alpha_o \mid\!\sim \alpha_1, \alpha_1 \mid\!\sim \alpha_2, \ldots, \alpha_k \mid\!\sim \alpha_0}{\alpha_0 \mid\!\sim \alpha_k}$$

- For the system CL $=$ C $+$ Loop and cumulative models with transitive preference relations, we could prove another *representation theorem*.

The Or Rule

Or rule:

$$\frac{\alpha \mid\!\sim \gamma,\ \beta \mid\!\sim \gamma}{\alpha \vee \beta \mid\!\sim \gamma}$$

Not true in C. Counterexample:

$$\begin{aligned} W &= \langle S, l, \prec \rangle \\ S &= \{s_1, s_2, s_3\}, s_i \not\prec s_j\ \forall s_i, s_j \in S \\ l(s_1) &= \{\{a,b,c\},\{a,\neg b,c\}\} \\ l(s_2) &= \{\{b,a,c\},\{b,\neg a,c\}\} \\ l(s_3) &= \{\{a,b,\neg c\},\{a,\neg b,\neg c\},\{\neg a,b,\neg c\}\} \end{aligned}$$

$a \mid\!\sim_W c$, $b \mid\!\sim_W c$ but $a \vee b \not\mid\!\sim_W c$.
Note: Or is not valid in DL.

System P

- System P contains all rules of C and the Or rule.
- Derived rules in P:
  - Hard half of deduction theorem (S):

$$\frac{\alpha \wedge \beta \mid\!\sim \gamma}{\alpha \mid\!\sim \beta \rightarrow \gamma}$$

  - Proof by case analysis (D):

$$\frac{\alpha \wedge \neg\beta \mid\!\sim \gamma,\ \alpha \wedge \beta \mid\!\sim \gamma}{\alpha \mid\!\sim \gamma}$$

- D and Or are equivalent in the presence of the rules in C.

Preferential Models

Definition

A cumulative model $W = \langle S, l, \prec \rangle$ such that $\prec$ is a *strict partial order* (irreflexive and transitive) and $|l(s)| = 1$ for all $s \in S$ is a preferential model.

Preferential Models

Theorem (Soundness)

The consequence relation $\mid\sim_W$ induced by a preferential model is preferential.

Proof.

Since W is cumulative, we only have to verify that Or holds. Note that in preferential models we have $\widehat{\alpha \vee \beta} = \widehat{\alpha} \cup \widehat{\beta}$. Suppose $\alpha \mid\sim_W \gamma$ and $\beta \mid\sim_W \gamma$. Because of the above equation, each minimal state of $\widehat{\alpha \vee \beta}$ is minimal in $\widehat{\alpha} \cup \widehat{\beta}$. Since γ is satisfied in all minimal states in $\widehat{\alpha} \cup \widehat{\beta}$, γ is also satisfied in all minimal states of $\widehat{\alpha \vee \beta}$. Hence $\alpha \vee \beta \mid\sim_W \gamma$. □

Preferential Models

Theorem (Representation)

A consequence relation is preferential iff it is induced by a preferential model.

Proof.

Similar to the one for **C**. □

Summary of Consequence Relations

System	*Models*
C	
Reflexivity	States: sets of worlds
Left Logical Equivalence	Preference relation: arbitrary
Right Weakening	Models must be smooth
Cut	
Cautious Monotonicity	
CL	
+ Loop	Preference relation: strict partial order
P	
+ Or	States: singletons

Strengthening the Consequence Relation

- System C and System P do not produce many of the inferences one would hope for:

  Given $K = \{\text{Bird} \mid\!\sim \text{Flies}\}$ one cannot conclude $\text{Red} \wedge \text{Bird} \mid\!\sim \text{Flies}$!!
- In general, adding information that is irrelevant cancels the plausible conclusions. $\Longrightarrow$ Cumulative and Preferential consequence relations are too nonmonotonic.
- The plausible conclusions have to be strengthened!!

Strengthening the Consequence Relations

- The rules so far seem to be reasonable and one cannot think of rules of the same form (if we have some plausible implications, other plausible implications should hold) that could be added.
- However, there are other types of rules one might want add.
- Disjunctive Rationality:

$$\frac{\alpha \not\mid\!\sim \gamma,\ \beta \not\mid\!\sim \gamma}{\alpha \vee \beta \not\mid\!\sim \gamma}$$

- Rational Monotonicity:

$$\frac{\alpha \mid\!\sim \gamma,\ \alpha \not\mid\!\sim \neg\beta}{\alpha \wedge \beta \mid\!\sim \gamma}$$

- Note: Consequence relations obeying these rules are not closed under intersection, which is a problem.

Probabilistic View of Plausible Consequences

- Consider probability distributions P on the set $\mathcal{M}$ of all propositional interpretations $m \in \mathcal{M}$ of our language.
- $P(m)$ is the probability of the possible world m.
- Extend this to probability of formulae:

$$P(\alpha) = \sum \{P(m) | m \in \mathcal{M}, m \models \alpha\}$$

- Conditional probability is defined in the standard way.

$$P(\beta|\alpha) = \frac{P(\alpha \wedge \beta)}{P(\alpha)}$$

ε-Entailment

Definition

$\alpha \mid\!\sim \beta$ is ε-entailed by a set K iff for all $\varepsilon > 0$ there is $\delta > 0$ such that $P(\beta|\alpha) \geq 1 - \varepsilon$ for all probability distributions P such that $P(\beta'|\alpha') \geq 1 - \delta$ for all $\alpha' \mid\!\sim \beta' \in K$.

ε-Entailment: Example

One probability distribution P such that $P(f|b) \geq 0.9$, $P(\neg f|p) \geq 0.9$ and $P(b|p) \geq 0.9$ is the following.

	p	b	f	P
w_1	0	0	0	0.00
w_2	0	0	1	0.00
w_3	0	1	0	0.00
w_4	0	1	1	0.99
w_5	1	0	0	0.00
w_6	1	0	1	0.00
w_7	1	1	0	0.01
w_8	1	1	1	0.00

$$P(f|b) = \frac{P(w_4)+P(w_8)}{P(w_3)+P(w_4)+P(w_7)+P(w_8)} = \frac{0.99}{1.00}$$

$$P(\neg f|p) = \frac{P(w_5)+P(w_7)}{P(w_5)+P(w_6)+P(w_7)+P(w_8)} = \frac{0.01}{0.01}$$

$$P(b|p) = \frac{P(w_7)+P(w_8)}{P(w_5)+P(w_6)+P(w_7)+P(w_8)} = \frac{0.01}{0.01}$$

Properties of ε-Entailment

Theorem

$\alpha \mid\!\sim \beta$ *is in all preferential consequence relations that include* K *if and only if* $\alpha \mid\!\sim \beta$ *is* ε*-entailed by* K.

So, System P provides a proof system that exactly corresponds to ε-entailment.

Weakness of ε-Entailment

- Question: Why is Eagle $\mid\!\sim$ Flies not an ε-consequence of $K = \{$Eagle $\mid\!\sim$ Bird, Bird $\mid\!\sim$ Flies$\}$?
- Answer: Because there are probability distributions that simultaneously assign very high probabilities to $P($Bird$|$Eagle$)$ and $P($Flies$|$Bird$)$ and a low probability to $P($Flies$|$Eagle$)$.
- K does not justify the low probability of $P($Flies$|$Eagle$)$: there are exactly as many worlds satisfying Bird $\wedge$ Eagle $\wedge$ Flies and Bird $\wedge$ Eagle $\wedge \neg$Flies, and the worlds satisfying Bird $\wedge$ Flies have a much higher probability that those satisfying Bird $\wedge \neg$Flies.
  Why should the probabilities for eagles be the other way round?
- We would like to restrict to probability distributions that are not biased toward non-flying eagles without a reason.

Entropy of a Probability Distribution

Definition

The entropy of a probability distribution P is

$$H(P) = -\sum_{m \in \mathcal{M}} P(m) \log P(m)$$

The probability distribution with the highest entropy is the one that assigns the same probability to every world.

ME-Entailment

Definition

$\alpha \mathrel{|\!\sim} \beta$ is ME-entailed by a set K iff for all $\varepsilon > 0$ there is $\delta > 0$ such that $P(\beta|\alpha) \geq 1 - \varepsilon$ for the distribution P that has the maximum entropy among distributions satisfying $P(\beta'|\alpha') \geq 1 - \delta$ for all $\alpha' \mathrel{|\!\sim} \beta' \in K$.

Entropy of a Probability Distribution: Example

The distribution P that has the maximum entropy among distributions such that $P(b|e) \geq 0.9$ and $P(f|b) \geq 0.9$ is the following.

	e	b	f	P
w_1	0	0	0	0.1875
w_2	0	0	1	0.1875
w_3	0	1	0	0.0292
w_4	0	1	1	0.1875
w_5	1	0	0	0.0204
w_6	1	0	1	0.0204
w_7	1	1	0	0.0292
w_8	1	1	1	0.3380

$$P(f|b) = \frac{P(w_4)+P(w_8)}{P(w_3)+P(w_4)+P(w_7)+P(w_8)} = \frac{0.5255}{0.5839}$$

$$P(b|e) = \frac{P(w_7)+P(w_8)}{P(w_5)+P(w_6)+P(w_7)+P(w_8)} = \frac{0.3672}{0.4080}$$

$$P(f|e) = \frac{P(w_6)+P(w_8)}{P(w_5)+P(w_6)+P(w_7)+P(w_8)} = \frac{0.3584}{0.4080} = 0.8784$$

ME-Entailment: Examples

1. $\{$Eagle $\mathrel{|\!\sim}$ Bird, Bird $\mathrel{|\!\sim}$ Flies$\}$ ME-entails Eagle $\mathrel{|\!\sim}$ Flies
2. $\{$Penguin $\mathrel{|\!\sim}$ Bird, Bird $\mathrel{|\!\sim}$ Flies, Penguin $\mathrel{|\!\sim}$ $\neg$Flies$\}$ ME-entails Bird $\wedge$ Penguin $\mathrel{|\!\sim}$ $\neg$Flies
3. $\{$Eagle $\mathrel{|\!\sim}$ Bird$\}$ ME-entails $\neg$Bird $\mathrel{|\!\sim}$ $\neg$Eagle

Summary

- Instead of *ad hoc* extensions of the logical machinery, analyze the properties of nonmonotonic consequence relations.
- Correspondence between rule system and models for System C, and for System P also wrt a probabilistic semantics.
- Irrelevant information poses a problem. Solution approaches: rational monotony, maximum entropy.

Literature

Sarit Kraus, Daniel Lehmann, and Menachem Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.

Introduces Cumulative consequence relations.

Daniel Lehman and Menachem Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55:1–60, 1992. Introduces rational consequence relations.

Dov M. Gabbay. Theoretical foundations for non-monotonic reasoning in expert systems. In K. R. Apt, editor, *Proceedings NATO Advanced Study Institute on Logics and Models of Concurrent Systems*, pages 439–457. Springer-Verlag, Berlin, Heidelberg, New York, 1985.

First to consider abstract properties of nonmonotonic consequence relations.

Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann Publishers, 1988. One section on ε-semantics and maximum entropy.

Yoav Shoham. *Reasoning about Change*. MIT Press, Cambridge, MA, 1988. Introduces the idea of preferential models.