UNIVERSITÄT ZU LÜBECK
INSTITUT FÜR INFORMATIONSSYSTEME

Özgür L. Özçep

Ontology Change 1

Lecture 9: AGM Belief Revision

20 January, 2016

Foundations of Ontologies and Databases for Information Systems

CS5130 (Winter 2015)

Recap of Lecture 8

OBDA

- Ontology-Based Data Access in classical sense
- Rewriting: Reasoning services provided by rewriting them into query without TBox
- Complete (and correct) rewriting guaranteed for lightweight logics
- Unfolding: Transform (rewritten) query into query of backend source w.r.t. mappings

End of Recap

Solutions for Exercise 6

Exercise 6.1 (2 Points)

Prove that DL-Lite$_{\mathcal{F}}$ can have ontologies having only infinite models (using, e.g., the example mentioned in the lecture)

Solution:

- We consider ontology $\mathcal{O}$ from the lecture
  - $\mathit{Nat} \sqsubseteq \exists \mathit{hasSucc}$, $\exists \mathit{hasSucc}^- \sqsubseteq \mathit{Nat}$, $(\mathit{funct}\ \mathit{hasSucc}^-)$,
  - $\mathit{Zero} \sqsubseteq \mathit{Nat}$, $\mathit{Zero} \sqsubseteq \neg\exists \mathit{hasSucc}^-$, $\mathit{Zero}(0)$
- We prove by induction on $n \in N$:
  for all n there is a non-cyclic *hasSucc* path with start point 0.
  - $n = 0$: there is zero path from 0 to 0.
  - $n \mapsto n+1$: Assume there is a non-cyclic n-path P from 0. Let d_n denote the last node in the past. It must have successor d_{n+1}. But this one can not be one of the nodes in P as otherwise one node would have two predecessors. Hence we can add the *hasSucc* edge (d_n, d_{n+1}) to P, reaching a non-cyclic path of length $n+1$.
- A finite model does not allow for paths of arbitrary lengths

Exercise 6.2 (3 Points)

The anonymization function in the PerfRew algorithm is allowed to be applied only to unbound variables that are not distinguished (that is are not answer variables). Give an example why this restriction makes sense.

Solution:

- Consider the following ontology $\mathcal{O} = (\mathcal{T}, \mathcal{A})$ and query
  - $\mathcal{T} = \{A \sqsubseteq \exists R, B \sqsubseteq \exists S\}$
  - $\mathcal{A} = \{A(a), B(a)\}$
  - $Q(x) = \exists y.R(x, y) \wedge S(x, y)$
- $a \notin cert(Q, \mathcal{O})$, as the following model $\mathcal{I} \models \mathcal{O}$ demonstrates
  - $\Delta^{\mathcal{I}} = \{a, b, c\}$
  - $(a)^{\mathcal{I}} = a$
  - $A^{\mathcal{I}} = \{a\}, B^{\mathcal{I}} = \{a\}$
  - $R^{\mathcal{I}} = \{(a, b)\}$, $S^{\mathcal{I}} = (a, c)\}$ (Note that $b \neq c$)
- Would we anonymize the y in Q we would get the query
- $Q'(x) = R(x, \_) \wedge S(x, \_)$
- Applying the TBox axioms would result in

  $Q''(x) = A(x) \wedge B(x)$, but $a \in cert(Q'', \mathcal{A})$

Exercise 6.3 (3 Points)

Explain the notion of reification, and show (with an example) why it is needed for (classical) OBDA.

Solution:

- Reification denotes a method to represent semantical objects such as sentences or relations as objects in the domain.
- Reification is necessary if one, e.g., wants to represent ternary predicates in a language allowing maximally binary predicates (such as DLs as used on OBDA).

Exercise 6.4 (4 Points)

Many relevant DL reasoning services can be reduced to ontology satisfiability in DL-Lite. Show that subsumption w.r.t. a DL-Lite TBox can be reduced to (un)satisfiability test of a DL-Lite ontology!

Hint: Use the general fact of entailment that $\psi \models \phi$ iff $\psi \wedge \neg\phi$ is unsatisfiable. Then think of how the latter can be formulated in a DL-Lite ontology (introducing perhaps new symbols).

Solution:
We have to find an equivalent representation for $\mathcal{T} \models C \sqsubseteq D$. We know that $\mathcal{T} \models C \sqsubseteq D$ holds iff (abusing notation): $\mathcal{T} \cup \neg(C \sqsubseteq D)$ is unsatisfiable, i.e., if there is an c such that $\mathcal{T} \cup \{C(c) \wedge \neg D(c)\}$ is unsatisfiable.
As we are allowed to use only atomic symbols in the ABox, we represent $\{C(c) \wedge \neg D(c)\}$ as $\{A \sqsubseteq C, A \sqsubseteq \neg D, A(c)\}$. So we have to show formally the reduction (with symbols A, c not occurring in $\mathcal{T}$):

$$\mathcal{T} \models C \sqsubseteq D \text{ iff } \mathcal{O} := (\mathcal{T} \cup \{A \sqsubseteq C, A \sqsubseteq \neg D\}, \{A(c)\}) \text{ is unsatisfiable}$$

"$\Rightarrow$": Assume that $\mathcal{O}$ is satisfiable by $\mathcal{I}$. But $\mathcal{I} \models \mathcal{T}$ and $(c)^{\mathcal{I}} \in C^{\mathcal{I}}$ but $(c)^{\mathcal{I}} \notin D^{\mathcal{I}}$.

"$\Leftarrow$" : Assume that $\mathcal{O}$ is un-satisfiable and assume for contradiction that not $\mathcal{T} \models C \sqsubseteq D$. Then there must be a model $\mathcal{I} \models \mathcal{T}$ and $d \in \Delta^{\mathcal{I}}$ with $d \in C^{\mathcal{I}}$ but $d \notin D^{\mathcal{I}}$. We can now extend $\mathcal{I}$ to a model $\mathcal{I}'$ which is the same as $\mathcal{I}$ for all symbols except for A and c. We let $A^{\mathcal{I}'} = \{d\}$, $c^{\mathcal{I}'} = d$. But then $\mathcal{I}' \models \mathcal{O}$, contradicting the assumption from the beginning.

References

- Eduardo Ferme: Belief Revision from 1985 to 2013
  Slides of IJCAI 2013-Tutorial
  `http://www.ijcai13.org/files/tutorial_slides/ta4.pdf`
- **Lit:** P. Gärdenfors. Knowledge in Flux: Modeling the Dynamics of Epistemic States. The MIT Press, Bradford Books, Cambridge, MA, 1988.
- **Lit:** S. O. Hansson. A Textbook of Belief Dynamics. Kluwer Academic Publishers, 1999.

Motivation

Ontology-Level Integration

- So far: Two (different) types of integration
  - Data exchange: directed schema-level integration over finite DBs
  - OBDA: directed schema-level-to-ontology integration
- We consider now: ontology-level integration (in these lectures: mainly directed integration)
- Required in different ontology change scenarios where multiple (versions of) ontologies: exist ontology import, merge, versioning, development, alignment, articulation etc.

  Lit: G. Flouris et al. Ontology change: classification and survey. The Knowledge Engineering Review, 23(2):117–152, 2008.
- Main problem to tackle in all of them: Joined ontology may be incompatible (incoherent, inconsistent)

Example (Incompatible ontologies)

$\mathcal{O}_A$

A1 $Article \equiv \exists publ.Journal$

A2 $Journal \sqsubseteq \neg Proceedings$

A3 $(func\ publ)$

$\mathcal{O}_B$

B1 $Article \equiv \exists publ.Journal \sqcup Proceedings$

B2 $publish(ab, procXY)$

B3 $Proceedings(procXY)$

- $\mathcal{O}_A \cup \mathcal{O}_B$ is inconsistent
- How to repair this?
  - Find all culprits (group) (Here one group: $\mathcal{O}_A \cup \mathcal{O}_B$)
  - If culprit group has more than one sentence, which to eliminate? (Here: Eliminate A1 or ... or B3?)

$\Longrightarrow$ Research field Ontology Change (OC)

- This lecture: Research field Belief Revision (BR)
- Next lecture: Extensions of BR w.r.t. OC and OC in detail

Belief Revision (BR)

- 31 years aged interdisciplinary research field in philosophy, cognitive science, CS
- Landmark paper by AGM (Alchourrón, Gärdenfors, Makinson)
  Lit: C.E. Alchourrón, P. Gärdenfors, and D. On the logic of theory change: partial meet contraction and revision functions. Journal of Symbolic Logic, 50:510–530, 1985.
- BR deals with operators for revising theories under possible inconsistencies
  - Investigates concrete revision operators
  - Principles that these must fulfill
  - Representation theorems
- Recent research how to adapt these for non-classical logics/ontologies, mappings, programs.

Terminology

- Unfortunately the field of Belief Revision is called after the particular class of revision operators
- But it handles other types of changing beliefs/theories: expansion, update, and contraction
- We stick to this folklore use and hide it behind the acronym BR

AGM Postulates

Consequence Operator

- AGM framework based on general notion of logic in polish tradition

  Lit: R. Wójcicki. Theory of Logical Calculi. Kluwer Academic Publishers, Dordrecht, 1988.

- Logic: $(\mathcal{L}, Cn)$
  - $\mathcal{L}$: Set of well-formed sentences
  - Cn: Consequence operator $Pow(\mathcal{L}) \longrightarrow Pow(\mathcal{L})$
    Note: No distinction between syntax and semantics

Definition (Tarskian consequence operator)

For all $X, X_1, X_2 \subseteq \mathcal{L}$:

1. $X \subseteq Cn(X)$ (Inclusion)
2. If $X_1 \subseteq X_2$, then $Cn(X_1) \subseteq Cn(X_2)$. (Monotonicity)
3. $Cn(X) = Cn(Cn(X))$ (Idempotence)

Definition (Tarskian consequence operator)

For all $X, X_1, X_2 \subseteq \mathcal{L}$:

1. $X \subseteq Cn(X)$ (Inclusion)
2. If $X_1 \subseteq X_2$, then $Cn(X_1) \subseteq Cn(X_2)$. (Monotonicity)
3. $Cn(X) = Cn(Cn(X))$ (Idempotence)

Wake-Up questions

- How would one define an entailment relation based on Cn—and vice versa?
- In natural language speak explain what the following mean
  - $Cn(X) = \mathcal{L}$
  - $\alpha \in Cn(\emptyset)$
  - $\neg\alpha \in Cn(\emptyset)$

AGM Consequence Operator

Definition (Tarskian consequence operator)

For all $X, X_1, X_2 \subseteq \mathcal{L}$:

1. $X \subseteq Cn(X)$ (Inclusion)
2. If $X_1 \subseteq X_2$, then $Cn(X_1) \subseteq Cn(X_2)$. (Monotonicity)
3. $Cn(X) = Cn(Cn(X))$ (Idempotence)

- AGM additionally demands that Cn fulfills
  - **Supra-classicality:** If α can be derived from X by propositional logic, then $\alpha \in Cn(X)$
  - **Compactness:** If $\alpha \in Cn(X)$m then $\alpha \in Cn(X')$ for some finite $X' \subseteq X$.
  - **Deduction:** $\beta \in Cn(X \cup \{\alpha\})$ iff $(\alpha \rightarrow \beta) \in Cn(X)$

Belief Sets

Definition (Belief Set)

- Belief set (BS) for $(\mathcal{L}, Cn)$ is a set of the form $Cn(X)$ for $X \subseteq \mathcal{L}$.
- $\mathcal{BS}_{\mathcal{L}}$ = Set of all belief sets for $(\mathcal{L}, Cn)$

- Idealization of the beliefs of a rational agent

- AGM consider (inter-related) operators for changing BSs into new BSs under a single trigger sentence $\in \mathcal{L}$
- Types of AGM change operators $\mathcal{BS}_{\mathcal{L}} \times \mathcal{L} \longrightarrow \mathcal{BS}_{\mathcal{L}}$
  - **Expansion:** add trigger and closed up w.r.t. Cn
  - **Contraction:** delete trigger from BS
  - **Revision:** add trigger and eliminate inconsistencies

AGM Postulates for Expansion

(E1) $K + \alpha \in \mathcal{BS}_{\mathcal{L}}$ (Closure)

(E2) $\alpha \in K + \alpha$ (Success)

(E3) $K \subseteq K + \alpha$ (Inclusion)

(E4) If $\alpha \in K$, then $K = K + \alpha$. (Vacuity)

(E5) If $K \subseteq X$, then $K + \alpha \subseteq X + \alpha$. (monotonicity)

(E6) $K + \alpha$ is the smallest belief set fulfilling (E1)–(E5).

Note:

- Postulates defined for fixed belief set K.
- Postulates specify properties of intended BR operators
- In general, many structurally different operators may fulfill the postulates, but ...

AGM Postulates for Expansion

(E1) $K + \alpha \in \mathcal{BS}_{\mathcal{L}}$ (Closure)

(E2) $\alpha \in K + \alpha$ (Success)

(E3) $K \subseteq K + \alpha$ (Expansion 1)

(E4) If $\alpha \in K$, then $K = K + \alpha$. (Expansion 2)

(E5) If $K \subseteq X$, then $K + \alpha \subseteq X + \alpha$. (Monotonicity)

(E6) $K + \alpha$ is the smallest belief set fulfilling (E1)–(E5).

- ... (E1)–(E6) are such specific that they uniquely identify $+$

Proposition

An operator $+$ fulfills (E1)–(E6) iff for α: $K + \alpha = Cn(K \cup \alpha)$

- This is a **representation result**

AGM Postulates for Contraction

(C1) $K \div \alpha \in \mathcal{BS}_{\mathcal{L}}$ (Closure)

(C2) $K \div \alpha \subseteq K$ (Inclusion)

(C3) If $\alpha \notin K$, then $K = K \div \alpha$ (Vacuity)

(C4) If $\alpha \notin Cn(\emptyset)$, then $\alpha \notin K \div \alpha$. (Success)

(C5) If $\alpha \in K$, then $K \subseteq (K \div \alpha) + \alpha$. (Recovery)

(C6) If $\alpha \leftrightarrow \beta \in Cn(\emptyset)$, then $K \div \alpha = K \div \beta$. ((Right) Extensionality)

(C7) $K \div \alpha \cap K \div \beta \subseteq K \div (\alpha \wedge \beta)$ (Conjunction 1)

(C8) If $\alpha \notin K \div (\alpha \wedge \beta)$, then $K \div (\alpha \wedge \beta) \subseteq K \div \alpha$. (Conjunction 2)

AGM Postulates for Revision

(R1) $K * \alpha \in \mathcal{BS}_{\mathcal{L}}$ (Closure)

(R2) $\alpha \in K * \alpha$ (Success)

(R3) $K * \alpha \subseteq K + \alpha$ (Expansion 1/Inclusion)

(R4) If $\neg\alpha \notin K$, then $K + \alpha \subseteq K * \alpha$. (Expansion 2/Vacuity)

(R5) If $\bot \in Cn(K * \alpha)$, then $\neg\alpha \in Cn(\emptyset)$. (Consistency)

(R6) If $\alpha \leftrightarrow \beta \in Cn(\emptyset)$, then $K * \alpha = K * \beta$.
((Right) Extensionality)

(R7) $K * (\alpha \wedge \beta) \subseteq (K * \alpha) + \beta$ (Conjunction 1)

(R8) If $\neg\beta \notin K * \alpha$, then $(K * \alpha) + \beta \subseteq K * (\alpha \wedge \beta)$.
(Conjunction 2)

Mutual Interdefinability

- Intuitively, contraction is the more primitive operation. Indeed:

Theorem

The revision operator defined by the ***Levi Identity***

$$K * \alpha = (K \div \neg\alpha) + \alpha$$

fulfills (R1)-(R8) if $\div$ fulfills (C1)–(C8).

- But technically also contraction is definable by revision

Theorem

The contraction operator defined by the ***Harper Identity***

$$K \div \alpha = K \cap (K * \neg\alpha)$$

fulfills (C1)-(C8) if $*$ fulfills (R1)–(R8).

AGM Operators

Operators for Revision and Contraction Postulates

- We still did not see concrete revision and contraction operators
- We seek for models of Postulates (R1)–(R8) and (C1)–(C8).

- In contrast to $+$, the postulates do not fix a single operator but a whole class
- But: Postulates are so specific that the classes can be characterized by constructions principles.

- There are various construction principles leading to different classes
  - Partial meet
  - Safe/kernel
  - Epistemic entrenchment
  - Possible worlds
  - Sphere-based

Remainder Set

- Main construct underlying partial meet operators
- Describe maximal possible scenarios that are compatible with the negation of the trigger

Definition (Remainder Set Informally)

The **remainder set** $X \perp \alpha$ **of** X **by** α consists of all maximal subsets of X not entailing α.

The sets in $X \perp \alpha$ are called **remainders**.

Remainder Set

- Main construct underlying partial meet operators
- Describe maximal possible scenarios that are compatible with the negation of the trigger

Definition (Remainder Set formally)

The **remainder set** $X \perp \alpha$ **of** X **by** α consists of all sets X' s.t.:

1. $X' \subseteq X$;
2. $\alpha \notin Cn(X')$;
3. There is no X'', such that $X' \subsetneq X'' \subseteq K$ and $\alpha \notin Cn(X'')$.

Example (Hansson Dynamics of Belief, Exercise 26a,f)

- $\{p, q\} \perp (p \wedge q) = \{\{p\}, \{q\}\}$
- $\{p \vee r, p \vee \neg r, q \wedge s, q \wedge \neg s\} \perp p \wedge q =$
  $\{\{p \vee r, p \vee \neg r\}, \{p \vee r, q \wedge s\}, \{p \vee r, q \wedge \neg s\},$
  $\{p \vee \neg r, q \wedge s\}, \{p \vee \neg r, q \wedge \neg s\}\}$

Wake Up

Definition (Remainder Set Formally)

The **remainder set** $X \perp \alpha$ **of** X **by** α consists of all sets X' s.t.:

1. $X' \subseteq X$;
2. $\alpha \notin Cn(X')$;
3. There is no X'', s.t. $X' \subsetneq X'' \subseteq K$ and $\alpha \notin Cn(X'')$.

Wake-up Questions

- Show that the remainders for a belief set are by themselves remainders.

Selection Function

- Handle multiplicity of scenarios (remainder sets) with fairness condition
  $\Longrightarrow$ Apply selection function

Definition (Selection Function)

An **AGM-selection function** $\gamma : Pow(\mathcal{BS}_{\mathcal{L}}) \longrightarrow Pow(\mathcal{BS}_{\mathcal{L}})$ for K fulfills for all α:

1. If $K \perp \alpha \neq \emptyset$, then $\emptyset \neq \gamma(K \perp \alpha) \subseteq K \perp \alpha$;
2. $\gamma(\emptyset) = \{K\}$.

- γ is defined for a given K

Partial Meet

Definition

For a selection function γ define on K

- $K \div_\gamma \alpha = \bigcap \gamma(K \perp \alpha)$ **(Partial meet contraction)**
- $K *_\gamma \alpha = (K \div_\gamma \neg\alpha) + \alpha$ **(Partial meet revision)**
- **Maxi-Choice** = partial meet with $|\gamma(X)| = 1$.
- **Full meet** = partial meet change with $\gamma(X) = X$ (for $X \neq \emptyset$).

- Maxi-choice and full-meet are two extremes of partial meet change

Properties Maxi-Choice and Full-Meet

- Maxi-choice revision is all-too deterministic: It decides the status of any sentence

Theorem

Let $*_\gamma$ be a maxi-choice revision operator. Then, for any (!) $\beta \in \mathcal{L}$ either $\beta \in K *_\gamma \alpha$ or $\neg\beta \in K *_\gamma \alpha$

- Full-meet revision is too skeptical.

Theorem

Let $*_\gamma$ be a full-meet revision operator. Then for all α with $\neg\alpha \in K$: $K *_\gamma \alpha = Cn(\alpha)$.

Representation Theorem

- The basic axioms for AGM revision and contraction characterize the class of partial meet revision and partial meet contraction operators

Theorem

An operator $\div$ on belief set K fulfills (C1)–(C6) iff there is a selection function γ such that for all α:

$$K \div \alpha = K \div_\gamma \alpha$$

An operator $*$ on belief set K fulfills (R1)–(R6) iff there is a selection function γ such that for all α:

$$K * \alpha = K *_\gamma \alpha$$

- Partial-Meet operators do not necessarily fulfill the additional postulates (R7,8), (C7,8), resp.
- For this one considers γ with additional properties

Representation Theorems

- Representation theorem in a general sense
  - Given a class A of structures satisfying a set of axioms
  - Output: A class of structures B (adhering to some simple construction) such that any A-structure is isomorphic to some B-structure
  - Example: Stone's result that every boolean algebra is isomorphic to an algebra of sets
- Representation Theorems in BR are special cases
  - Domains of operators are fixed
  - Equality instead of isomorphism

Other Constructions for Concrete Operators

- Other equally powerful constructions exist that lead to representation theorems for AGM postulates

- **Kernel revision**
  - Consider duals to remainder set: kernels
  - kernel = Minimal set responsible for inconsistency (culprit group)
  - Revision: Revise by eliminating from every kernel at least one element

- **Rank** based revision (such as **epistemic entrenchment**)
  - Idea: Specify (partial) order on sentences w.r.t. a belief set
  - Revision: Eliminate the least epistemically entrenched ones

- **Possible Worlds** (see following slides)

AGM: criticism, extensions and more

AGM: the Core of BR Research

- AGM change operators have been criticized on different grounds again and again
- This shows importance of AGM rather than weakness
- We discuss criticisms of AGM, extensions, and alternative operators ...
- ... mainly with respect to use of BR for CS and ontology change

General Criticism: Recovery

Example

- Belief set K contains
  - Cleopatra had a son. (α)
  - Cleopatra had a daughter (β)
  - Cleopatra had a child. ($\alpha \vee \beta$)
- Contract with $\alpha \vee \beta$
- Then add $\alpha \vee \beta$.
- Why should one still hold to facts α and β?

- Recovery somehow wrongly implements minimality

General Criticisms: No Minimality

Example

- AGM postulates allow amnestic revision of form

$$K * \alpha = Cn(\alpha)$$

- This is not minimal in a genuine sense

- Lead to invention of relevance postulates
- Allow the elimination only of those parts which are relevant for the trigger

  Lit: R. Parikh. Beliefs, belief revision, and splitting languages. In Logic, Language and Computation, vol. 2, pages 266–278,1999.
- But there are also considerations why "dogma of minimality" is not satisfiable

  Lit: H. Rott. Two dogmas of belief revision. The Journal of Philosophy, 97(9):503–522, 2000.

General Criticism: Success postulate

Example

- Child: "There was a dinosaur in our flat who broke the vase"
- One wants to trust only some parts of information (a glass was broken) but not other parts (it was a dinosaur)

- Lead to non-prioritized belief revision: no priority for trigger

  Lit: S. O. Hansson. A survey of non-prioritized belief revision. Erkenntnis, 50(2-3):413–427, 1999.

- Types
  1. Revise only with credible triggers
  2. Delete elements from belief base or the trigger
  3. Delete elements from belief base or from closure of trigger
  4. Extend with trigger and then delete inconsistencies
  5. Decide which part $f(\alpha)$ to delete from trigger

Requirement of Finite Belief Sets

- CS cannot handle infinite belief sets
- Objects (data base, knowledge base, ontology etc.) are finite or finitely representable

- Three possible approaches
  1. Change operators for **finitely generated belief sets** $Cn(X)$ with X finite (see textbook of Hansson)
  2. Change operators for finite **belief bases**
     Belief base = not necessarily closed subset of $\mathcal{L}$
  3. Change operators for models of finite Belief Bases

Syntax-sensitive Belief Base Revision

- Hansson's Approach: use syntax sensitivity in order to represent additional justification information

Example

- $B_1 = \{p, q\}$
  Belief in p and q with independent justifications for p and q
- $B_2 = \{p \wedge q\}$
  Belief in p and q but with common justification for p and q
- $B_1 \equiv B_2$

- $B_1 \div p$ may reasonably contain q
- $B_2 \div p$ leads to $\emptyset$

Syntax-sensitive Belief Base Revision

- Similar constructions and postulates as in AGM
- Main difference: expansion now reads as $B + \alpha = B \cup \{\alpha\}$
- Additional phenomena and revision operators due to handling of inconsistency
  - First prevent inconsistency then add trigger
    $B *_{internal} \alpha = (B \div \neg\alpha) + \alpha$ (as in AGM)
  - First add trigger then handle inconsistency
    $B *_{external} \alpha = (B + \alpha) \div \bot$ (New)

Semantical Belief-Base Revision

- Semantical belief-revision demands syntax insensitiviy in both arguments: trigger and also the belief base

- In this scenario: belief bases = **knowledge bases**

Schema for semantical belief revision

$$B * \alpha = \mathit{FinRepr}(\mathit{Mod}(B) *_{\mathit{sem}} \mathit{Mod}(\alpha))$$

- $\mathit{Mod}(X)$ = Models of X
- $*_{\mathit{sem}}$ a semantical revision operator operating on pairs of sets of models
- $\mathit{FinRep}(M)$ = Formulate or finite set of formulae that hold in all models in M

Approach 1 to Semantical Revision: Generalization

- Generalize (weaken) your belief base B' minimally s.t. enlarged set of models G_i intersects with Models of trigger

- Dalal's approach
  - Defined for propositional logic models • •
  - G_i = models with Hamming distance $\leq i$ to models in $Mod(B)$

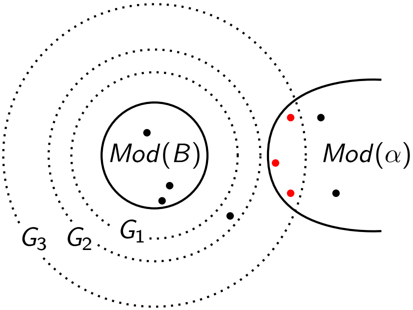

Lit: M. Dalal. Investigations into a theory of knowledge base revision: preliminary report. In AAAI-88, pages 475–479, 1988.

Approach 1 to Semantical Revision: Generalization

- Generalize (weaken) your belief base B' minimally s.t. enlarged set of models G_i intersects with Models of trigger

- Groves's approach: spheres
  - Defined on possible worlds • •
  - Possible world = maximally consistent set w.r.t. logic $(\mathcal{L}, Cn)$
  - G_i = sphere = set of possible worlds

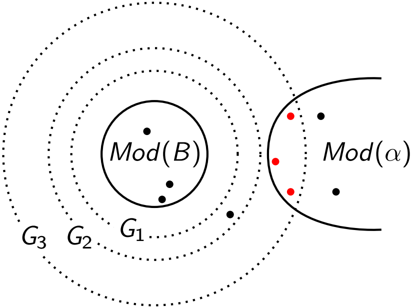

- **Note:** Maximal consistent sets correspond to models
- "Semantics" also possible in logics defined by $(\mathcal{L}, Cn)$

Lit: A. Grove. Two modellings for theory change. Journal of Philosophical Logic, 17:157–170, 1988.

Approach 2 to Semantical Revision: Minimal distance

- Dual but more general approach to generalization: minimality
- Find trigger models with "minimal distance" to $Mod(B)$
  $B * \alpha = FinRep\big(Min_{\leq_{Mod(B)}}(Mod(\alpha))\big)$

- Various ways to specify minimal distance
  - incorporating order, cardinality, etc.

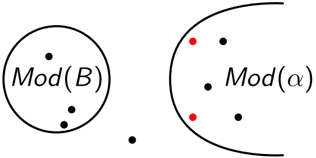

Lit: K. Satoh. Nonmonotonic reasoning by minimal belief revision. In FGCS-88, 455–462, 1988.

Lit: A. Borgida. Language features for flexible handling of exceptions in information systems. ACM Trans. Database Syst., 10(4):565–603, 1985.

Lit: A. Weber. Updating propositional formulas. In Expert Database Conf., pp. 487–500, 1986.

Lit: M. Winslett. Updating Logical Databases. Cambridge University Press, 1990.

Lit: K. D. Forbus. Introducing actions into qualitative simulation. In IJCAI-89, 1273–1278, 1988.

Complexity of Revision

- A main requirement in implementing BR operators: Feasibility of testing: $B * \alpha \models \beta$.
- Even if B is a finite propositional belief base, complexity is not really feasible
- Reason: Consistency testing is hard and you have potentially all subsets as culprit candidates
- Roughly the complexities are between NP and the second level of the polynomial hierarchy (so in PSPACE)

  Lit: T. Eiter and G. Gottlob. On the complexity of propositional knowledge base revision, updates, and counterfactuals. Artif. Intell., 57:227–270, October 1992.
- How to react to this?
  - Restrict logic to be used
  - Restrict the set of culprits: E.g., allow only culprits in ABox
  - Restrict other relevant parameters: treewidth, common variables

    Lit: A. Pfandler et al. On the parameterized complexity of belief revision. In IJCAI-15, pages 3149–3155, 2015.

Update vs. Revision

- Early CS work related to BR in Database Theory

  Lit: A. M. Keller and M. Winslett. On the use of an extended relational model to handle changing incomplete information. IEEE Transactions on Software Engineering, 11(7):620–633, 1985.
- Problem: Preserve integrity constraints when DB is **updated**
- Two main differences to BR
  - In DB : Not a theory to update but a structure
  - Update operators $\diamond$ fulfill different postulates
- Reason is: different conflict diagnostics
  - Revision: Conflict caused by false information
  - Update: Conflict caused by outdated information
  - Side note: In ontology change even a third diagnostics is possible: different terminology

Lit: H. Katsuno and A. Mendelzon. On the difference between updating a knowledge base and revising it. In KR-91, pages 387–394,1991.

Example (Winslett 1988)

- Input belief set: There is either a book on the table or a magazine $Cn(\alpha \leftrightarrow \neg\beta))$
- Trigger information: A book is put on the table α
- Apply revision operator fulfilling Postulates (R3) and (R4)

  (R3): $K * \alpha \subseteq K + \alpha$

  (R4): If $\neg\alpha \notin K$, then $K + \alpha \subseteq K * \alpha$. (Vacuity)
- Output belief set: There is a book on the table and no magazine.

  $$Cn(\alpha \leftrightarrow \neg\beta) \cup \{\alpha\}) = Cn(\alpha \wedge \neg\beta)$$

- Alternative postulate instead of vacuity

  If $\alpha \in K$, then $K \diamond \alpha = K$

Lit: M. Winslett. Reasoning about action using a possible models approach. In Proc. of the 7th National Conference on Artificial Intelligence (AAAI-88), pp. 89–93, 1988.

Further Requirements

- Trigger is by itself a belief base: **Multiple Belief Revision**
- There is no a single trigger, but a whole sequence: **Iterated revision**
- Learning ontologies: need non-amnestic (dynamic) iterated belief revision (connections to inductive learning)
- Need different logics (not fulfilling, e.g., Deduction property): Revision for ontologies in DLs
- Need to revise mappings

Exercise 7

Exercise 7.1 (2 Points)

Show that postulates (R1)–(R5) entail the following fact:
If $\alpha \in K$, then $K * \alpha = K$.

Exercise 7.2 (2 Points)

Show that $*$ is not commutative, i.e., there are K, α, β such that:

$$(K * \alpha) * \beta \neq (K * \beta) * \alpha$$

Exercise 7.3 (2 Points)

Show that Postulates (R1)–(R8) entail the following fact:
$K * \alpha = K * \beta$ iff $\alpha \in K * \beta$ and $\beta \in K * \alpha$

Exercise 7.4 (6 Points)

Show the following refined version of the theorem for the Levi-Identity:

If $*$ is defined by the Levi identity $K * \alpha = (K \div \neg\alpha) + \alpha$, then it fulfills Postulates (R*1)–(R*6) if $+$ fulfills Postulates (E1)-(E6) and $\div$ fulfills postulates (C1)–(C4) and (C6).

Exercise 7.5 (3 Points)

Calculate the following remainder sets:

1. $\{p, q, r\} \perp p \wedge q$
2. $\{q\} \perp p \wedge q$
3. $\emptyset \perp p \wedge q$