Principles of Knowledge Representation and Reasoning

Dynamics of belief

Albert-Ludwigs-Universität Freiburg

UNI FREIBURG

Bernhard Nebel, Stefan Wölfl, and Julien Hué

Winter Semester 2013/2014

1 Introduction

Introductory example

Gärdenfors - 1988

Oscar used to believe that he had given Victoria a gold ring at their wedding. He had bought their two rings at a jewellery in Casablanca. He thought it was a bargain. The merchant had claimed that the rings were made of 24 carat gold. They certainly looked like gold, but to be on the safe side Oscar had taken the rings to the jeweller next door who has testified to their gold content. However, some time after the wedding, Oscar was repairing his boat and he noticed that the sulphuric acid he was using stained his ring. He remembered from his school chemistry that the only acid that affected gold was aqua regia. Somewhat surprised, he verified that the ring was also stained by the acid.

Principles

Propositional logic flaws:

- The world is not always static.
- The knowledge about the world is sometimes uncertain or imprecise

Therefore:

- Need to incorporate new (possibly contradictory) beliefs
- Need to take into account change in the world

Belief base revision

- How to react to new information? K is the old information, A the new information

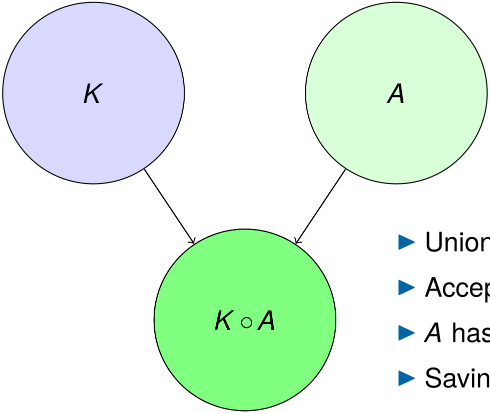

- Union $\rightarrow$ inconsistency
- Accept loss of beliefs
- A has priority over K
- Saving the most from K

The Gettier argument

Plato - Theaetetus: knowledge is

1. justified true belief

Agrippa's trilemma - A problem with the justification:

1. Either the justification stops at some unjustified belief;
2. The justification is infinite (Socrates' clouds);
3. The justification is supported by affirmations it is supposed to justify (Baron Münchhausen's hair).

Foundationalism and coherentism

Three solutions:

Foundationalism: allow for unjustified beliefs
$\rightarrow$ Formalization issues
$\rightarrow$ Humans don't keep track of sources
$\rightarrow$ TMS System

"Infinitism": allow for infinite justification
$\rightarrow$ Does it really make sense?

Coherentism: allow for circular justifications
$\rightarrow$ What is a solid belief?
$\rightarrow$ Belief revision/update

Revision or update?

- We have a theory about the world, and the new information is meant to correct our theory

$\rightsquigarrow$ belief revision: change your belief state minimally in order to accommodate the new information

- We have a (supposedly) correct theory about the current state of the world, and the new information is meant to record a change in the world

$\rightsquigarrow$ belief update: incorporate the change by assuming that the world has changed minimally

Overview of an operation

What are the criteria for definition of a belief revision operation?

Gärdenfors and Rott - belief revision (1995):

1. How are beliefs represented?
2. What is the relation between beliefs represented explicitly in the belief base and beliefs which can be derived from them?
3. In the face of a contradiction, how to deal with both new and old information?

Social choice theory: Arrow’s theorem

Arrow’s impossibility theorem - there is no voting system which respects:

- Non-dictatorship
  (all voters should be taken into account)
- Universality
  (complete and deterministic ranking)
- Independance of irrelevant alternatives
  (ranking between x and y depends only on x and y)
- Pareto efficiency
  (if all preferences states $x < y$, then so must the results)

Consequence

There is no perfect belief operation

Belief base, belief set or interpretation?

Introduction
Belief revision
A little bit of update
Bibliography

General assumptions:

- A belief set is a deductively closed theory, i.e., $K = \mathrm{Cn}(K)$ with Cn the consequence operator
- $\mathcal{L}$: logical language (propositional logic)
- $\mathrm{Th}_{\mathcal{L}}$: set of deductively closed theories (or belief sets) over $\mathcal{L}$

Belief change operations

Monotonic addition: $+ : \mathrm{Th}_{\mathcal{L}} \times \mathcal{L} \to \mathrm{Th}_{\mathcal{L}}$

$$K + \psi = \mathrm{Cn}(K \cup \{\psi\})$$

Revision: $\dot{+} : \mathrm{Th}_{\mathcal{L}} \times \mathcal{L} \to \mathrm{Th}_{\mathcal{L}}$

Consider $K = \{a, b\}$ and $K' = \{a \wedge b\}$. What is happening to $K \dot{+} \{\neg a\}$?

Semantic

- No difference between K and K'

a	b	$\mathcal{I}$
0	0	0
0	1	0
1	0	0
1	1	1

Syntactic

- $X = \{b\}$ is the only maximal subset of K s.t. $X \cup \{\neg a\}$ is consistent.
- $X' = \emptyset$ is the only maximal subset of K' s.t. $X' \cup \{\neg a\}$ is consistent.

2 Belief revision

- Syntactic approaches
- Semantic approaches

What is a good revision operator?

- Consistency: a revision has to produce a consistent set of beliefs
- Minimality of change: a revision has to change as few beliefs as possible
- Priority to the new information: the ’new’ information is considered more important than the ’old’ one

To characterize good operators, Alchourron, Gärdenfors, and Makinson identified postulates a “good” revision operator should respect.

The expansion operation

Theory expansion corresponds to the introduction of a formula into the theory, without modification of the initial theory.

Definition

The expansion of the theory T by the formula A is defined as

$$T + A = \mathrm{Cn}(T \cup \{A\})$$

Example

$$\begin{aligned} T &= \{a, b \rightarrow c\} \\ A &= \{b\} \\ T + A &= \{a, b, b \rightarrow c, c\} \end{aligned}$$

The contraction operation

Theory contraction corresponds to the removal of a formula from the theory.

Definition

The result of the contraction of the theory T by the formula A, denoted by $T - A$, is defined as a subset T' of T such that $T' \not\models A$

Example

$$
\begin{aligned}
T &= \{a, b, b \rightarrow c\} \\
A &= \{c\} \\
T - A &= \mathrm{Cn}(\{a, b \rightarrow c\}) \\
T - A &= \mathrm{Cn}(\{a, b\}) \\
T - A &= \mathrm{Cn}(\{a\}) \\
T - A &= \mathrm{Cn}(\emptyset)
\end{aligned}
$$

The Levi identity

Revision can be defined in terms of two suboperations.

- $+$ (expansion) denotes the simple union of beliefs;
- $-$ (contraction) denotes the removal of information contradicting the input.

The Levi identity

$$K \dot{+} \varphi \equiv Cn[(K - \neg\varphi) + \varphi]$$

Example

$$\begin{aligned} T &= \{a, b, b \rightarrow c\} \\ A &= \{\neg c\} \\ T \dot{+} A &= \text{Cn}(\{a, b \rightarrow c, \neg c\}) \\ T \dot{+} A &= \text{Cn}(\{a, b, \neg c\}) \\ T \dot{+} A &= \text{Cn}(\{a, \neg c\}) \\ T \dot{+} A &= \text{Cn}(\{\neg c\}) \end{aligned}$$

Remainder Set

Definition

Let K be a collection of formulae and A be a formula. The A-remainder set of K, denoted by $K \perp A$, is the collection of subsets Γ of $\mathcal{L}$ such that:

1. $\Gamma \subseteq K$
2. $A \notin Cn(\Gamma)$
3. There is no set Γ' such that $\gamma \subset \Gamma' \subseteq K$ and $A \notin Cn(\Gamma')$

Full-meet contraction

Definition

Full-meet contraction is defined by $K - \varphi = \bigcap(K \perp \varphi)$.

Is full-meet contraction reasonable?

- No! It is far too cautious.
- It can nevertheless be used as a lower bound to any reasonable operator.

$K \dot{+} \varphi = \bigcap(K \perp \neg\varphi) + \varphi$ is referred to as the full-meet revision.

Being Reasonable?

AGM postulates:

($\dot{+}$1) $K \dot{+} \varphi \in \mathrm{Th}_{\mathcal{L}}$

($\dot{+}$2) $\varphi \in K \dot{+} \varphi$;

($\dot{+}$3) $K \dot{+} \varphi \subseteq K + \varphi$

($\dot{+}$4) If $\neg\varphi \notin K$, then $K + \varphi \subseteq K \dot{+} \varphi$

($\dot{+}$5) $K \dot{+} \varphi = \mathrm{Cn}(\bot)$ only if $\vdash \neg\varphi$

($\dot{+}$6) If $\vdash \varphi \leftrightarrow \psi$ then $K \dot{+} \varphi = K \dot{+} \psi$

Proposition

Full-meet revision respects all AGM postulates.

Proof

($\dot{+}1$) and ($\dot{+}2$) are true by construction

($\dot{+}3$) Two cases: (1) If $K+\varphi$ is consistent then $K-\varphi = K$ and $K \dot{+} \varphi = K+\varphi$. (2) If $K+\varphi$ is inconsistent then $K+\varphi = \mathrm{Cn}(\bot)$ and $K \dot{+} \varphi \subseteq K+\varphi$.

($\dot{+}4$) Because $K \nvdash \neg\varphi$ then $K \bot \varphi = \{K\}$ and thus $K \dot{+} \varphi = K+\varphi$.

($\dot{+}5$) $K \dot{+} \varphi = \mathrm{Cn}\left(\cap_{\alpha \in (K\bot\varphi)} \alpha \cup \varphi\right)$. But $\forall \alpha, \alpha \cup \varphi \nvdash \bot$, therefore $\cap_{\alpha \in (K\bot\varphi)} \alpha \cup \varphi \nvdash \bot$ (as PL is monotonic).

($\dot{+}6$) Lets assume that $\alpha \in K\bot\varphi$ but $\alpha \notin K\bot\Psi$. Two cases: (1) $\alpha \cup \Psi \vdash \bot \overset{(\varphi\leftrightarrow\Psi)}{\longrightarrow} \alpha \cup \varphi \vdash \bot$ which is not possible. (2) $\exists \beta$ s.t. $\alpha \subsetneq \beta$ and $\beta \cup \Psi \nvdash \bot \overset{(\varphi\leftrightarrow\Psi)}{\longrightarrow} \beta \cup \varphi \nvdash \bot$ which is not possible.

Complexity result

Formula-based approaches

The question of whether Ψ belongs to $K \dot{+} \varphi$ (if $\dot{+}$ is a full-meet revision operator) is $\Delta_2^p - (\Sigma_1^p \cup \Pi_1^p)$ provided that NP $\neq$ co-NP.

Proof

If $\dot{+}$ is a full-meet revision, $\Psi \in \mathrm{Cn}(K) \dot{+} \varphi$ can be solved by the following algorithm: if $K \not\models \neg\Psi$, then $K \cup \Psi \models \varphi$ else $\Psi \models \varphi \longrightarrow$ Membership in Δ_2^p.
Furthermore, SAT can be polynomially transformed to full-meet revision by solving $\Psi \in \mathrm{Cn}(\Psi) \dot{+} \top$ and UNSAT can be polynomially transform to full-meet revision by solving $\bot \in \mathrm{Cn}(\emptyset) \dot{+} \Psi$. Hence, assuming that full-meet revision belongs to both NP and co-NP would lead to NP = co-NP.

Maxi-choice revision

On the other side, one can ask for the principle of minimality to be strictly respected.

Definition

A selection function for K is a function γ such that for all sentences φ:

1. If $K \perp \neg\varphi$ is non-empty, then $\gamma(K \perp \neg\varphi)$ is a non-empty subset of $K \perp \neg\varphi$, and
2. If $K \perp \neg\varphi$ is empty, then $\gamma(K \perp \neg\varphi) = \{K\}$.

Definition

Maxichoice contraction is defined as $K - \varphi = \gamma(K \perp \neg\varphi)$ where γ is a selection function.

Partial-meet revision

Maxi-choice can be too bold: there is sometimes no reason to trust one piece more than one another.

Definition

A partial-meet revision operation is an operation defined as:

$$K \dot{+} \varphi = \bigcap \gamma(K \perp \neg\varphi) + \varphi$$

Seems to be a good compromise between full-meet and maxi-choice

Example

$$K = \left\{ \begin{array}{ccc} a & a \rightarrow f & d \\ a \rightarrow g & d \vee e & c \vee e \\ f \rightarrow h & g \rightarrow h & \end{array} \right\} \qquad A = \{\neg h\}$$

$$K \perp \neg A = \left\{ \begin{array}{l} \{a \rightarrow f, d, a \rightarrow g, d \vee e, c \vee e, f \rightarrow h, g \rightarrow h\} \\ \{a, d, d \vee e, c \vee e, f \rightarrow h, g \rightarrow h\} \\ \{a, d, a \rightarrow g, d \vee e, c \vee e, f \rightarrow h\} \\ \{a, a \rightarrow f, d, d \vee e, c \vee e, g \rightarrow h\} \\ \{a, a \rightarrow f, d, a \rightarrow g, d \vee e, c \vee e\} \end{array} \right\}$$

$$K \perp \neg A = \left\{ \begin{array}{l} \{a \rightarrow f, {\color{red}d}, a \rightarrow g, {\color{red}d \vee e}, {\color{red}c \vee e}, f \rightarrow h, g \rightarrow h\} \\ \{a, {\color{red}d}, {\color{red}d \vee e}, {\color{red}c \vee e}, f \rightarrow h, g \rightarrow h\} \\ \{a, {\color{red}d}, a \rightarrow g, {\color{red}d \vee e}, {\color{red}c \vee e}, f \rightarrow h\} \\ \{a, a \rightarrow f, {\color{red}d}, {\color{red}d \vee e}, {\color{red}c \vee e}, g \rightarrow h\} \\ \{a, a \rightarrow f, {\color{red}d}, a \rightarrow g, {\color{red}d \vee e}, {\color{red}c \vee e}\} \end{array} \right\}$$

Full-meet contraction:

Reformulation of AGM postulates

Katsuno and Mendelzon refomulation of AGM: knowledge bases are here represented as formulas.

Definition

Let φ, μ and χ be formulas

(R1) $\varphi \circ \mu$ implies μ

(R2) If $\varphi \wedge \mu$ is satisfiable, then $\varphi \circ \mu \equiv \varphi \wedge \mu$

(R3) If μ is satisfiable, then so is $\varphi \circ \mu$

(R4) If $\varphi_1 \equiv \varphi_2$ and $\mu_1 \equiv \mu_2$, then $\varphi_1 \circ \mu_1 \equiv \varphi_2 \circ \mu_2$

(R5) $(\varphi \circ \mu) \wedge \chi$ implies $\varphi \circ (\mu \wedge \chi)$

(R6) If $(\varphi \circ \mu) \wedge \chi$ is satisfiable, then $\varphi \circ (\mu \wedge \chi)$ implies $(\varphi \circ \mu) \wedge \chi$

Remark

Not to be confused with Katsuno and Mendelzon postulates for update

Faithful Assignment

Definition

A preorder $\leq$ over $\mathcal{I}$ is a reflexive and transitive relation on $\mathcal{I}$.

- $\leq$ is total if $\forall I, I' \in \mathcal{I}, I \leq I'$ or $I' \leq I$
- Assume that to each φ, there is an assigned preorder $\leq_\varphi$

Definition

The assignment $\varphi \mapsto \leq_\varphi$ is faithful iff

1. $I, I' \in \text{Mod}(\varphi)$ implies $I \not< _\varphi I'$
2. $I \in \text{Mod}(\varphi)$ and $I' \notin \text{Mod}(\varphi)$ implies $I <_\varphi I'$
3. $\varphi \leftrightarrow \chi$ implies $\leq_\varphi = \leq_\chi$

Faithful Assignment versus Postulates

Theorem (From Katzuno-Mendelzon)

A Revision operator $\circ$ satisfies (R1)-(R6) iff there exists a faithful assignment that maps each sentence φ into a total preorder $\leq_\varphi$ such that:

$$Mod(\varphi \circ \mu) = \min(Mod(\mu), \leq_\varphi)$$

- Epistemic states versus belief sets/bases/interpretations

Distance-based revision operations

Definition

The Dalal revision operation, denoted by $\circ_D$, is defined by:

$$K \circ_D \varphi = \min(\text{Mod}(\varphi), \leq_K)$$

where d_H is the Hamming Distance and

$$\alpha \leq_K \beta \qquad \text{iff}$$

$$\exists \omega \in \text{Mod}(K), \forall \omega' \in \text{Mod}(K), d_H(\alpha, \omega) \leq d_H(\beta, \omega')$$

Example

	a	b	c
I_{φ_1}	0	0	0
I_{φ_2}	0	0	1
	0	1	0
I_{K_1}	0	1	1
	1	0	0
I_{K_2}	1	0	1
	1	1	0
I_{K_3}	1	1	1

Let $\varphi = \{\neg a, \neg b\}$ and $K = \{(a \vee b) \wedge c\}$:

$d(I_{\varphi_1}, I_{K_1}) = 2 \quad d(I_{\varphi_2}, I_{K_1}) = 1$
$d(I_{\varphi_1}, I_{K_2}) = 2 \quad d(I_{\varphi_2}, I_{K_2}) = 1$
$d(I_{\varphi_1}, I_{K_3}) = 3 \quad d(I_{\varphi_2}, I_{K_3}) = 2$

Spohn's Ordinal Conditional Function

Ordinal Conditional Function associates an ordinal number to each interpretation.

- This number, denoted by $\kappa(A)$, represents a degree of disbelief (0 being the most plausible)
- For some formula A, either $\kappa(A) = 0$ or $\kappa(\neg A) = 0$
- $\kappa(A) = \min_{I \in \mathrm{Mod}(A)} \kappa(I)$
- $\kappa(A \vee B) = \min(\kappa(A), \kappa(B))$
- A is accepted if $\kappa(\neg A) > 0$

Spohn's Ordinal Conditional Function

- The belief base is a set of ranked models
- The added information is a pair (μ, m)

$$\kappa_{(\mu,m)}(I) = \begin{cases} \kappa(I) - \kappa(\mu) & \text{if } I \in \mathrm{Mod}(\mu) \\ \kappa(I) - \kappa(\neg\mu) + m & \text{if } I \notin \mathrm{Mod}(\mu) \end{cases}$$

Spohn's Ordinal Conditional Function

Example

$$\varphi = \{a, b, (a \wedge b) \rightarrow c\}$$

has one model $I = \{a, b, c\}$ with ranking $\kappa(I) = 0$. The other ones have a ranking of 1. $\mu = \{\neg c\}$ with a post-revision degree of 3.

$$\text{Mod}(\mu) = \left\{ \begin{array}{ccc} J_1 = \{a, b, \neg c\} & , & J_2 = \{a, \neg b, \neg c\}, \\ J_3 = \{\neg a, b, \neg c\} & , & J_4 = \{\neg a, \neg b, \neg c\} \end{array} \right\}$$

The result of the revision process is:

$$\begin{array}{lll} \kappa_{\mu,3}(I) = 2, & \kappa_{\mu,3}(J_1) = 0, & \kappa_{\mu,3}(J_2) = 0, \\ \kappa_{\mu,3}(J_3) = 0, & \kappa_{\mu,3}(J_4) = 0, & \forall I' \neq I, J_i, \kappa_{\mu,3}(I') = 3 \end{array}$$

3 A little bit of update

Update and revision are different

Assume the new information is consistent with our old beliefs.

- In case of belief revision, we would like to add the new information monotonically to our old beliefs.
- For belief update this is not necessarily the case.
  - Assume we know that the **door is open or the window is open**.
  - Assume we learn that the world has changed and the **door is now closed**.
  - In this case, we do not want to add this information monotonically to our theory, since we would be forced to conclude that **the window is open**.

General understanding of update and revison

Revision: The worlds has not changed but we have a better information than previously

- For the revision of φ by μ: choose the models of μ which are the closest to φ

Update: The world has changed

- For the update of φ by μ: choose for each models of φ, the models of μ which are the closest

Katsuno and Mendelzon Postulates

Definition

Let φ, μ, χ be formulas.

(U1) $\varphi \diamond \mu$ implies μ

(U2) If φ implies μ then $\varphi \diamond \mu \equiv \varphi$

(U3) If φ and μ are satisfiable then $\varphi \diamond \mu$ is also satisfiable

(U4) If $\varphi_1 \equiv \varphi_2$ and $\mu_1 \equiv \mu_2$ then $\varphi_1 \diamond \mu_1 \equiv \varphi_2 \diamond \mu_2$

(U5) $(\varphi \diamond \mu) \wedge \chi$ implies $\varphi \diamond (\mu \wedge \chi)$

(U6) If $\varphi \diamond \mu_1$ implies μ_2 and $\varphi \diamond \mu_2$ implies μ_1 then $\varphi \diamond \mu_1 \equiv \varphi \diamond \mu_2$

(U7) If φ is complete then $(\varphi \diamond \mu_1) \wedge (\varphi \diamond \mu_2)$ implies $\varphi \diamond (\mu_1 \vee \mu_2)$

(U8) $\varphi \diamond (\mu_1 \vee \mu_2) \equiv (\varphi \diamond \mu_1) \wedge (\varphi \diamond \mu_2)$

A formula φ is complete, if for any formula μ, φ implies μ or φ implies $\neg \mu$.

Faithful update assignment

Definition

An update operator $\diamond$ satisfies conditions (U1)-(U8) iff there exists a faithful update assignment that maps each interpretation I to a partial pre-order (or order) $\leq_I$ such that:

$$\mathrm{Mod}(\Psi \diamond \mu) = \bigcup_{I \in \mathrm{Mod}(\Psi)} \min(\mathrm{Mod}(\mu), \leq_I)$$

a formula φ is complete, if for any formula μ, φ implies μ or φ implies $\neg\mu$.

Update - Example

The next example has been taken from Katsuno and Mendelzon.

Example

5-bit unchanger register

$$\varphi = 10000 \vee 00111$$

$$\mu = 11111 \vee 00000$$

- revision of φ by μ gives 00000
- update of φ by μ gives $00000 \vee 11111$

4 Bibliography

Literature I

Peter Gärdenfors and Hans Rott,
Belief revision,
Handbook of Logic in AI and LP, 1995.

Carlos E. Alchourròn, Peter Gärdenfors, David Makinson,
On the Logic of Theory Change: Partial Meet Contraction and Revision Functions,
Journal of Symbolic Logic, 1985.

Bernhard Nebel
Base Revision Operations and Schemes: Semantics, Representation and Complexity
ECAI, 1994.