

Section 7
Belief Revision

Subsection 7.1
Introduction

Introduction

Nonmonotonic reasoning:

make plausible conjectures in the absence of complete information.

Intelligent systems:

- adept in changing the knowledge base when new knowledge arrives.
- resolving potential conflicts of old and new knowledge

Belief revision:

- is a process that can be used to **modify a knowledge base** when new information is acquired
- needs to **cope with inconsistencies** due to the addition of new knowledge (Recall: anything follows logically from an inconsistent knowledge base; rendering it useless for reasoning.)

Goal of this chapter:

lay down the principles of foundations of principled mechanisms for **modifying a knowledge base** in a rational and coherent way.

Our abstract setting

The setting considered in this chapter:

- agent (or information system)
- knowledge base (KB)
 - represents the collection of facts (/information) the agent believes in
 - represented by a theory T
- some language \mathcal{L} that is
 - a set of closed formulas (sentences) over a fixed signature.
 - used for content of KB and the fact to be added
- one change φ
one fact at a time is added or retracted
- modifications of the KB are addition or removal of facts. It is not the transformation of formulas (e.g.: “all birds, but Tweety fly.”)

Considerations

Belief Revision is a mechanism for modeling rational decisions concerning modifications of the KB. When incorporating new knowledge that is inconsistent with the knowledge base, the agent must decide which knowledge to give up.

A rational agent would modify the KB as little as possible to incorporate new information. Usually there are several options.

Example on retraction choices

Example 7.1 (Tweety again)

Consider the beliefs: “all birds fly”

“Tweety is a bird”

Consequence: “Tweety can fly”

New fact: “Tweety cannot fly”

How is the KB to be modified?

Retract “Tweety can fly”, but also all beliefs that are in conflict with the new fact.

There are several options:

1. retract “Tweety is a bird”
2. retract “all birds fly”
3. retract both

Which option is chosen depends on the preferences of the agent!

Agent might be less confident about one fact and prefer to retract that.

The choice for the ‘right’ option depends on aspects outside of logic.

We employ a preference ordering of the facts in the KB to realize a choice function.

Principles for modifications

Aim: study rational modifications to KBs guided by the “Principle of Minimal Change”

Notions of rationality and minimality are not so easy to capture formally!

Principle of rationality:

only facts are retracted that cause the inconsistencies.

Principle of minimal change:

as much as possible of the underlying information should be kept (in accordance of the preference relation).

Preference relation captures

- the information content of the KB
- the agent’s commitment to this information
- how the information should behave under change

Considering Example 7.1 again

Minimality?

Reconsider Example 7.1:
Are choices 1. and 2. more minimal than 3.?
They give up a smaller number of facts.

But: Interdependencies among beliefs might necessitate to give up more than the minimal number of beliefs.

How to measure the [magnitude of change](#)?
Is it better to discard one strongly held belief or several weakly held beliefs?

The AGM framework

The basis for most belief revision approaches (or systems) is the AGM framework

The AGM framework:

- developed by Carlos Alchourrón, Peter Gärdenfors and David Makinson in the early eighties.
- is a formal framework for modeling ideal and rational changes of KBs under the Principle of Minimal Change
- provides coherent retraction and incorporation of information
- for our Tweety example it admits all three choices for retraction

Two approaches to belief revision

The approaches to belief revision we consider:

1. **Axiomatic**

Describe change functions axiomatically⁷ by use of postulates

2. **Constructive**

By use of the information encoded in a preference relation, called epistemic entrenchment relation, one can construct the change functions.

We discuss the **change operations** for KBs:
expansion, contraction, withdrawal, and revision.

⁷ Similarly to our way of showing properties of inference relations.

Subsection 7.2

Axiomatizing change functions

Change function: expansion

Expansion⁸ of a knowledge base:

- accepting new information
- not removing any of the previously accepted information
- can lead to inconsistencies in the KB

Definition 7.2 (Expansion)

Let T be a theory and $\varphi \in \mathcal{L}$. The **expansion** of a theory T w.r.t. a sentence φ is the logical closure of $T \cup \varphi$.

The set of all theories of a language \mathcal{L} is $\mathcal{K}_{\mathcal{L}}$.

An **expansion function** $^+$ is a function $^+ : \mathcal{K}_{\mathcal{L}} \times \mathcal{L} \mapsto \mathcal{K}_{\mathcal{L}}$ mapping (T, φ) to T_{φ}^+ , where $T_{\varphi}^+ = Th(T \cup \varphi)$.

⁸Not to be confused with the expansions of autoepistemic theories.

Expansion operation

Note, that the expansion operation ...

- is a **monotonic operation**,
i.e. $T \subseteq T_{\varphi}^{+}$ and if $\neg\varphi \in T$, then T_{φ}^{+} is inconsistent.
- of T by φ , if $\neg\varphi \notin T$, has T_{φ}^{+} as the smallest change to incorporate φ in T —realizing the **Principle of Minimal Change**.
- is a **unique operation!** (No choice here.)

Change function: contraction

A contraction of a KB w.r.t. a formula φ causes the removal of a set of formulas from the KB, so that φ is no longer implied (unless φ is a tautology).

The challenge is to determine which sentences should be given up.

Definition 7.3 (Contraction)

Let T be a theory and $\varphi \in \mathcal{L}$.

A contraction function $\bar{\cdot}$ is a function $\bar{\cdot} : \mathcal{K}_{\mathcal{L}} \times \mathcal{L} \mapsto \mathcal{K}_{\mathcal{L}}$ mapping (T, φ) to T_{φ}^{-} , where T_{φ}^{-} satisfies the following postulates (for any $\varphi, \psi \in \mathcal{L}$ and any $T \in \mathcal{K}_{\mathcal{L}}$):

- (1) $T_{\varphi}^{-} \in \mathcal{K}_{\mathcal{L}}$
- (2) $T_{\varphi}^{-} \subseteq T$
- (3) If $\varphi \notin T$ then $T \subseteq T_{\varphi}^{-}$
- (4) If $\not\vdash \varphi$ then $\varphi \notin T_{\varphi}^{-}$
- (5) $T \subseteq (T_{\varphi}^{-})_{\varphi}^{+}$ (recovery)
- (6) If $\vdash (\varphi \leftrightarrow \psi)$ then $T_{\varphi}^{-} = T_{\psi}^{-}$
- (7) $T_{\varphi}^{-} \cap T_{\psi}^{-} \subseteq T_{(\varphi \wedge \psi)}^{-}$
- (8) If $\varphi \notin T_{(\varphi \wedge \psi)}^{-}$ then $T_{(\varphi \wedge \psi)}^{-} \subseteq T_{\varphi}^{-}$

On the contraction postulates

The AGM postulates for contraction

- incorporate the 'Principle of minimal change'
- act as integrity constraints for change functions
- do **not** determine a **unique change function**, but rather identify the set of possible KBs resulting from retracting information
- identify a **class of functions for a KB**
- are motivated by the criterion of informational economy

Intuition for the individual contraction postulates I

First postulate ($\bar{1}$):

The result of contraction is a theory, new KB is logically closed.

Second postulate ($\bar{2}$)

Contracting a theory involves only removal of old information, never the addition of new information. No spurious information is added.

Third postulate ($\bar{3}$)

When information φ is not accepted then taking it away has no effect.

($\bar{2}$) and ($\bar{3}$) say that “if φ is not in T , then taking it away has no effect”

Fourth postulate ($\bar{4}$)

If φ is not a tautology, then it must be removed in the contraction T_{φ}^{-} .

Intuition for the individual contraction postulates II

Fifth postulate ($\bar{5}$), a.k.a recovery

If φ is retracted and then replace it by using expansion, then T is restored.

Together with the last 4 postulates says that if $\varphi \in T$, then $T = (T)_{\varphi}^{+}$ i.e., no more information is lost than can be reintroduced by an expansion w.r.t. the information retracted.

Sixth postulate ($\bar{6}$)

The same result is obtained if contraction is performed w.r.t. equivalent sentences; contraction is oblivious to syntax.

Seventh postulate ($\bar{7}$)

The theory obtained from retracting w.r.t. $\varphi \wedge \psi$ should never be smaller than taking the intersection of T_{φ}^{-} and T_{ψ}^{-} .

Eighth postulate ($\bar{8}$)

If φ is not contained in the contraction w.r.t. $(\varphi \wedge \psi)$ then the contraction w.r.t. this conjunction cannot be larger than the theory obtained from contraction of φ alone.

The recovery postulate (̄5) and withdrawal operation

The recovery postulate is controversial!

Two strange effects of recovery:

1. A consequence of theory closure and recovery is that for all $\psi \in \mathcal{L}$ holds: if a sentence φ is contained in theory T , then $(\psi \longrightarrow \varphi) \in T_{\psi \vee \varphi}^-$ and the consequences $\varphi \in (T_{\psi \vee \varphi}^-)_{\psi}^+$ and $\neg\psi \in (T_{\psi \vee \varphi}^-)_{\neg\psi}^+$ are not always desirable.
2. Assume φ and ψ are in T , then so is $(\varphi \vee \psi)$. If $(\varphi \vee \psi)$ is learned to be wrong and retracted by the agent, but later on re-learned, then it should be possible not to add φ and ψ again.

Change function: withdrawal

The withdrawal operation is similar to contraction, but need not satisfy recovery. A [withdrawal function](#) satisfies the postulates (̄1) to (̄4) and (̄6) to (̄8), but not necessarily (̄5).

Change function: revision

Revision attempts to change a KB as little as possible to accommodate new information. If the KB turns inconsistent, then information has to be retracted to regain consistency. Thus revision functions are nonmonotonic.

Definition 7.4 (Revision)

Let T be a theory and $\varphi \in \mathcal{L}$.

A **revision function** $*$ is a function $*$: $\mathcal{K}_{\mathcal{L}} \times \mathcal{L} \mapsto \mathcal{K}_{\mathcal{L}}$ mapping (T, φ) to T_{φ}^* , where T_{φ}^* satisfies the following postulates (for any $\varphi, \psi \in \mathcal{L}$ and any $T \in \mathcal{K}_{\mathcal{L}}$):

- (*1) $T_{\varphi}^* \in \mathcal{K}_{\mathcal{L}}$
- (*2) $\varphi \in T_{\varphi}^*$
- (*3) $T_{\varphi}^* \subseteq T_{\varphi}^+$
- (*4) If $\neg\varphi \notin T$ then $T_{\varphi}^+ \subseteq T_{\varphi}^*$
- (*5) If $T_{\varphi}^* = \perp$ then $\vdash \neg\varphi$
- (*6) If $\vdash (\varphi \leftrightarrow \psi)$ then $T_{\varphi}^* = T_{\psi}^*$
- (*7) $T_{(\varphi \wedge \psi)}^* \subseteq (T_{\varphi}^*)_{\psi}^+$
- (*8) If $\neg\psi \notin T_{\varphi}^*$ then $(T_{\varphi}^*)_{\psi}^+ \subseteq T_{\varphi \wedge \psi}^*$

Intuition for the revision postulates I

First postulate (*1):

Revising a theory results in a theory, new KB is logically closed.

Second postulate (*2)

Information to be added by a revision is always successfully added to the KB.

Third postulate (*3)

Revising a theory can never incorporate more information than an expansion operation.

Fourth postulate (*4)

If $\neg\varphi$ is not contained in T then the expansion of T w.r.t. φ is contained in the revision w.r.t. φ .

Intuition for the revision postulates II

Fifth postulate (*5)

The only way to obtain an inconsistent theory is to revise w.r.t. an inconsistent sentence

Sixth postulate (*6)

Revision functions are syntax independent.

Seventh postulate (*7)

The theory that results from revising w.r.t. $\varphi \wedge \psi$ should never contain more than the revision w.r.t. φ followed by the expansion w.r.t. ψ .

Eighth postulate (*8)

If $\neg\varphi$ is contained in the revision w.r.t. φ then the revision w.r.t. φ followed by the expansion w.r.t. ψ should not contain more information than the theory that results from revising w.r.t. $\varphi \wedge \psi$

If the first 6 postulates hold. Then (*7) and (*8) imply that

$T_{\varphi \wedge \psi}^*$ is equivalent to T_{φ}^* , T_{ψ}^* or $T_{\varphi}^* \cap T_{\psi}^*$

Subsection 7.3

Interrelations between change functions

The change functions are interdefinable

Relating contraction and expansion to revision

Theorem 7.5

If $-$ is a contraction function and $+$ a expansion function, then $$ defined by the *Levi Identity* below defines a revision function.*

$$T_{\varphi}^* = (T_{\neg\varphi}^-)^+$$

Proof: blackboard.

Intuitively, the Levi Identity says:

A revision operator can be defined in terms of contraction as follows:
to revise T by φ ,

1. contract T by $\neg\varphi$
(removing anything that may contradict the new information)
2. expand the resulting theory with φ .

Relating expansion and expansion to contraction

There is also a similar way of defining contraction in terms of revision:

Theorem 7.6

If $$ is a revision function, then $\bar{\cdot}$ defined by the Harper Identity below defines a contraction function.*

$$T_{\varphi}^{-} = T \cap T_{\neg\varphi}^{*}$$

Proof: blackboard.

Intuitively, the Harper Identity says:

A contraction operator can be defined in terms of revision as follows:
to contract T by φ ,

1. revise T w.r.t. $\neg\varphi$
(remove everything from T that is used to derive φ .)
2. restrict to formulas in T
(do not add new knowledge derived from $\neg\varphi$.)

On the Levi and the Harper identity

The **Levi Identity** of producing revision functions covers the whole range of AGM revision functions, i.e. for every AGM revision function $*$ there is a contraction function $\bar{\cdot}$ that produces $*$ by means of the Levi Identity.

Likewise, the **Harper Identity** is a sound and complete method for constructing contraction functions, i.e. the functions generated from an AGM revision function by means of the Harper Identity satisfies (7.1) to (7.8).

By combining both identities, one makes the full circle:

$$\begin{aligned} T &= T \cap ((T_{\neg\varphi}^-)^+_{\varphi}) \\ &= (T \cap T_{\neg\varphi}^*)^+_{\varphi} \end{aligned}$$

Subsection 7.4

Selection functions

Selection functions by Epistemic Entrenchment Orderings

Considerations

Goal: · constructive models for revision and contraction
(from the axiomatic specification)

Due to Levi and Harper Identity it suffices to devise one constructive model for revision or for contraction.

Since the AGM postulates describe classes of revision or of contraction functions, it takes a [preference relation](#) to single out one. The choice depends on

- belief state
- assumptions on epistemic dependence
- assumptions on causal dependence
- independence of the reasons for adopting beliefs

Epistemic entrenchment orderings (by Gärdenfors)

Epistemic entrenchment orderings are based on a relative ranking of information according to its importance.

This function can uniquely determine a change function by providing a selection criterion that can be used to identify those sentences to be retracted (or retained or acquired) during change.

Intuition:

when a choice needs to be taken, those sentences having the lowest degree of epistemic entrenchment are discarded.

Defining epistemic entrenchment orderings

Definition 7.7

Let T be a theory of \mathcal{L} . An **epistemic entrenchment** related to T is any binary relation \leq on \mathcal{L} satisfying the conditions:

- (EE1) If $\varphi \leq \psi$ and $\psi \leq \chi$, then $\varphi \leq \chi$.
- (EE2) For all $\varphi, \psi \in \mathcal{L}$, if $\varphi \vdash \psi$ then $\varphi \leq \psi$.
- (EE3) For all $\varphi, \psi \in \mathcal{L}$, $\varphi \leq (\varphi \wedge \psi)$ or $\psi \leq (\varphi \wedge \psi)$.
- (EE4) When $T \neq \perp$, $\varphi \notin T$ iff $\varphi \leq \psi$ for all $\psi \in \mathcal{L}$.
- (EE5) If $\psi \leq \varphi$ for all $\psi \in \mathcal{L}$, then $\vdash \varphi$.

If $\varphi \leq \psi$, then ψ is **at least entrenched** as φ . Define $\varphi < \psi$ as $\varphi \leq \psi$ and $\psi \not\leq \varphi$. If $\varphi \leq \psi$ and $\psi \leq \varphi$, then φ and ψ are **equally entrenched**.

Intuition of the EE conditions

First condition (EE1)

requires that an epistemic entrenchment ordering be transitive.

Second condition (EE2)

says that if φ is logically stronger than ψ , then ψ is at least as entrenched as φ .

E.g.: $\varphi \models (\varphi \wedge \psi)$ and (EE2) says that $(\varphi \wedge \psi)$ is at least as entrenched as φ .

Third condition (EE3)

Condition (EE3) together with (EE1) and (EE2) imply that a conjunction is ranked at the same level as its least ranked conjunct.

Fourth condition (EE4)

sentences not in the theory T are minimal.

Fifth condition (EE5)

tautologies are maximal.

The set cut_{\leq}

Definition 7.8

Let \leq be an epistemic entrenchment ordering and φ a sentence.

Define $cut_{\leq}(\varphi) = \{\psi \mid \varphi \leq \psi\}$.

The set $cut_{\leq}(\varphi)$ contains all those sentences that are at least as entrenched as φ .

Epistemic entrenchments are total preorders (i.e. a reflexive and transitive binary relation) of sentences in \mathcal{L} .

Theorem 7.9

If \leq is an epistemic entrenchment, then for any sentence φ , $\text{cut}_{\leq}(\varphi)$ is a theory.

Proof: exercise.

Theorem 7.10

Let T be a theory of \mathcal{L} .

For every contraction function $\bar{\cdot}$ for T there exists an epistemic entrenchment \leq related to T s.t. (E^-) is true for every $\varphi \in \mathcal{L}$.

Conversely, for every epistemic entrenchment \leq related to T , there exists a contraction function $\bar{\cdot}$ s.t. (E^-) is true for every $\varphi \in \mathcal{L}$.

$$(E^-) \quad T_{\varphi}^- = \begin{cases} \{\psi \in T \mid \varphi < (\varphi \vee \psi)\} & \text{if } \not\vdash \varphi \\ T, & \text{otherwise} \end{cases}$$

Proof sketch: blackboard

Consequences of Theorem 7.10

Theorem 7.10 is a **constructive approach** to modeling belief contraction, in contrast to the AGM postulates (*1) to (*8).

Theorem 7.10 is also a **representation result**:

- Epistemic entrenchment \leq related to T 'holds all the information' needed to devise a contraction function.
- Every contraction function can be constructed from some epistemic entrenchment ordering—by condition (E^-).

Does it help?

Example 7.11

Reconsider Example 7.1: "all birds fly", "Tweety is a bird"

Consequence: "Tweety can fly"

New fact: "Tweety cannot fly"

- If "Tweety can fly \vee all birds fly" is strictly more entrenched than "Tweety can fly", then "all birds fly" will remain after contraction of "Tweety can fly".
- Conversely, if "Tweety can fly \vee all birds fly" and "Tweety can fly" are equally entrenched, then "all birds fly" will be retracted.

An analogous result for revision

Theorem 7.12

Let T be a theory.

For every function $*$ for T there exists an epistemic entrenchment \leq related to T s.t. (E^*) , below, is true for every $\varphi \in \mathcal{L}$.

Conversely, for every epistemic entrenchment \leq related to T , there exists a revision function $*$ for T s.t. (E^*) s is true for every $\varphi \in \mathcal{L}$.

$$(E^*) \quad T_{\varphi}^* = \begin{cases} \{\psi \in \mathcal{L} \mid \neg\varphi < (\neg\varphi \vee \psi)\} & \text{if } \not\vdash \neg\varphi \\ \perp, & \text{otherwise} \end{cases}$$

Proof: exercise.

Interconnections between nonmonotonic reasoning and belief revision

Nonmonotonic reasoning and belief revision are closely related:

Both rely on the property of minimality.

Both semantics can be described by a mechanism that selects minimal models (a.k.a. "most expected", "most preferred" according to the agent's understanding).

One can view revision as a nonmonotonic inference operation (see Exercise 11.1).

If $\neg\varphi \notin T$, then the inference operation is **monotonic**.

If $\neg\varphi \in T$ it is **nonmonotonic**.

Translation between revision and nonmonotonic inference

From a revision function to an inference relation:

- Theorem 7.12: epistemic entrenchment related to a theory gives a revision function. Thus an inference operation C can be defined from the epistemic entrenchment \leq where T represents the background information from which default information is derived by $*$.
- Conversely, inference relation $C_{<}$ can be constructed from preferential model structures (Definition 6.11). In particular, $\{\varphi\} \sim \psi$ holds in minimal models containing φ .
- Epistemic entrenchment orderings can be used to build preferential model structures whose minimal models containing φ are precisely the models for T_{φ}^* .

Thus the inference $\{\varphi\} \sim \psi$ can be described using a revision function.