

Principles of Knowledge Representation and Reasoning

Nonmonotonic Reasoning III: Cumulative Logics

Bernhard Nebel, Malte Helmert and Stefan Wölfel

Albert-Ludwigs-Universität Freiburg

June 3, 2008

Principles of Knowledge Representation and Reasoning

June 3, 2008 — Nonmonotonic Reasoning III: Cumulative Logics

Definition of Cumulative Logics

- Motivation

- Properties

- Derived Rules in **C**

- Undesirable Properties

Reasoning

- Cumulative Closure

Semantics

- Cumulative Models

- Consequence Relations

Preferential Reasoning

- Preferential Relations

Probabilistic Semantics

- ϵ -Semantics

- Maximum Entropy

Literature

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 $\Theta \vdash \varphi$, we do not necessarily have $\Theta \cup \{\psi\} \vdash \varphi$.
- ▶ 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** $\alpha \sim \beta$: if α is all we know, can we conclude β ?
- ▶ $\alpha \sim \beta$ does not imply $\alpha \wedge \alpha' \sim \beta$!
Compare to conditional probability: $P(\beta|\alpha) \neq P(\beta|\alpha, \alpha')$!
- ▶ Find rules characterizing \sim : for example, if $\alpha \sim \beta$ and $\alpha \sim \gamma$, then $\alpha \sim \beta \wedge \gamma$.
- ▶ Write down all such rules!
- ▶ Perhaps we find a **semantic characterization** of \sim .

Desirable Properties 1: Reflexivity

► Reflexivity:

$$\overline{\alpha \sim \alpha}$$

- **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 \sim_{Δ} that are defined in terms of Default Logic.
 $\alpha \sim_{\langle D, W \rangle} \beta$ means that β is a skeptical conclusion of $\langle D, W \cup \{\alpha\} \rangle$.

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:**

$$\frac{\models \alpha \leftrightarrow \beta, \alpha \vdash \gamma}{\beta \vdash \gamma}$$

- **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:**

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

- **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:

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

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. \square

Desirable Properties 4: Cut

► **Cut:**

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

- **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 $\Delta' = \langle D, W \cup \{\alpha \wedge \beta\} \rangle$.

Consistency of justifications of defaults is 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:**

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

- **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 $\alpha \vdash \beta$, 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 $\alpha \vdash \beta$. Assume further that $\alpha \vdash \gamma$. With **rule 5** (CM), we have $\alpha \wedge \beta \vdash \gamma$. Similarly, from $\alpha \wedge \beta \vdash \gamma$ by **rule 4** (Cut) we get $\alpha \vdash \gamma$.

Hence the plausible conclusions from α and $\alpha \wedge \beta$ are the same.

\Leftarrow . Assume **Cumulativity** and $\alpha \vdash \beta$. Now we can derive **rules 4 and 5**. □

The System C

1. Reflexivity

$$\frac{}{\alpha \sim \alpha}$$

2. Left Logical Equivalence

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

3. Right Weakening

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

4. Cut

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

5. Cautious Monotonicity

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

Derived Rules in C

► *Equivalence:*

$$\frac{\alpha \vdash \beta, \beta \vdash \alpha, \alpha \vdash \gamma}{\beta \vdash \gamma}$$

► *And:*

$$\frac{\alpha \vdash \beta, \alpha \vdash \gamma}{\alpha \vdash \beta \wedge \gamma}$$

► *MPC:*

$$\frac{\alpha \vdash \beta \rightarrow \gamma, \alpha \vdash \beta}{\alpha \vdash \gamma}$$

Proofs

Equivalence

Assumption: $\alpha \vdash \beta, \beta \vdash \alpha, \alpha \vdash \gamma$

Cautious Monotonicity: $\alpha \wedge \beta \vdash \gamma$

Left L Equivalence: $\frac{\beta \wedge \alpha}{\beta} \vdash \gamma$

Cut: $\frac{\beta \vdash \gamma}{\beta} \vdash \gamma$

And

Assumption: $\alpha \vdash \beta, \alpha \vdash \gamma$

Cautious Monotonicity: $\alpha \wedge \beta \vdash \gamma$

propositional logic: $\alpha \wedge \beta \wedge \gamma \models \beta \wedge \gamma$

Supraclassicality: $\alpha \wedge \beta \wedge \gamma \vdash \beta \wedge \gamma$

Cut: $\frac{\alpha \wedge \beta \vdash \beta \wedge \gamma}{\alpha \wedge \beta} \vdash \beta \wedge \gamma$

Cut: $\frac{\alpha \wedge \beta \vdash \beta \wedge \gamma}{\alpha} \vdash \beta \wedge \gamma$

MPC is an Exercise.

Undesirable Properties 1: Monotonicity and Contraposition

▶ Monotonicity:

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

- ▶ **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:

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

- ▶ **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

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



Undesirable Properties 1: Contraposition

$\alpha \vdash \beta$ but not $\neg\beta \vdash \neg\alpha$ pictorially:



Undesirable Properties 2: Transitivity & EHD

► **Transitivity:**

$$\frac{\alpha \vdash \beta, \beta \vdash \gamma}{\alpha \vdash \gamma}$$

► **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):**

$$\frac{\alpha \vdash \beta \rightarrow \gamma}{\alpha \wedge \beta \vdash \gamma}$$


Undesirable Properties 2: Transitivity

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



Undesirable Properties 2: EHD

$\alpha \vdash \beta \rightarrow \gamma$ but not $\alpha \wedge \beta \vdash \gamma$ pictorially:



Undesirable Properties 3

Theorem

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

Proof.

Monotonicity \Rightarrow *EHD*:

- ▶ $\alpha \vdash \beta \rightarrow \gamma$ (assumption)
- ▶ $\alpha \wedge \beta \vdash \beta \rightarrow \gamma$ (monotonicity)
- ▶ $\alpha \wedge \beta \vdash \alpha \wedge \beta$ (reflexivity)
- ▶ $\alpha \wedge \beta \vdash \beta$ (right weakening)
- ▶ $\alpha \wedge \beta \vdash \gamma$ (MPC)

Monotonicity \Leftarrow *EHD*:

- ▶ $\models \alpha \rightarrow \beta, \beta \vdash \gamma$ (assumption)
- ▶ $\beta \vdash \alpha \rightarrow \gamma$ (right weakening)
- ▶ $\beta \wedge \alpha \vdash \gamma$ (EHD)
- ▶ $\alpha \vdash \gamma$ (left logical equivalence)

□

Undesirable Properties 4

Theorem

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

Proof.

Monotonicity \Rightarrow *transitivity*:

- ▶ $\alpha \vdash \beta, \beta \vdash \gamma$ (assumption)
- ▶ $\alpha \wedge \beta \vdash \gamma$ (monotonicity)
- ▶ $\alpha \vdash \gamma$ (cut)

Monotonicity \Leftarrow *transitivity*:

- ▶ $\models \alpha \rightarrow \beta, \beta \vdash \gamma$ (assumption)
- ▶ $\alpha \models \beta$ (deduction theorem)
- ▶ $\alpha \vdash \beta$ (supraclassicality)
- ▶ $\alpha \vdash \gamma$ (transitivity)



Undesirable Properties 5

Theorem

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

Proof.

1. $\models \alpha \rightarrow \beta, \beta \vdash \gamma$ (assumption)
2. $\neg \gamma \vdash \neg \beta$ (contraposition)
3. $\models \neg \beta \rightarrow \neg \alpha$ (classical contraposition)
4. $\neg \gamma \vdash \neg \alpha$ (right weakening)
5. $\alpha \vdash \gamma$ (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 \vdash from φ to ψ ?
- ▶ **Assumption**: We have a set **K** of **conditional statements** of the form $\alpha \vdash \beta$.
The question is: Assuming the statements in **K**, is it plausible to conclude ψ given φ ?
- ▶ **Idea**: We consider **all** cumulative consequence relations that contain **K**.
- ▶ **Remark**: It suffices to consider only the **minimal** cumulative consequence relations containing **K**.

Cumulative Closure 2

Lemma

The set of cumulative consequence relations is closed under intersection.

Proof.

Let \vdash_1 and \vdash_2 be cumulative consequence relations. We have to show that $\vdash_1 \cap \vdash_2$ 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 \vdash_1 and \vdash_2 , then the consequence is trivially also satisfied by both. \square

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 $\mathbf{K}_1, \dots, \mathbf{K}_m$. Then $\mathbf{K} = \mathbf{K}_1 \cap \dots \cap \mathbf{K}_m$ 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: $\alpha \sim \beta$ 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

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

- ▶ 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 $\langle S, I, \prec \rangle$ such that
 1. S is a set of **states**,
 2. I is a mapping $I : S \rightarrow 2^{\mathcal{U}}$, 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 ($s \models \alpha$) iff $m \models \alpha$ for all propositional interpretations $m \in I(s)$.
The set of states satisfying α is denoted by $\hat{\alpha}$.
- ▶ **Smoothness condition:** A cumulative model satisfies this condition iff for all formulae α , $\hat{\alpha}$ is **smooth**.

Consequence Relation Induced by a Cumulative Model

A cumulative model W induces a consequence relation \vdash_W as follows:

$$\alpha \vdash_W \beta \text{ iff } s \models \beta \text{ for every minimal } s \text{ in } \widehat{\alpha}.$$

Example

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

$$I(s_1) = \{ \{ \neg p, b, f \} \}$$

$$I(s_2) = \{ \{ p, b, \neg f \} \}$$

$$I(s_3) = \{ \{ \neg p, \neg b, f \}, \{ \neg p, \neg b, \neg f \} \}$$

Does W satisfy the smoothness condition?

$$\neg p \wedge \neg b \vdash f? \quad \text{N} \quad \text{Also: } \neg p \wedge \neg b \not\vdash \neg f!$$

$$p \vdash \neg f? \quad \text{Y}$$

$$\neg p \vdash f? \quad \text{Y}$$

Soundness 1

Theorem

If W is a cumulative model, then \vdash_W is a cumulative consequence relation.

Proof.

- ▶ *Reflexivity*: satisfied \checkmark .
- ▶ *Left logical equivalence*: satisfied \checkmark .
- ▶ *Right weakening*: satisfied \checkmark .
- ▶ *Cut*: $\alpha \vdash \beta, \alpha \wedge \beta \vdash \gamma \Rightarrow \alpha \vdash \gamma$. 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 $\alpha \vdash_W \gamma$.

□

Soundness 2

Proof continues...

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

Assume $\alpha \vdash_W \beta$ and $\alpha \vdash_W \gamma$. We have to show: $\alpha \wedge \beta \vdash_W \gamma$, i.e., $s \models \gamma$ 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 $s' \models \beta$, 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 $s \models \gamma$. Because this is true for all minimal elements in $\widehat{\alpha \wedge \beta}$, we get $\alpha \wedge \beta \vdash_W \gamma$.

□

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: $\alpha \vdash \beta \Rightarrow \neg\beta \vdash \neg\alpha$

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

W is a cumulative model with $a \vdash_W b$ but $\neg b \not\vdash_W \neg a$.

Completeness?

- ▶ Each cumulative model W *induces* a cumulative consequence relation \vdash_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:

$$\frac{\alpha_0 \sim \alpha_1, \alpha_1 \sim \alpha_2, \dots, \alpha_k \sim \alpha_0}{\alpha_0 \sim \alpha_k}$$

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

The Or Rule

Or rule:

$$\frac{\alpha \vDash \gamma, \beta \vDash \gamma}{\alpha \vee \beta \vDash \gamma}$$

Not true in **C**. Counterexample:

$$W = \langle S, I, \prec \rangle$$

$$S = \{s_1, s_2, s_3\}, s_i \not\prec s_j \forall s_i, s_j \in S$$

$$I(s_1) = \{\{a, b, c\}, \{a, \neg b, c\}\}$$

$$I(s_2) = \{\{a, b, c\}, \{\neg a, b, c\}\}$$

$$I(s_3) = \{\{a, b, \neg c\}, \{a, \neg b, \neg c\}, \{\neg a, b, \neg c\}\}$$

$a \vDash_W c, b \vDash_W c$ but $a \vee b \not\vDash_W c$.

Note: Or is not valid in DL.

System P

- ▶ System **P** contains all rules of **C** and the **Or** rule.
- ▶ A consequence relation that satisfies **P** is called **preferential**.
- ▶ Derived rules in **P**:
 - ▶ Hard half of deduction theorem (**S**):

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

- ▶ Proof by case analysis (**D**):

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

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

Preferential Models

Definition

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

Preferential Models

Theorem (Soundness)

The consequence relation \vdash_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 $\alpha \vdash_W \gamma$ and $\beta \vdash_W \gamma$. 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 $\alpha \vee \beta \vdash_W \gamma$. □

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

C

Reflexivity

Left Logical Equivalence

Right Weakening

Cut

Cautious Monotonicity

CL

+ Loop

P

+ Or

Models

States: sets of worlds

Preference relation: arbitrary

Models must be smooth

Preference relation: strict partial order

States: singletons

Strengthening the Consequence Relation

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

*Given $K = \{Bird \sim Flies\}$ one cannot conclude
 $Red \wedge Bird \sim Flies!$*

- ▶ In general, adding information that is **irrelevant** cancels the plausible conclusions. \implies 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:**

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

- ▶ **Rational Monotonicity:**

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

- ▶ **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:

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

- ▶ **Conditional probability** is defined in the standard way.

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

ϵ -Entailment

Definition

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

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

$\alpha \vdash \beta$ is in all preferential consequence relations that include K if and only if $\alpha \vdash \beta$ 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 $\text{Eagle} \sim \text{Flies}$ not an ϵ -consequence of $K = \{\text{Eagle} \sim \text{Bird}, \text{Bird} \sim \text{Flies}\}$?
- ▶ **Answer:** Because there are probability distributions that simultaneously assign very high probabilities to $P(\text{Bird}|\text{Eagle})$ and $P(\text{Flies}|\text{Bird})$ and a low probability to $P(\text{Flies}|\text{Eagle})$.
- ▶ K does not justify the low probability of $P(\text{Flies}|\text{Eagle})$: there are exactly as many worlds satisfying $\text{Bird} \wedge \text{Eagle} \wedge \text{Flies}$ and $\text{Bird} \wedge \text{Eagle} \wedge \neg\text{Flies}$, and the worlds satisfying $\text{Bird} \wedge \text{Flies}$ have a much higher probability than those satisfying $\text{Bird} \wedge \neg\text{Flies}$. 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

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

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

$$\begin{aligned}
 P(f|b) &= \frac{P(w_4)+P(w_8)}{P(w_3)+P(w_4)+P(w_7)+P(w_8)} \\
 &= \frac{0.5255}{0.5839} = 0.9000
 \end{aligned}$$

$$\begin{aligned}
 P(b|e) &= \frac{P(w_7)+P(w_8)}{P(w_5)+P(w_6)+P(w_7)+P(w_8)} \\
 &= \frac{0.3672}{0.4080} = 0.9000
 \end{aligned}$$

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

ME-Entailment: Examples

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

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 monotonicity, maximum entropy.

Literature I



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.

Literature II



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.