TECHNISCHE UNIVERSITÄT DRESDEN

Computational Logic Group

Sebastian Rudolph

(based on slides by Jonas Karge)

Faculty of Computer Science, Institute of Artificial Intelligence, Computational Logic Group

Inconsistency Handling

Lecture 9, 11th Dec 2022 // Foundations of Knowledge Representation, WS 2023/24

Inconsistency: Motivational Example

Suppose that you have a knowledge base that contains, among other things, the following pieces of information (in some formal language):

a : All European swans are white.
b : The bird caught in the trap is a swan.
c : The bird caught in the trap comes from Sweden.
d : Sweden is part of Europe.

From $a - d$, infer e: The bird caught in the trap is white.

Now suppose that: The bird caught in the trap turns out to be black.

Problem: By adding $\neg e$ the knowledge base becomes inconsistent.

Inconsistency Handling – Overview

We can distinguish between two general approaches to handle inconsistency:

(1) Paraconsistent Reasoning, i.e. reasoning from inconsistent knowledge bases:

- One option is to consider only consistent subsets of the database *(i.e. maximal consistent subsets)*.
- Another option is to leave the database inconsistent, but to prohibit the logics from deriving trivial inferences *(i.e. paraconsistent logics)*.

(2) Belief Revision, integrate knowledge into a given knowledge base and preserve consistency.

Inconsistency Handling – Outline

Paraconsistent Reasoning
- Maximal Consistent Subsets
- Four-Valued Logic

Belief Revision
- Motivation and Preliminaries
- AGM
- Contraction
- Epistemic Entrechment

Paraconsistent Reasoning

Maximal Consistent Subsets

A very natural approach to draw plausible conclusions from an inconsistent knowledge base considers the maximal consistent subsets of the base.

Consistency

A set Γ of propositions is **consistent** iff
there is no φ such that $\Gamma \vdash \varphi$ and $\Gamma \vdash \neg\varphi$.

Maximal Consistent

A set Γ of propositions is **maximal consistent** iff

- Γ is consistent, and
- $\Gamma \subseteq \Gamma'$ and Γ' consistent $\Rightarrow \Gamma = \Gamma'$.

Maximal Consistent Subsets

Goal: Define an inference relation based on an inference mechanism that is based on maximal consistent subsets.

Consider the following ones:

(1) Skeptical inference, where a formula is a consequence of the base iff it is entailed by each maximal consistent subset,

(2) Credulous inference, where a formula is a consequence of the base iff it is entailed by at least one maximal consistent subset,

(3) Argumentative inference, where a formula is a consequence of the base iff it is entailed by at least one maximal consistent subset and its negation is not.

Maximal Consistent Subsets – Formally

Denote the set of all maximal consistent subsets of Γ with $MC(\Gamma)$.

Then:

Skeptical inference:

$$\Gamma \vdash_{skep} \varphi \text{ iff } \Sigma \vdash \varphi \text{ for all } \Sigma \in MC(\Gamma)$$

Credulous inference:

$$\Gamma \vdash_{cred} \varphi \text{ iff } \Sigma \vdash \varphi \text{ for at least one } \Sigma \in MC(\Gamma)$$

Argumentative inference:

$$\Gamma \vdash_{arg} \varphi \text{ iff } \Sigma \vdash \varphi \text{ for at least one } \Sigma \in MC(\Gamma) \text{ and}$$
$$\text{there is no } \Sigma \in MC(\Gamma) \text{ such that } \Sigma \vdash \neg\varphi$$

Maximal Consistent Subsets – Example

Suppose that we have the following (contradictory) pieces of information:

Example (Inconsistent Database K)

We have six pieces of information ($\varphi_1, \ldots, \varphi_6$):

$\varphi_1 = a \land b$

$\varphi_2 = a \land \neg d$

$\varphi_3 = a \land \neg c \land e$

$\varphi_4 = a \land (\neg c \to \neg e)$

$\varphi_5 = a \land c$

$\varphi_6 = a \land \neg a$

We can derive the following maximal consistent subsets:

(K1) = $\{\varphi_1, \varphi_2, \varphi_3\}$,
(K2) = $\{\varphi_1, \varphi_2, \varphi_4, \varphi_5\}$.

Maximal Consistent Subsets – Example

Given the maximal consistent subsets

(K1) $= \{\varphi_1, \varphi_2, \varphi_3\}$,
(K2) $= \{\varphi_1, \varphi_2, \varphi_4, \varphi_5\}$

we can make the following inferences:

Using skeptical inference: from all these maximal consistent subsets, φ_1 and φ_2 can be derived as a conclusion.

Using credulous inference: from all these maximal consistent subsets, $\varphi_1, \ldots, \varphi_5$ can be derived as a conclusion.

Using argumentative inference: from all these maximal consistent subsets, φ_1 and φ_2 can be derived as a conclusion;

Maximal Consistent Subsets – Example

Given the maximal consistent subsets

(K1) $= \{\varphi_1, \varphi_2, \varphi_3\}$,
(K2) $= \{\varphi_1, \varphi_2, \varphi_4, \varphi_5\}$

we can make the following inferences:

Using skeptical inference: from all these maximal consistent subsets, φ_1 and φ_2 can be derived as a conclusion.

Using credulous inference: from all these maximal consistent subsets, $\varphi_1, \ldots, \varphi_5$ can be derived as a conclusion.

Using argumentative inference: from all these maximal consistent subsets, φ_1 and φ_2 can be derived as a conclusion;
so can $\psi = a \land b \land \neg d \land e$, since $\{\varphi_1, \varphi_2, \varphi_3\} \vdash \psi$, while both $\{\varphi_1, \varphi_2, \varphi_3, \psi\}$ and $\{\varphi_1, \varphi_2, \varphi_4, \varphi_5, \psi\}$ are consistent.

Maximal Consistent Subsets – Example

Given the maximal consistent subsets

(K1) $= \{\varphi_1, \varphi_2, \varphi_3\}$,
(K2) $= \{\varphi_1, \varphi_2, \varphi_4, \varphi_5\}$

we can make the following inferences:

Using skeptical inference: from all these maximal consistent subsets, φ_1 and φ_2 can be derived as a conclusion.

Using credulous inference: from all these maximal consistent subsets, $\varphi_1, \ldots, \varphi_5$ can be derived as a conclusion.

Using argumentative inference: from all these maximal consistent subsets, φ_1 and φ_2 can be derived as a conclusion;
so can $\psi = a \land b \land \neg d \land e$, since $\{\varphi_1, \varphi_2, \varphi_3\} \vdash \psi$, while both $\{\varphi_1, \varphi_2, \varphi_3, \psi\}$ and $\{\varphi_1, \varphi_2, \varphi_4, \varphi_5, \psi\}$ are consistent.
On the other hand, $\{\varphi_1, \varphi_2, \varphi_4, \varphi_5, \neg\psi\}$ is also consistent;
thus ψ is not a skeptical consequence.

Paraconsistent Logics

There is a range of so-called paraconsistent logics for reasoning with inconsistency. Selecting an appropriate paraconsistent logic for an application depends on the requirements of the application.

Types of paraconsistent logics that have proved to be of use for knowledge representation and reasoning:

(1) Weakly-negative logics

(2) Four-valued logics

(3) Signed systems

(4) Quasi-classical logic

For more information on systems (1), (3), and (4) see:
Anthony Hunter: Paraconsistent Logics, 1998

Four-Valued Logic

Idea: Use a subset of the classical language and a subset of the classical proof theory, together with an intuitive four-valued semantics.

Language

The language for four-valued logic is a subset of classical logic. Let P be the usual set of formulae of classical logic that is formed using the connectives $\neg$, $\wedge$, and $\vee$. Then the set of formulae of the language, denoted Q, is $P \cup \{\alpha \rightarrow \beta \mid \alpha, \beta \in P\}$.

$\rightsquigarrow$ Implication cannot be nested: $A \rightarrow B$ vs. $\neg A \vee B$.

Truth Values

A formula in the language can be one of "true", "false", "both" or "neither", which we denote by the symbols **T**, **F**, **B**, and **N**, respectively.

Four-Valued Logic – Example

Imagine an artificial reasoner (AR) to operate as follows:

(1) The AR receives assertions and denials of atomic sentences.
Upon receiving an assertion, it is to mark the item "told true".
Upon receiving a denial, it is to mark the item "told false".

(2) The AR has exactly four possibilities for any particular atomic sentence:
Told true but never told false, i.e., {"told true"}
Told false but never told true, i.e., {"told false"};
Never told true and never told false, i.e., {};
Told true and told false, i.e., {"told true", "told false"}.

(3) The four possibilities correspond to our four values: **T**, **F**, **N**, **B**.

Four-Valued Logic – Example

(4) The AR receives questions, which it has to answer. Facing a question of the form "p?", it answers as follows:

If the item has the value T then answer "yes";
If the item has the value F then answer "no";
If the item has the value N then answer "don't know"
If the item has the value B then answer "yes and no".

(5) If it receives additional information, it incorporates it into the information it already has. It does this by representing what it has been told by a setup.

Four-Valued Logic – Example

Setup

A setup maps atomic sentences into the set $\mathbf{4} = \{\mathbf{T}, \mathbf{F}, \mathbf{N}, \mathbf{B}\}$, according to the following rules:

If an atomic formula p is affirmed, the setup is revised as follows:

If the current value of p is **N**, it is mapped to **T**.
If the current value of p is **B**, it is mapped to **B**.
If the current value of p is **T**, it is mapped to **T**.
If the current value of p is **F**, it is mapped to **B**.

If an atomic formula p is denied, the setup is revised as follows:

What do you think? Left as an exercise.

Four-Valued Logic – Example

Thus we envision that the database we are working with can be represented by a list of statements.

Consider the 2000 presidential election in the US and a hypothetical setup for the states *(NC), (FL), (CA)* as well as the candidates *Bush*, *Gore*, and *Nader*.

Statement	Truth Value Assigned
Bush won NC	{told true} = **T**
Gore won NC	{told false} = **F**
Nader won NC	{told false} = **F**
Bush won FL	{told true, told false} = **Both**
Gore won FL	{told true, told false} = **Both**
Nader won FL	{told false} = **F**
Bush won CA	{} = **None**
Gore won CA	{} = **None**
Nader won CA	{} = **None**

Four-Valued Logic – Natural Orderings

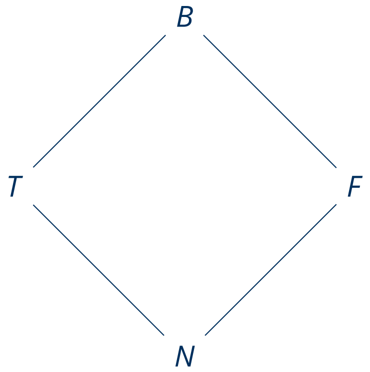

Information Lattice

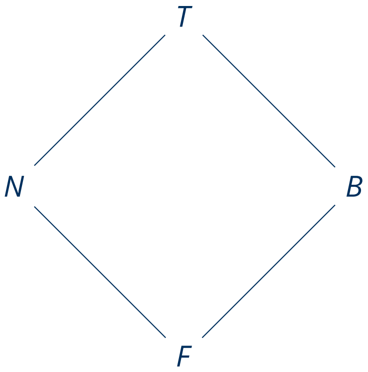

Logical Lattice
(Truth Ordering)

Four-Valued Logic – Semantics

Basic Assumption: For the semantics we use the truth ordering.

Semantic Assignment Function, Properties

The semantic assignment function observes monotonicity and complementation in the logical lattice. So $\neg x$ is the complement of x, $x \wedge y$ is the meet of $\{x, y\}$ and $x \vee y$ is the join of $\{x, y\}$, giving the following truth tables for the $\neg, \wedge, \vee$ connectives:

Four-Valued Logic – Semantics

Table: Truth table for conjunction

$\wedge$	N	F	T	B
N	N	F	N	F
F	F	F	F	F
T	N	F	T	B
B	F	F	B	B

Table: Truth table for disjunction

$\vee$	N	F	T	B
N	N	N	T	T
F	N	F	T	B
T	T	T	T	T
B	T	B	T	B

Table: Truth table for negation

α	N	F	T	B
$\neg\alpha$	N	T	F	B

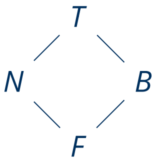

Four-Valued Logic – Semantics

Intuitive Idea: An inference in our four-valued logic never takes one farther away from the truth:

Inference never obtains a conclusion that is "less true" than the premises.

Inference

Let α and β be formulas of four-valued logic over atoms A.

The inference $\alpha \rightarrow \beta$ is **valid** iff $v(\alpha) \leq_t v(\beta)$ for all four-valued truth value assignments $v : A \rightarrow \mathbf{4}$, where $\leq_t$ is the truth ordering.

Four-Valued Logic – Semantics

Intuitive Idea: An inference in our four-valued logic never takes one farther away from the truth:

Inference never obtains a conclusion that is "less true" than the premises.

Inference

Let α and β be formulas of four-valued logic over atoms A.

The inference $\alpha \rightarrow \beta$ is **valid** iff $v(\alpha) \leq_t v(\beta)$ for all four-valued truth value assignments $v : A \rightarrow \mathbf{4}$, where $\leq_t$ is the truth ordering.

Example

Consider the classical tautology $(A \wedge \neg A) \rightarrow B$. Suppose $v(A) = \mathbf{B}$ and $v(B) = \mathbf{F}$.
Then $v(A \wedge \neg A) = \mathbf{B} \wedge \neg\mathbf{B} = \mathbf{B}$.
Since $\mathbf{B} \not\leq_t \mathbf{F}$ the inference $(A \wedge \neg A) \rightarrow B$ is not valid.

Belief Revision

Belief Revision

Consider again the following situation:

a : All European swans are white.
b : The bird caught in the trap is a swan.
c : The bird caught in the trap comes from Sweden.
d : Sweden is part of Europe.

From a – d, we infer e: The bird caught in the trap is white.

Now suppose that: The bird caught in the trap turns out to be black.

Problem: Upon adding $\neg e$, the database becomes inconsistent.

This is a typical belief revision scenario: A rational agent receives new information that makes it change its beliefs.

Methodological Problems and Kinds of Belief Change

Methodological Problems:

(1) How to represent the beliefs in the database?

(2) What is the relation between the elements explicitly represented in the database and the beliefs that may be derived from these elements?

(3) How to choose which beliefs to retract?

Kinds of Belief Change

(1) **Expansion**: A new sentence is added to a belief system $\mathcal{B}$ together with the logical consequences of the addition (regardless of whether the larger set so formed is consistent).

(2) **Revision**: A new sentence that is inconsistent with a belief system $\mathcal{B}$ is added, but, in order to maintain consistency in the resulting belief system, some of the old sentences in $\mathcal{B}$ are deleted.

(3) **Contraction**: Some sentence in $\mathcal{B}$ is retracted without adding any new facts. In order for the resulting system to be closed under logical consequences some other sentences from $\mathcal{B}$ must be given up.

Formal Preliminaries

For our framework, we are working with a language $\mathcal{L}$ based on first-order logic.

$\mathcal{L}$ is closed under application of boolean operators:
$\neg$ (negation), $\wedge$ (conjunction), $\vee$ (disjunction), $\rightarrow$ (implication)

$\mathcal{L}$ is identified by its consequence relation $\vdash$:

(1) $\vdash \varphi$ for all truth-functional tautologies.

(2) If $\vdash (\varphi \rightarrow \psi)$ and $\vdash \varphi$, then $\vdash \psi$.

(3) $\vdash$ is consistent, i.e. $\nvdash \bot$.

(4) $\vdash$ satisfies the deduction theorem:
$\{\varphi_1, \varphi_2, \ldots, \varphi_n\} \vdash \psi$ iff $\vdash (\varphi_1 \wedge \varphi_2 \wedge \ldots \wedge \varphi_n) \rightarrow \psi$.

(5) $\vdash$ is compact.

The AGM Paradigm

In the AGM paradigm, beliefs are represented as sentences of $\mathcal{L}$ and belief sets as theories of $\mathcal{L}$ that are closed under logical consequence.

The process of belief revision is modeled as a function $*$ mapping a belief set $\mathcal{B}$ and a sentence φ to a new belief set $\mathcal{B}*\varphi$.

We want $*$ to capture the notion of rational belief revision correctly. Hence, we need to impose certain constraints. These constraints (next slide) are based on the principle of minimal change:

Principle of minimal change:

A rational agent ought to change its beliefs as little as possible in order to (consistently) accommodate the new information.

The AGM Postulates for Belief Revision

($\mathcal{B} * 1$) $\varphi \in \mathcal{B} * \varphi$.

($\mathcal{B} * 2$) $\mathcal{B} * \varphi \subseteq \mathcal{B} + \varphi$.

($\mathcal{B} * 3$) If $\neg\varphi \notin \mathcal{B}$ then $\mathcal{B} + \varphi \subseteq \mathcal{B} * \varphi$.

($\mathcal{B} * 4$) If φ is consistent then $\mathcal{B} * \varphi$ is also consistent.

($\mathcal{B} * 5$) If $\vdash \varphi \leftrightarrow \psi$ then $\mathcal{B} * \varphi = \mathcal{B} * \psi$.

($\mathcal{B} * 6$) $\mathcal{B} * (\varphi \wedge \psi) \subseteq (\mathcal{B} * \varphi) + \psi$.

($\mathcal{B} * 7$) If $\neg\psi \notin \mathcal{B} * \varphi$ then $(\mathcal{B} * \varphi) + \psi \subseteq \mathcal{B} * (\varphi \wedge \psi)$.

The AGM Postulates – Explanation

($\mathcal{B} * 1$): Says that the new information on φ should always be included in the new belief set.

($\mathcal{B} * 2$) and ($\mathcal{B} * 3$): Together, they state that whenever the new information φ does not contradict the initial belief set $\mathcal{B}$, there is no reason to remove any of the original beliefs at all; the new belief state will contain $\mathcal{B}$, the new information, and what follows from the logical closure.

($\mathcal{B} * 4$): Says that the agent should aim for consistency.

($\mathcal{B} * 5$): Says that the syntax of the new information has no effect on the revision process.

($\mathcal{B} * 6$) and ($\mathcal{B} * 7$): They say that for any two sentences φ and ψ, if in revising the initial belief set $\mathcal{B}$ by φ one can reach a belief set $\mathcal{B} * \varphi$ that is consistent with ψ, then to produce $\mathcal{B}(\varphi \wedge \psi)$ all that one needs to do is to expand $\mathcal{B} * \varphi$ with ψ.

Belief Contraction

Contraction can be described as the process of rationally removing a certain belief φ from a belief set $\mathcal{B}$.

It is also formally defined as a function $\dot{-}$ mapping a belief set $\mathcal{B}$ and a sentence φ to a new belief set $\mathcal{B} \dot{-} \varphi$.

Contraction typically occurs when an agent loses faith in φ and decides to give it up.

Problem: It does not suffice to only take out φ from $\mathcal{B}$. Other sentences in $\mathcal{B}$ could reproduce φ trough logical closure.

Example

Consider the belief set $\mathcal{B} = Cn(\{p \rightarrow q, p, q\})$ and assume that we want to contract $\mathcal{B}$ by q. Then, not only do we have to remove q from $\mathcal{B}$, but we also need to give up (at least) one of $p \rightarrow q$ or p, for otherwise q will resurface.

Contraction Postulates

($\mathcal{B}\dot{-}1$) $\mathcal{B}\dot{-}\varphi \subseteq \mathcal{B}$.

($\mathcal{B}\dot{-}2$) If $\varphi \notin \mathcal{B}$ then $\mathcal{B}\dot{-}\varphi = \mathcal{B}$.

($\mathcal{B}\dot{-}3$) If $\nvdash \varphi$ then $\varphi \notin \mathcal{B}\dot{-}\varphi$.

($\mathcal{B}\dot{-}4$) If $\varphi \in \mathcal{B}$, then $\mathcal{B} \subseteq (\mathcal{B}\dot{-}\varphi) + \varphi$.

($\mathcal{B}\dot{-}5$) If $\vdash \varphi \leftrightarrow \psi$ then $\mathcal{B}\dot{-}\varphi = \mathcal{B}\dot{-}\psi$.

($\mathcal{B}\dot{-}6$) $(\mathcal{B}\dot{-}\varphi) \cap (\mathcal{B}\dot{-}\psi) \subseteq \mathcal{B}\dot{-}(\varphi \wedge \psi)$.

($\mathcal{B}\dot{-}7$) If $\psi \notin \mathcal{B}\dot{-}(\varphi \wedge \psi)$ then $\mathcal{B}\dot{-}(\varphi \wedge \psi) \subseteq \mathcal{B}\dot{-}\varphi$.

Contraction Postulates – Explanation

($\mathcal{B} \dot{-} 1$): Says that contraction produces a belief set smaller than or equal to the original one.

($\mathcal{B} \dot{-} 2$): Says that if φ is not in the initial belief set $\mathcal{B}$ to start with, then there is no reason to change anything at all.

($\mathcal{B} \dot{-} 3$): Tells us that all sentences but tautologies can be removed from the initial beliefs $\mathcal{B}$.

($\mathcal{B} \dot{-} 4$): Says that contracting and then expanding by φ will give us back (at least) the initial theory $\mathcal{B}$.

Contraction Postulates – Explanation

($\mathcal{B} \dot{-} 5$): Tells us that contraction by logically equivalent sentences produces the same result. That is, contraction is not syntax-sensitive.

($\mathcal{B} \dot{-} 6$): Tells us that if there is some belief $\chi \in \mathcal{B}$ that is neither related to φ nor ψ, it should not be affected by the contraction of $\mathcal{B}$ by $\varphi \wedge \psi$.

($\mathcal{B} \dot{-} 7$): Since $\mathcal{B} \dot{-} \varphi$ is the minimal change of $\mathcal{B}$ to remove φ, it follows that $\mathcal{B} \dot{-} (\varphi \wedge \psi)$ cannot be larger than $\mathcal{B} \dot{-} \varphi$. This postulate makes it smaller or equal to it.

Relation between revision and contraction (1)

Idea: Define revision in terms of contraction as follows:

To revise $\mathcal{B}$ by φ, first contract $\mathcal{B}$ by $\neg\varphi$ (thus removing anything that may contradict the new information) and then expand the resulting theory with φ.

This is known as the *Levi Identity*:

Levi Identity

$$\mathcal{B} * \varphi = (\mathcal{B} \dot{-} \neg\varphi) + \varphi$$

Relation between revision and contraction (2)

Now: Define contraction in terms of revision. The idea is that:

A sentence ψ is accepted in the contraction $\mathcal{B} \dot{-} \varphi$ iff ψ is accepted both in $\mathcal{B}$ and in $\mathcal{B} * \neg\varphi$.

This is known as the *Harper Identity*:

Harper Identity

$$\mathcal{B} \dot{-} \varphi = (\mathcal{B} * \neg\varphi) \cap \mathcal{B}$$

Epistemic Entrechment

Motivation: Agents can perceive their individual beliefs to have different epistemic values. Consider, for instance:

ψ All swans are white. (law-like belief)

χ Lucy is a swan. (simple belief)

Idea: Introduce the notion of *epistemic entrenchment*.
Intuitively, this notion can be understood as follows:

Epistemic Entrechment:
The epistemic entrenchment of a belief ψ is the degree of resistance that ψ exhibits to change: the more entrenched ψ is, the less likely it is to be removed during contraction by somer other belief φ.

Epistemic Entrenchment – Formally

Epistemic entrenchment is defined as a preorder $\leqslant$ on L encoding the "retractibility" of individual beliefs.

Preorder $\leqslant$:

$\chi \leqslant \psi$ iff the agent is at least as (or more) reluctant to give up ψ than it is to give up χ.

Constraints on $\leqslant$ in relation to a given $\mathcal{B}$:

(EE1) If $\varphi \leqslant \psi$ and $\psi \leqslant \chi$ then $\varphi \leqslant \chi$.

(EE2) If $\varphi \vdash \psi$ then $\varphi \leqslant \psi$.

(EE3) $\varphi \leqslant \varphi \wedge \psi$ or $\psi \leqslant \varphi \wedge \psi$.

(EE4) When $\mathcal{B}$ is consistent, $\varphi \notin \mathcal{B}$ iff $\varphi \leqslant \psi$ for all $\psi \in L$.

(EE5) If $\psi \leqslant \varphi$ for all $\psi \in L$, then $\vdash \varphi$.

Epistemic Entrenchment Postulates – Explanation

(EE1) States that $\leqslant$ is transitive.

(EE2) Says that the stronger a belief is logically, the less entrenched it is.

(EE3) Says that if one wants to retract $\varphi \wedge \psi$ from $\mathcal{B}$, this can only be achieved by giving up either φ or ψ.

(EE4) Says that in the principle case where $\mathcal{B}$ is consistent, all non-beliefs (all sentences that are not in $\mathcal{B}$) are minimally entrenched.

(EE5) Says that tautologies are the only maximal elements of $\leqslant$ and therefore the hardest to remove.

Epistemic Entrenchment and Contraction

In a next step, we define contraction in terms of epistemic entrenchment by defining the condition (C-):

(C-)

$\psi \in \mathcal{B} \dot{-} \varphi$ iff $\psi \in \mathcal{B}$ and either $\varphi < \varphi \vee \psi$ or $\vdash \varphi$.

With that definition, we get the following *representation theorem*:

Theorem

Let $\mathcal{B}$ be a theory of L. If $\leqslant$ is a preorder in L that satisfies the axioms (EE1)–(EE5) then the function defined by (C-) is an AGM contraction function. Conversely, if $\dot{-}$ is an AGM contraction function, then there is a preorder $\leqslant$ in L that satisfies the axioms (EE1)–(EE5) as well as condition (C-).

Conclusion

- A knowledge base K is **inconsistent** if both $K \vdash \varphi$ and $K \vdash \neg\varphi$ for some φ.
- To avoid trivial conclusions from inconsistent KBs, various techniques can be applied:
- Reasoning (credulously/skeptically/argumentatively) from **maximally consistent subsets**.
- Using **four-valued logic** to represent inconsistency directly in the truth values.
- Modify the knowledge base (maintaining logical closure) using techniques of belief change:
  - **Expansion**: Add a sentence
  - **Revision**: Add a sentence and achieve consistency
  - **Contraction**: Retract a sentence