Nonmonotonic Reasoning

James Delgrande

Simon Fraser University

jim@cs.sfu.ca

Overview

1. Introduction
2. Closed-World Reasoning
3. Default Logic
4. Circumscription
5. Nonmonotonic Inference Relations
6. Other Issues

Introduction

Classical Logic and KR

As Robert Moore observed, classical logic is terrific for representing *incomplete* information. For example:

- $\exists x\, Loves(mary, x)$. *But who?*
  $\forall x\, Duck(x) \supset Bird(x)$. *What is the set of ducks?*
  $On(A, B) \vee On(A, table)$. *But which?*
  $\neg AtSchool(mary)$ *But where is she?*

Classical Logic and KR

As Robert Moore observed, classical logic is terrific for representing *incomplete* information. For example:

- $\exists x\, Loves(mary, x)$. *But who?*
  $\forall x\, Duck(x) \supset Bird(x)$. *What is the set of ducks?*
  $On(A, B) \vee On(A, table)$. *But which?*
  $\neg AtSchool(mary)$ *But where is she?*
- But FOL is limited in the forms of inference that it permits.
  - E.g. ask: Is Ralph, a raven, black?
  - To *derive* this information, we can (effectively) only reason from facts about Ralph, or general knowledge about ravens.

Classical Logic and KR

As Robert Moore observed, classical logic is terrific for representing *incomplete* information. For example:

- $\exists x\, Loves(mary, x)$. *But who?*
  $\forall x\, Duck(x) \supset Bird(x)$. *What is the set of ducks?*
  $On(A, B) \lor On(A, table)$. *But which?*
  $\neg AtSchool(mary)$ *But where is she?*
- But FOL is limited in the forms of inference that it permits.
  - E.g. ask: Is Ralph, a raven, black?
  - To *derive* this information, we can (effectively) only reason from facts about Ralph, or general knowledge about ravens.
- Commonsense knowledge and reasoning are not like this.
  - Often we want to obtain *plausible* conclusions, . . .
  - . . . that fill in our incomplete information.

Generic Statements

Observe: Most of the properties of objects or topics in everyday life hold *normally* or *usually* or *in general*.

For example:

- "Ravens are black".

  Every raven? Albinos? A raven you're told isn't black?

Generic Statements

Observe: Most of the properties of objects or topics in everyday life hold *normally* or *usually* or *in general*.

For example:

- "Ravens are black".
  Every raven? Albinos? A raven you're told isn't black?
- "Medication x is used to treat ailment y"
  Always? What if the patient is allergic to x?.

Generic Statements

Observe: Most of the properties of objects or topics in everyday life hold *normally* or *usually* or *in general*.

For example:

- "Ravens are black".
  Every raven? Albinos? A raven you're told isn't black?
- "Medication x is used to treat ailment y"
  Always? What if the patient is allergic to x?.
- "John goes for coffee at 10:00"
  Invariably? Even if he is sick or has a meeting?

Generic Statements

Observe: Most of the properties of objects or topics in everyday life hold *normally* or *usually* or *in general*.

For example:

- "Ravens are black".
  Every raven? Albinos? A raven you're told isn't black?
- "Medication x is used to treat ailment y"
  Always? What if the patient is allergic to x?.
- "John goes for coffee at 10:00"
  Invariably? Even if he is sick or has a meeting?
- And similarly for everyday topics including trees, pens, games, weddings, coffee, temporal persistence, etc.

  ☞ In fact, in commonsense domains, there are almost no interesting conditionals that hold universally.

Types of Defaults

Call a statement of the form "P's are Q's" that allows exceptions a *default*.

Types of defaults:

- Normality: *Birds normally fly*.
- Prototypicality: *The prototypical apple is red*.
- Statistical: *Most students know CPR*.
- Conventional: *Stop for a red light*.
- Persistence: Things tend to remain the same unless something causes a change.
- And many others.

Nonmonotonic Reasoning

General Goal:

Given that P's are normally Q's, want to conclude $Q(a)$ given $P(a)$, unless there is a good reason not to.

Nonmonotonic Reasoning

General Goal:

Given that P's are normally Q's, want to conclude $Q(a)$ given $P(a)$, unless there is a good reason not to.

- Classical deduction clearly isn't sufficient.
  - For example, listing exceptional conditions doesn't work:

$$\forall x\,(P(x) \land \neg Ex_1(x) \land \cdots \land \neg Ex_n(x) \supset Q(x))$$

Nonmonotonic Reasoning

General Goal:

Given that P's are normally Q's, want to conclude $Q(a)$ given $P(a)$, unless there is a good reason not to.

- Classical deduction clearly isn't sufficient.
  - For example, listing exceptional conditions doesn't work:

$$\forall x\,(P(x) \wedge \neg Ex_1(x) \wedge \cdots \wedge \neg Ex_n(x) \supset Q(x))$$

    - We can't list all exceptional conditions $Ex_1, \ldots, Ex_n$, and
    - We don't want to have to prove $\neg Ex_1(a), \ldots, \neg Ex_n(a)$ in order to conclude $Q(a)$.

Nonmonotonic Reasoning

General Goal:

Given that P's are normally Q's, want to conclude $Q(a)$ given $P(a)$, unless there is a good reason not to.

- Classical deduction clearly isn't sufficient.
  - For example, listing exceptional conditions doesn't work:

$$\forall x\,(P(x) \wedge \neg Ex_1(x) \wedge \cdots \wedge \neg Ex_n(x) \supset Q(x))$$

  - We can't list all exceptional conditions $Ex_1, \ldots, Ex_n$, and
  - We don't want to have to prove $\neg Ex_1(a), \ldots, \neg Ex_n(a)$ in order to conclude $Q(a)$.

☞ Hence need theories of how *plausible* conclusions may be drawn from uncertain, partial evidence.

Nonmonotonic Reasoning

In the notation of FOL:

Monotonic: If $\Gamma \vdash \alpha$ then $\Gamma, \Delta \vdash \alpha$.
Non-monotonic: If $\Gamma \vdash \alpha$, possibly $\Gamma, \Delta \nvdash \alpha$.

- Classical logic is monotonic.
  - For nonmonotonic reasoning we will have to alter the classical notions of validity and of proof.
- In nonmonotonic theories, an inference may depend on *lack* of information.
  - Hence a nonmonotonic inference may involve the theory as a whole.
- A rule like *P's are (normally, usually) Q's* is commonly referred to as a *default*, and the goal is to account for *default reasoning* (not to be confused with *Default Logic*, which is a specific approach).

Nonmonotonic Reasoning: Approaches

We'll cover the following approaches:

Closed World Assumption
: Formalise the assumption that a fact is false if it cannot be shown to be true.

Default Logic
: Augment classical logic with rules of the form $\frac{\alpha : \beta}{\gamma}$.
  Intuitively: If α is true and β is consistent with what's known then conclude γ.

Circumscription
: Formalise the notion that a predicate applies to as few individuals as possible.
  Then can write $\forall x(P(x) \wedge \neg Ab(x) \supset Q(x))$.

Nonmonotonic Inference Relations
: Formalise a notion of nonmonotonic inference $\alpha \mid\!\sim \beta$.
  Also expressed via a *conditional logic*, where a default $\alpha \Rightarrow \beta$ is an object in a theory.

Closed World Reasoning

Closed World Assumption

Observation: In a knowledge base, typically the number of positive facts are overwhelmed by the negative facts.

Closed World Assumption

Observation: In a knowledge base, typically the number of positive facts are overwhelmed by the negative facts.

- Thus, if we know the positive facts we don't need to store the negative facts.

Closed World Assumption

Observation: In a knowledge base, typically the number of positive facts are overwhelmed by the negative facts.

- Thus, if we know the positive facts we don't need to store the negative facts.
- E.g. in an airline database, store facts like

  $$DirectConnection(vancouver, frankfurt)$$

  but not

  $$\neg DirectConnection(vancouver, dresden).$$

Closed World Assumption

Observation: In a knowledge base, typically the number of positive facts are overwhelmed by the negative facts.

- Thus, if we know the positive facts we don't need to store the negative facts.
- E.g. in an airline database, store facts like

  $DirectConnection(vancouver, frankfurt)$

  but not

  $\neg DirectConnection(vancouver, dresden)$.

Closed-World Assumption (CWA) [Reiter, 1978]

> *If an atomic sentence is not known to be true, it can be assumed to be false.*

CWA: Formalisation

Define a new version of entailment:

$$KB \models_{cwa} \alpha \;\; \textit{iff} \;\; CWA(KB) \models \alpha, \quad \textit{where}$$

$$CWA(KB) = KB \cup \{\neg p \mid KB \not\models p \textit{ where } p \textit{ is atomic.}\}$$

CWA: Formalisation

Define a new version of entailment:

$$KB \models_{cwa} \alpha \;\; \textit{iff} \;\; CWA(KB) \models \alpha, \quad \textit{where}$$

$$CWA(KB) = KB \cup \{\neg p \mid KB \not\models p \textit{ where } p \textit{ is atomic.}\}$$

- Example: In a blocks world we might have:

$$KB = \{On(a, b, s), On(b, table, s), On(c, table, s)\}$$

With the CWA we can infer

$$\neg On(a, a, s), \; \neg On(b, a, s) \text{ and } \neg On(table, a, s).$$

CWA and DCA

With the CWA and for

$$KB = \{On(a, b, s), On(b, table, s), On(c, table, s)\},$$

we cannot infer

$$\forall x \neg On(x, a, s).$$

- Reason: There may be some (unnamed) x that is on a.

CWA and DCA

With the CWA and for

$$KB = \{On(a, b, s), On(b, table, s), On(c, table, s)\},$$

we cannot infer

$$\forall x \neg On(x, a, s).$$

- Reason: There may be some (unnamed) x that is on a.

Domain-closure assumption (DCA): Often we can assume (or we know) that the only objects are the named objects.

- In the above, this would amount to
  $$\forall x\, [Block(x) \equiv (x = a \vee x = b \vee x = c)]$$
- With the DCA we can infer
  $$\forall x \neg On(x, a, s).$$

☞ Note that we would not want to apply the DCA to s.

Query evaluation with the CWA+DCA

- With the CWA and DCA, entailment becomes easy.
- Let $\models_{cd}$ be entailment under the CWA and DCA, and let α and β be in negation normal form.
  Then
  - $KB \models_{cd} \alpha \wedge \beta$ iff $KB \models_{cd} \alpha$ and $KB \models_{cd} \beta$
  - $KB \models_{cd} \alpha \vee \beta$ iff $KB \models_{cd} \alpha$ or $KB \models_{cd} \beta$
  - $KB \models_{cd} \forall x\, \alpha$ iff $KB \models_{cd} \alpha[x/c]$ for every c in the KB.
  - $KB \models_{cd} \exists x\, \alpha$ iff $KB \models_{cd} \alpha[x/c]$ for some c in the KB.
- Reduces to $KB \models_{cd} \ell$ where ℓ is a literal.
- If atoms are stored in a table, this reduces to table lookup.
- To handle equality, need the *unique names assumption (UNA)*:
  For distinct constants c, d, assume that $(c \neq d)$.

Consistency of the CWA

- Consider where $KB \models p \vee q$ but $KB \not\models p$ and $KB \not\models q$.

Consistency of the CWA

- Consider where $KB \models p \vee q$ but $KB \not\models p$ and $KB \not\models q$.
- Then $CWA(KB) = KB \cup \{\neg p, \neg q\}$
  - But this is inconsistent!

Consistency of the CWA

- Consider where $KB \models p \vee q$ but $KB \not\models p$ and $KB \not\models q$.
- Then $CWA(KB) = KB \cup \{\neg p, \neg q\}$
  - But this is inconsistent!
- One solution: Generalised closed world assumption (GCWA).

$$
\begin{aligned}
GCWA(KB) = KB \cup \{\neg p \mid & \\
& \text{if } KB \models p \vee q_1 \vee \cdots \vee q_n \\
& \qquad \text{then } KB \models q_1 \vee \cdots \vee q_n \; \}
\end{aligned}
$$

- Obtain:
  - $GCWA(KB)$ is consistent if KB is.
  - If $GCWA(KB) \models \alpha$ then $CWA(KB) \models \alpha$.

Complexity

- Propositional CWA deduction can be done with $O(\log m)$ calls to an NP oracle.
  - Hence the problem is in $P^{NP[O(\log n)]}$.
- Propositional GCWA deduction can be done with $O(\log m)$ calls to Σ_2^P oracle.
  - Hence the problem is in $P^{\Sigma_2^P[O(\log n)]}$.
- Reference: [Eiter and Gottlob, 1993].

CWA: Concluding Points

- FO CWA is noncomputable (since it appeals to $\not\models$).
- We have the theorem:

  If KB is Horn and consistent, then CWA(KB) is consistent.

- CWA (and DCA) rely on the syntactic form of the theory.
  - E.g. replace *On* by *Off* in the block's world example, and you get exactly the opposite assertions.
- CWA (+ DCA and UNA) is fundamental in deductive and (implicitly) relational database theory, as well as in logic programming.

Default Logic

Default Logic

Default Logic (DL) [Reiter, 1980] is probably the best-known and most studied approach to NMR.

- Reiter's intuition:

  > *Default reasoning "corresponds to the process of deriving conclusions based on patterns of inference of the form 'in the absence of information to the contrary, assume . . . ' ".*

- Informally:
  - With the CWA, negated ground atoms are added to a KB.
  - In DL, formulas are added to a KB based on what's known and not known.

Default Rules

- In classical logic, inference rules sanction the derivation of a formula based on other formulas that have been derived.
- Defaults in DL are like domain-specific inference rules, but with an added consistency condition.
- E.g.: "University students are normally adults" can be expressed by

$$\frac{UnivSt(x) \ : \ Adult(x)}{Adult(x)}$$

First approximation: If $UnivSt(c)$ is true for ground term c and $Adult(c)$ is consistent, then $Adult(c)$ can be derived "by default".

Default Rules and Extensions

Problem: How to characterize default consequences?

Consider a default rule $\frac{\alpha : \beta}{\gamma}$.

- Intuition:

  γ can be derived if α has been derived and β is consistent.

Default Rules and Extensions

Problem: How to characterize default consequences?

Consider a default rule $\frac{\alpha : \beta}{\gamma}$.

- Intuition:

  γ can be derived if α has been derived and β is consistent.

- Question:

  Consistent with what?

Default Rules and Extensions

Problem: How to characterize default consequences?

Consider a default rule $\frac{\alpha : \beta}{\gamma}$.

- Intuition:

  γ can be derived if α has been derived and β is consistent.
- Question:

  Consistent with what?
- Reiter's answer:

  Consistent with the full set of formulas that can be justified by classical reasoning and application of default rules.
- Such a set of sentences is called an *extension*.

Basic Definitions

A *default* is an expression of the form

$$\frac{\alpha : \beta_1, \ldots, \beta_n}{\gamma}$$

where α, β_i, γ are formulas of first order (or propositional) logic.

- α is the *prerequisite*
- $\beta_1, \ldots, \beta_n$ are *justifications*
  - We'll stick with $n = 1$.
- γ is the *consequent*.

A *default theory* is a pair (W, D) where W is a set of sentences of first order (or propositional) logic and D is a set of defaults.

More Basic Definitions

A default is *closed* if it contains no free variables among its formulas; otherwise it is *open*.

- An open default will stand for its set of ground instances.
- So we can assume that we are (effectively) dealing with a closed default theory.

Default Extensions

A default theory (W, D) induces a set of *extensions*, where an extension is a "reasonable" set of beliefs based on (W, D).

Reiter lists the following desirable properties of any extension E:

Default Extensions

A default theory (W, D) induces a set of *extensions*, where an extension is a "reasonable" set of beliefs based on (W, D).

Reiter lists the following desirable properties of any extension E:

1. Since W is certain, we require $W \subseteq E$.

Default Extensions

A default theory (W, D) induces a set of *extensions*, where an extension is a "reasonable" set of beliefs based on (W, D).

Reiter lists the following desirable properties of any extension E:

1. Since W is certain, we require $W \subseteq E$.
2. E is deductively closed, that is, $E = Cn(E)$.

Default Extensions

A default theory (W, D) induces a set of *extensions*, where an extension is a "reasonable" set of beliefs based on (W, D).

Reiter lists the following desirable properties of any extension E:

1. Since W is certain, we require $W \subseteq E$.
2. E is deductively closed, that is, $E = Cn(E)$.
3. A maximal set of defaults is applied.

   So for $\frac{\alpha : \beta}{\gamma} \in D$, if $\alpha \in E$ and $\neg\beta \notin E$ then $\gamma \in E$.

Default Extensions

A default theory (W, D) induces a set of *extensions*, where an extension is a "reasonable" set of beliefs based on (W, D).

Reiter lists the following desirable properties of any extension E:

1. Since W is certain, we require $W \subseteq E$.
2. E is deductively closed, that is, $E = Cn(E)$.
3. A maximal set of defaults is applied.

   So for $\frac{\alpha : \beta}{\gamma} \in D$, if $\alpha \in E$ and $\neg\beta \notin E$ then $\gamma \in E$.

Unfortunately minimality wrt 1–3 doesn't give a satisfactory definition of an extension.

- E.g. for $(\emptyset, \{\frac{: \alpha}{\alpha}\})$, $E = Cn(\neg\alpha)$ satisfies 1–3.

Default Extensions: Definition

Reiter's definition:

Let (W, D) be a default theory. The operator Γ assigns to every set S of formulas the smallest set S' of formulas such that:

1. $W \subseteq S'$
2. $S' = Cn(S')$
3. If $\frac{\alpha : \beta}{\gamma} \in D$ and $\alpha \in S'$ and $\neg\beta \notin S$ then $\gamma \in S'$.

A set E is an extension for (W, D) iff $\Gamma(E) = E$.

☞ That is, E is a *fixed point* of Γ.

1. guarantees that the given facts are in the extension.
2. states that beliefs are deductively closed.
3. has the effect that as many defaults as possible (with respect to the extension) are applied.

Another Definition

Reiter gives an equivalent "pseudo-iterative" definition of an extension:

For default theory (W, D) define:

$$\begin{aligned} E_0 &= W \\ E_{i+1} &= Cn(E_i) \cup \left\{ \gamma \mid \frac{\alpha : \beta}{\gamma} \in D \text{ and } \alpha \in E_i \text{ and } \neg\beta \notin E \right\} \\ &\quad \text{for } i \geq 0 \end{aligned}$$

Then E is an extension for (W, D) iff $E = \bigcup_{i=0}^{\infty} E_i$.

- With this definition, it is straightforward to *verify* whether a given set of formulas constitutes an extension.

Example

Notation: For extension E of (W, D), let

$$\Delta_E = \{\gamma \mid \frac{\alpha : \beta}{\gamma} \in D, \alpha \in E, \neg\beta \notin E\}$$

- Consider:

$$W = \{Bird(tweety), Bird(opus), \neg Fly(opus)\}$$

$$D = \left\{ \frac{Bird(x) : Fly(x)}{Fly(x)} \right\}$$

- One extension E where $\Delta_E = \{Fly(tweety)\}$.

Another Example

- Consider:

$$W = \{Republican(dick), Quaker(dick)\}$$

$$D = \left\{ \frac{Republican(x) : \neg Pacifist(x)}{\neg Pacifist(x)}, \frac{Quaker(x) : Pacifist(x)}{Pacifist(x)} \right\}$$

- Two extensions:

$$\begin{aligned} \Delta_{E_1} &= \{\neg Pacifist(dick)\} \\ \Delta_{E_2} &= \{\ \ Pacifist(dick)\} \end{aligned}$$

Another Example

- Consider:

$$W = \{Republican(dick), Quaker(dick)\}$$

$$D = \left\{ \frac{Republican(x) : \neg Pacifist(x)}{\neg Pacifist(x)}, \frac{Quaker(x) : Pacifist(x)}{Pacifist(x)} \right\}$$

- Two extensions:

$$\Delta_{E_1} = \{\neg Pacifist(dick)\}$$
$$\Delta_{E_2} = \{\ Pacifist(dick)\}$$

- What to believe?

  First approximation:

  Credulous: Choose an extension arbitrarily
  Skeptical: Intersect the extensions.

Yet Another Example

- Consider:

$$W = \{Bat(tweety) \vee Bird(tweety)\}$$

$$D = \left\{ \frac{Bat(x) : Fly(x)}{Fly(x)}, \frac{Bird(x) : Fly(x)}{Fly(x)} \right\}$$

- One extensions $E = Cn(W)$.
- So, no reasoning by cases.

More Examples

- $W = \emptyset, \; D = \left\{ \dfrac{\top \; : \; a}{\neg a} \right\}$

  No extensions.

More Examples

- $W = \emptyset,\ D = \left\{ \dfrac{\top\ :\ a}{\neg a} \right\}$

  No extensions.

- "Closed world assumption" for predicate P:
  - Represent as $\frac{:\neg P(x)}{\neg P(x)}$.
  - If $W = \{P(a) \vee P(b)\}$,
    - DL yields 2 extensions;
    - CWA yields inconsistency.

Normal Default Theories

- Most often, default rules have the same justification and consequent.
- A rule of the form $\frac{\alpha : \beta}{\beta}$ is a *normal default rule*.
- Normal default theories have nice properties.

Normal Default Theories

- Most often, default rules have the same justification and consequent.
- A rule of the form $\frac{\alpha : \beta}{\beta}$ is a *normal default rule*.
- Normal default theories have nice properties.

  Let (W, D) be a normal default theory. Then:
  - (W, D) has an extension.

Normal Default Theories

- Most often, default rules have the same justification and consequent.
- A rule of the form $\frac{\alpha : \beta}{\beta}$ is a *normal default rule*.
- Normal default theories have nice properties.

  Let (W, D) be a normal default theory. Then:
  - (W, D) has an extension.
  - If (W, D) has extensions E_1, E_2 and $E_1 \neq E_2$, then $E_1 \cup E_2$ is inconsistent.

Normal Default Theories

- Most often, default rules have the same justification and consequent.
- A rule of the form $\frac{\alpha : \beta}{\beta}$ is a *normal default rule*.
- Normal default theories have nice properties.

  Let (W, D) be a normal default theory. Then:
  - (W, D) has an extension.
  - If (W, D) has extensions E_1, E_2 and $E_1 \neq E_2$, then $E_1 \cup E_2$ is inconsistent.
  - *Semi-monotonicity*: If E is an extension of (W, D) and D' is a set of normal defaults, then $(W, D \cup D')$ has an extension E' where $E \subseteq E'$.

Normal Default Theories

- Most often, default rules have the same justification and consequent.
- A rule of the form $\frac{\alpha : \beta}{\beta}$ is a *normal default rule*.
- Normal default theories have nice properties.

  Let (W, D) be a normal default theory. Then:
  - (W, D) has an extension.
  - If (W, D) has extensions E_1, E_2 and $E_1 \neq E_2$, then $E_1 \cup E_2$ is inconsistent.
  - *Semi-monotonicity*: If E is an extension of (W, D) and D' is a set of normal defaults, then $(W, D \cup D')$ has an extension E' where $E \subseteq E'$.
  - Also an extension can be specified iteratively.

Semi-Normal Defaults

So why not just stick with normal default theories?

Semi-Normal Defaults

So why not just stick with normal default theories?

- Problem:
    - typically university students are adults: $\frac{S(x) : A(x)}{A(x)}$
    - typically adults are employed: $\frac{A(x) : E(x)}{E(x)}$
    - typically university students are not employed: $\frac{S(x) : \neg E(x)}{\neg E(x)}$
- For $W = \{S(sue)\}$, get 2 extensions, one with $E(sue)$ and one with $\neg E(sue)$.
    - Want just the second extension, with $\neg E(sue)$.

Semi-Normal Defaults

So why not just stick with normal default theories?

- Problem:
  - typically university students are adults: $\frac{S(x) : A(x)}{A(x)}$
  - typically adults are employed: $\frac{A(x) : E(x)}{E(x)}$
  - typically university students are not employed: $\frac{S(x) : \neg E(x)}{\neg E(x)}$
- For $W = \{S(sue)\}$, get 2 extensions, one with $E(sue)$ and one with $\neg E(sue)$.
  - Want just the second extension, with $\neg E(sue)$.
- Solution: block transitivity with rule:

$$\frac{A(x) \; : \; \neg S(x) \wedge E(x)}{E(x)}$$

Semi-Normal Defaults

- A default of the form $\frac{\alpha : \beta \wedge \gamma}{\beta}$ is *semi-normal*.
- Semi-normal defaults are required for interacting defaults, as in the last example.
- For semi-normal defaults:
  - We may not have an extension
  - We lack semi-monotonicity
  - The proof theory appears considerably more complex

DL and Other Approaches

DL and Autoepistemic Logic:

- Autoepistemic Logic (AEL) [Moore, 1985] was developed as an account of how an ideal reasoner may form beliefs, reasoning about its beliefs and non-beliefs.
- Uses a *modal* approach: $B\alpha$ read as "α is believed".
- Belief set E of an agent should satisfy 3 properties:
  1. $Cn(E) = E$.
  2. If $\alpha \in E$ then $B\alpha \in E$.
  3. If $\alpha \notin E$ then $\neg B\alpha \in E$.

Autoepistemic Logic

- Leads to the notion of *(grounded) stable expansions*.
- E is a grounded stable extension of KB iff E is a minimal set wrt nonmodal formulas such that

$$\gamma \in E \quad \text{iff} \quad KB \cup \Delta \models \gamma \text{ where}$$
$$\Delta = \{B\alpha \mid \alpha \in E\} \cup \{\neg B\alpha \mid \alpha \notin E\}.$$

  - So, another fixed-point definition.

Autoepistemic Logic

- Leads to the notion of *(grounded) stable expansions*.
- E is a grounded stable extension of KB iff E is a minimal set wrt nonmodal formulas such that

$$\gamma \in E \;\text{ iff }\; KB \cup \Delta \models \gamma \text{ where}$$
$$\Delta = \{B\alpha \mid \alpha \in E\} \cup \{\neg B\alpha \mid \alpha \notin E\}.$$

  - So, another fixed-point definition.

- Shown in [Denecker *et al.*, 2003] to have deep connections to DL, wherein expansions correspond to extensions.
  - Roughly: AEL and DL can be generalised to sets of approaches, with a 1-1 correspondence between approaches.

DL and Other Approaches

DL and Answer Set Programming (ASP):

- Reference: [Gelfond, 2008]
- A (normal) answer set program is a set of rules of the form:

$$l_0 \leftarrow l_1, \ldots, l_n, \textit{not } l_{n+1}, \ldots, \textit{not } l_m$$

  where the l_i's are literals.
- An *answer set* for a program is (roughly) a minimal set of literals such that for every rule, if the positive part of the body is in the set and the negative part isn't, then the head is.
- ASP shows great promise in applications, and implementations are competitive with the best SAT solvers.

Answer Set Programming

- Let (W, D) be a a default theory where
  - each element of W is a ground fact and
  - each rule of D is of the form

$$\frac{l_1 \wedge \cdots \wedge l_n \ : \ l_{n+1}, \ldots, l_m}{l_0}$$

  where l_i, $0 \leq i \leq m$, is a literal.

- There is an AS program Π where rules as above are mapped to

$$l_0 \leftarrow l_1, \ldots, l_n, not\,\overline{l}_{n+1}, \ldots, not\,\overline{l}_m$$

  and $l \in W$ maps to $l \leftarrow$.

- Then ([Gelfond and Lifschitz, 1991])
  - For AS X of Π, $Cn(X)$ is an extension of (W, D)
  - For extension E of (W, D), the literals in E are an AS of Π.

Concluding Points

- For propositional DL:
  - Deciding extension existence is Σ_2^P-complete.
  - Deciding credulous inference is Σ_2^P-complete.
  - Deciding skeptical inference is Π_2^P-complete.
  - The latter 2 results hold for normal default theories.
  - Reference: [Gottlob, 1992].
- Previously, meaningful practical applications of DL have been lacking; this is changing with the advent of ASP.

Circumscription

Circumscription

Introduced by John McCarthy, with many results obtained by Vladimir Lifschitz.

- See [McCarthy, 1980], [McCarthy, 1986], [Lifschitz, 1994].

General Idea: Want to be able to say that the extension of a predicate is as small as possible.

Circumscription

Introduced by John McCarthy, with many results obtained by Vladimir Lifschitz.

- See [McCarthy, 1980], [McCarthy, 1986], [Lifschitz, 1994].

General Idea: Want to be able to say that the extension of a predicate is as small as possible.

- Then, for "university students are normally adults" can write:
  $$\forall x(S(x) \land \neg Ab(x) \supset A(x))$$
- Circumscribing Ab yields that Ab applies to as few individuals as possible.
- If we have $S(sue)$ and circumscribing Ab yields $\neg Ab(sue)$ we can conclude $A(sue)$.
- Circumscription can be specified semantically or syntactically.
  ☞ We'll focus on the semantic side.

Circumscription: Intuitions

- In classical logic, all models of a theory have the same status.
- In circumscribing P, we *prefer* those models of P with smaller extensions.

Circumscription: Intuitions

- In classical logic, all models of a theory have the same status.
- In circumscribing P, we *prefer* those models of P with smaller extensions.
  - E.g., if we knew only that
    $$\exists x P(x)$$
    we would expect the circumscription of P to entail
    $$\exists x \forall y (P(y) \equiv (x = y)).$$

Circumscription: Intuitions

- In classical logic, all models of a theory have the same status.
- In circumscribing P, we *prefer* those models of P with smaller extensions.
  - E.g., if we knew only that
    $$\exists x P(x)$$
    we would expect the circumscription of P to entail
    $$\exists x \forall y (P(y) \equiv (x = y)).$$
  - If we knew only that
    $$P(a) \vee P(b)$$
    we would expect the circumscription of P to entail
    $$(\forall x P(x) \equiv x = a) \ \vee \ (\forall x P(x) \equiv x = b).$$

Minimal Entailment

Let $\mathbf{P}$ be a set of predicates.
Let $\mathcal{I}_1 = (D_1, I_1)$, $\mathcal{I}_2 = (D_2, I_2)$ be two interpretations.

Define $\mathcal{I}_1 \leq_{\mathbf{P}} \mathcal{I}_2$, read $\mathcal{I}_1$ is *at least as preferred as* $\mathcal{I}_2$, if

1. $D_1 = D_2$,
2. $I_1[X] = I_2[X]$ for every predicate symbol X not in $\mathbf{P}$.
3. $I_1[P] \subseteq I_2[P]$ for every $P \in \mathbf{P}$.

Minimal Entailment

Let $\mathbf{P}$ be a set of predicates.
Let $\mathcal{I}_1 = (D_1, I_1)$, $\mathcal{I}_2 = (D_2, I_2)$ be two interpretations.

Define $\mathcal{I}_1 \leq_{\mathbf{P}} \mathcal{I}_2$, read $\mathcal{I}_1$ is *at least as preferred as* $\mathcal{I}_2$, if

1. $D_1 = D_2$,
2. $I_1[X] = I_2[X]$ for every predicate symbol X not in $\mathbf{P}$.
3. $I_1[P] \subseteq I_2[P]$ for every $P \in \mathbf{P}$.

$\mathcal{I}_1 <_{\mathbf{P}} \mathcal{I}_2$ iff: $\mathcal{I}_1 \leq_{\mathbf{P}} \mathcal{I}_2$ but not $\mathcal{I}_2 \leq_{\mathbf{P}} \mathcal{I}_1$.

Define a new version of entailment $\models_{\leq}$ by:

$KB \models_{\leq_{\mathbf{P}}} \alpha$ iff for every $\mathcal{I}$ where $\mathcal{I} \models KB$ and
$\nexists \mathcal{I}'$ s.t. $\mathcal{I}' <_{\mathbf{P}} \mathcal{I}$ and $\mathcal{I}' \models KB$,
then $\mathcal{I} \models \alpha$.

Examples

- $KB = \{ P(a) \wedge P(b) \}$

  $KB \models_{\leq_P} \forall x (P(x) \equiv (x = a \vee x = b))$

Examples

- $KB = \{ P(a) \wedge P(b) \}$

  $KB \models_{\leq_P} \forall x(P(x) \equiv (x = a \vee x = b))$

- $KB = \{ \forall x(Q(x) \supset P(x)) \}$

  $KB \models_{\leq_P} \forall x(Q(x) \equiv P(x))$

Problematic Example 1

$$
\begin{aligned}
KB \;=\; \{\, & \forall x(Bird(x) \wedge \neg Ab(x) \supset Fly(x)), \\
& \forall x(Penguin(x) \supset \neg Fly(x)), \\
& \forall x(Penguin(x) \supset Bird(x)) \,\}
\end{aligned}
$$

- Note that $KB \models \forall x(Penguin(x) \supset Ab(x))$
- Get:

$$KB \models_{\leq_{Ab}} \forall x(Ab(x) \equiv [Penguin(x) \vee (Bird(x) \wedge \neg Fly(x))])$$

- Can't conclude *Fly* by default for an individual.

Problematic Example 1: Solution

Intuition: Allow some predicates to vary (such as *Fly*) in minimising a predicate (such as *Ab*).

Problematic Example 1: Solution

Intuition: Allow some predicates to vary (such as *Fly*) in minimising a predicate (such as *Ab*).

Modify the definition:

Let **P**, **Q** be sets of predicates.

For $\mathcal{I}_1 = (D_1, I_1)$, $\mathcal{I}_2 = (D_2, I_2)$, define $\mathcal{I}_1 \leq_{\mathbf{P},\mathbf{Q}} \mathcal{I}_2$, if

1. $D_1 = D_2$,
2. $I_1[X] = I_2[X]$ for every predicate symbol X not in $\mathbf{P} \cup \mathbf{Q}$.
3. $I_1[P] \subseteq I_2[P]$ for every $P \in \mathbf{P}$.

Examples

- Now minimizing *Ab* and letting *Fly* vary gives

$$\forall x(Ab(x) \equiv Penguin(x))$$

  So, the only abnormal things are penguins.

Examples

- Now minimizing *Ab* and letting *Fly* vary gives

$$\forall x(Ab(x) \equiv Penguin(x))$$

  So, the only abnormal things are penguins.

- $$KB = \{\, \forall x(S(x) \land \neg Ab(x) \supset A(x)),\ S(sue),\ S(yi),\ \neg A(sue) \lor \neg A(yi)\,\}$$

  $KB \models_{\leq_{Ab}} A(sue) \lor A(yi)$.

  ☞ We don't get this result in the simpler formulation.

Problematic Example 2

$$
\begin{aligned}
KB \;=\; \{ & \forall x(Bird(x) \land \neg Ab_1(x) \supset Fly(x)), \\
& \forall x(Penguin(x) \land \neg Ab_2(x) \supset \neg Fly(x)), \\
& \forall x(Penguin(x) \supset Bird(x)), \\
& Penguin(opus) \; \}
\end{aligned}
$$

- Circumscribing with $\mathbf{P} = \{Ab_1, Ab_2\}$, $\mathbf{Q} = \{Fly\}$ we obtain $Ab_1(opus) \lor Ab_2(opus)$, and not $\neg Fly(opus)$.
  - ☞ So specificity is not handled.
- Solution [Lifschitz, 1985]: Prioritized circumscription.
  - Give a priority order for circumscription.
  - In the above, we would circumscribe Ab_2, then Ab_1.

Syntactic Characterisation

Circumscription can also be described syntactically.

- I.e. given a sentence KB, the circumscription produces a logically stronger sentence KB^*.
- Done in terms of a formula of second-order logic.
- We will just consider the basic case of circumscribing a single predicate.

Circumscription Schema

Notation: Let P and Q be predicates of the same arity.

$P \equiv Q$ abbreviates $\forall \bar{x}(P(\bar{x}) \equiv Q(\bar{x}))$.
$P \leq Q$ abbreviates $\forall \bar{x}(P(\bar{x}) \supset Q(\bar{x}))$.
$P < Q$ abbreviates $(P \leq Q) \wedge \neg(Q \leq P)$.

Circumscription Schema

Notation: Let P and Q be predicates of the same arity.

$$P \equiv Q \text{ abbreviates } \forall \bar{x}(P(\bar{x}) \equiv Q(\bar{x})).$$
$$P \leq Q \text{ abbreviates } \forall \bar{x}(P(\bar{x}) \supset Q(\bar{x})).$$
$$P < Q \text{ abbreviates } (P \leq Q) \land \neg(Q \leq P).$$

Let $KB(P)$ be a formula containing P, and let p be a predicate variable of same arity as P.

The *circumscription* of P in $KB(P)$ is the second-order formula:

$$KB(P) \land \neg\exists p(KB(p) \land p < P).$$

where $KB(p)$ is the result of replacing every occurrence of P in KB with p.

Circumscription Schema

For the circumscription of P in $KB(P)$

$$KB(P) \wedge \neg \exists p(KB(p) \wedge p < P),$$

we have that:

- $KB(P)$ guarantees that the circumscription has all the properties of the original formula;
- the conjunct $\neg \exists p(KB(p) \wedge p < P)$ says that there is no predicate p such that
  - p satisfies what P does, and
  - the extension of p is a proper subset of that of P.

I.e. P is minimal with respect to $KB(P)$.

Circumscription Schema: Notes

- The syntactic approach can be shown to capture the same results as minimal models.
- The definition can be extended to deal with sets of predicates, varying predicates, and priorities among predicates.
- Issue: Determining cases where the schema can be expressed as a formula of first-order logic.

Concluding Points

- The deduction problem for propositional circumscription,

  viz. does $Circ(A; \mathbf{P}; \mathbf{Q}) \models \alpha$?

  is Π_2^P-complete [Eiter and Gottlob, 1993].

Concluding Points

- The deduction problem for propositional circumscription,

  viz. does $Circ(A; \mathbf{P}; \mathbf{Q}) \models \alpha$?

  is Π_2^P-complete [Eiter and Gottlob, 1993].
- It is not clear that abnormality theories are adequate for dealing with defaults per se.

Concluding Points

- The deduction problem for propositional circumscription,

  viz. does $Circ(A; \mathbf{P}; \mathbf{Q}) \models \alpha$?

  is Π_2^P-complete [Eiter and Gottlob, 1993].
- It is not clear that abnormality theories are adequate for dealing with defaults per se.
- However, circumscription has found numerous applications, in areas such as
  - reasoning about action (and dealing with persistence) and
  - diagnosis.

Concluding Points

- Circumscription (like Default Logic) isn't a logic of defaults per se, but rather provide a mechanism wherein default reasoning may be encoded.
  - E.g. for variable circumscription, need to decide what predicates to allow to vary.
  - Hard to ensure that the "right" conclusions are obtained in all circumstances.

Defaults as Objects:

Nonmonotonic Inference Relations/ Conditional Logics

Motivation: In DL and circumscription, default theories have to be hand-coded.

- This suggests studying nonmonotonicity as an abstract phenomenon.

Introduction

Motivation: In DL and circumscription, default theories have to be hand-coded.

- This suggests studying nonmonotonicity as an abstract phenomenon.

Two broad approaches:

Nonmonotonic Inference Relations Analogously to classical inference, $\alpha \vdash \beta$, consider properties of a nonmonotonic inference relation $\alpha \mid\!\sim \beta$.

Conditional Logics Analogously to material implication, $\alpha \supset \beta$, consider properties of a default conditional $\alpha \Rightarrow \beta$ added to classical logic.

☞ These approaches basically coincide; we'll focus on the first.

Nonmonotonic Inference Relations [Kraus *et al.*, 1990]

Intuition:

- In classical logic, $\alpha \models \beta$ just when β is true in all models of α.
- The inference relation $\alpha \mid\!\sim \beta$ expresses that β is true in all *preferred* models of α.

Nonmonotonic Inference Relations [Kraus *et al.*, 1990]

Intuition:

- In classical logic, $\alpha \models \beta$ just when β is true in all models of α.
- The inference relation $\alpha \mid\sim \beta$ expresses that β is true in all *preferred* models of α.

Obvious question:

- How do we specify the notion of "preferred model"?

Nonmonotonic Inference Relations [Kraus *et al.*, 1990]

Intuition:

- In classical logic, $\alpha \models \beta$ just when β is true in all models of α.
- The inference relation $\alpha \mid\!\sim \beta$ expresses that β is true in all *preferred* models of α.

Obvious question:

- How do we specify the notion of "preferred model"?

Answer:

- This is given by a partial preorder over interpretations.
- Then $\alpha \mid\!\sim \beta$ just if β is true in the *minimal* models of α.

NMIR: Semantics

- L is the language of PC, with atomic sentences $\mathbf{P} = \{a, b, c, \ldots\}$ and the usual connectives.
- Ω is the set of interpretations of L.
  - Define $\|\alpha\| = \{w \in \Omega \mid w \models \alpha\}$.
- $\preceq$ is a *preference relation* on interpretations of L.
  - $\preceq$ is reflexive and transitive.
- Define

$$min(\|\alpha\|) = \left\{w \in \|\alpha\| \mid \not\exists w' \in \Omega \text{ s.t. } w' \prec w \text{ and } w' \models \alpha\right\}.$$

- Then

$$\alpha \mid\!\sim \beta \quad \text{just if} \quad min(\|\alpha\|) \subseteq \|\beta\|.$$

Proof Theory

Consider the following properties of NMIRs:

REF	$\alpha \mid\!\sim \alpha$.
LLE	If $\models \alpha \equiv \beta$ and $\alpha \mid\!\sim \gamma$ then $\beta \mid\!\sim \gamma$.
RW	If $\models \beta \supset \gamma$ and $\alpha \mid\!\sim \beta$ then $\alpha \mid\!\sim \gamma$.
AND	If $\alpha \mid\!\sim \beta$ and $\alpha \mid\!\sim \gamma$ then $\alpha \mid\!\sim \beta \wedge \gamma$.
OR	If $\alpha \mid\!\sim \gamma$ and $\beta \mid\!\sim \gamma$ then $\alpha \vee \beta \mid\!\sim \gamma$.
CM	If $\alpha \mid\!\sim \beta$ and $\alpha \mid\!\sim \gamma$ then $\alpha \wedge \beta \mid\!\sim \gamma$.

Obtain:

$\mid\!\sim$ is a preferential inference relation iff it satisfies REF– CM.

Aside: In a conditional logic, we would have axioms like:

$$(\alpha \Rightarrow \beta \wedge \alpha \Rightarrow \gamma) \supset \alpha \wedge \beta \Rightarrow \gamma.$$

in place of CM.

Examples

- Let $\Gamma = \{B \mid\!\sim F, \quad B \mid\!\sim W, \quad P \mid\!\sim B, \quad P \mid\!\sim \neg F\}$
- Γ is non-trivially satisfiable.
- From Γ can infer
  - $B \wedge W \mid\!\sim F$
  - $B \wedge P \mid\!\sim \neg F$
  - $B \mid\!\sim \neg P$

Examples

- Let $\Gamma = \{B \mid\sim F, \quad B \mid\sim W, \quad P \mid\sim B, \quad P \mid\sim \neg F\}$
- Γ is non-trivially satisfiable.
- From Γ can infer
  - $B \wedge W \mid\sim F$
  - $B \wedge P \mid\sim \neg F$
  - $B \mid\sim \neg P$
- However from Γ cannot infer
  - $B \wedge Gr \mid\sim F$
  - $P \mid\sim W$
- Basically at this point, while we have a "logic of defaults" we do not have an adequate system for default inference.

☞ Can't handle *irrelevant properties* like Gr in $B \wedge Gr \mid\sim F$.

Rational Closure

- As noted, we don't actually have a nonmonotonic system.
- [Lehmann and Magidor, 1992] defines the *rational closure* of a KB
  - Roughly: Given a KB, determine the preference relation where formulas are ranked "as low as possible".
- This is done wrt a stronger system, that incorporates *rational monotonity*.

  RM If $\alpha \mid\!\sim \gamma$ and $\alpha \not\mid\!\sim \neg\beta$ then $\alpha \wedge \beta \mid\!\sim \gamma$.
- Semantically this axiom enforces a total preorder on interpretations.

Rational Closure

- Define $\beta \prec \alpha$ iff $(\alpha \vee \beta) \mid\!\sim \neg\alpha$.
- Given an understood NM theory T, the *degree* of a formula is defined by:
  1. $deg(\alpha) = 0$ iff for no β do we have $\beta \prec \alpha$.
  2. $deg(\alpha) = i$ iff $deg(\alpha)$ is not less than i and for every β such that $\beta \prec \alpha$ we have $deg(\beta) < i$.
  3. $deg(\alpha) = \infty$ iff α is not assigned a degree above.

Rational Closure

- Define $\beta \prec \alpha$ iff $(\alpha \vee \beta) \mathrel{|\!\sim} \neg\alpha$.
- Given an understood NM theory T, the *degree* of a formula is defined by:
  1. $deg(\alpha) = 0$ iff for no β do we have $\beta \prec \alpha$.
  2. $deg(\alpha) = i$ iff $deg(\alpha)$ is not less than i and for every β such that $\beta \prec \alpha$ we have $deg(\beta) < i$.
  3. $deg(\alpha) = \infty$ iff α is not assigned a degree above.
- The *rational consequence* relation wrt T is given by:

$\alpha \mathrel{|\!\sim}_R \beta$ iff

$$deg(\alpha \wedge \beta) < deg(\alpha \wedge \neg\beta) \quad \text{or}$$
$$deg(\alpha) = \infty.$$

Example

For

$$B \mid\!\sim F,\ B \mid\!\sim W,\ P \mid\!\sim B,\ P \mid\!\sim \neg F$$

in the rational closure we have:

$$B \wedge Gr \mid\!\sim_R F,\ P \wedge Gr \mid\!\sim_R \neg F.$$

Limitations

Two major weaknesses with the rational closure:

- Can't inherit properties across exceptional subclasses.
  - E.g. can't conclude that $P \mid\!\sim_R W$ even though we have $P \mid\!\sim B$ and $B \mid\!\sim W$.

Limitations

Two major weaknesses with the rational closure:

- Can't inherit properties across exceptional subclasses.
  - E.g. can't conclude that $P \mid\!\sim_R W$ even though we have $P \mid\!\sim B$ and $B \mid\!\sim W$.
- Undesirable specificities are sometimes obtained. For example:
  - Add to our example $L \mid\!\sim C$ (large animals are calm).
  - Get that $deg(L) = deg(L \wedge \neg P) = 0$ and $deg(P) = deg(L \wedge P) = 1$.
  - Hence $deg(L \wedge \neg P) < deg(L \wedge P)$, and obtain that $L \mid\!\sim_R \neg P$.

Implementing the Rational Closure: System Z [Pearl, 1990]

Idea: A set of default conditionals R is partitioned into a list of mutually exclusive sets of rules $R_0, \ldots, R_n$.

- Lower ranked rules are more normal (or less specific) than higher ranked rules.
- Rules in higher-ranked sets conflict in some fashion with rules in lower-ranked sets.
- The ordering is determined by treating rules as material conditionals, and using standard propositional satisfiability.
- This ordering on rules induces an ordering on models.
- α *1-entails* β given R, written $\alpha \vdash_1 \beta$, if the least $\alpha \wedge \beta$ models are less than the least $\alpha \wedge \neg\beta$ models.
- 1-entailment corresponds with the rational closure.

Example

For

$$R = \{B \Rightarrow F,\; B \Rightarrow W,\; P \Rightarrow B,\; P \Rightarrow \neg F,\; P \wedge L \Rightarrow F\}$$

we obtain:

$$\begin{aligned} R_0 &= \{B \Rightarrow F,\; B \Rightarrow W\} \\ R_1 &= \{P \Rightarrow B,\; P \Rightarrow \neg F\} \\ R_2 &= \{P \wedge L \Rightarrow F\} \end{aligned}$$

Concluding Points

- Deciding membership in the rational closure can be done with $O(\log R)$ calls to an NP oracle.
  - Thus the problem is in $P^{NP[O(\log R)]}$.
- Despite the mentioned limitations, this work spurred a great deal of interest and research.
- While the focus has been on NMIR's, there are arguments in favour of using a conditional logic formulation.
  - E.g. In a NMIR, quantification is problematic, whereas there is no problem in principle with quantification in a conditional logic.

Concluding Remarks

Concluding Remarks

Further Issues

While research in “classical” nonmonotonic reasoning has decreased since it’s height in the late 1980’s and 1990’s, there are still plenty of open issues.

Defaults and the Real World

- Most NM approaches provide mechanisms whereby various phenomena can be encoded.
- We still don't have a comprehensive theory of defaults, as things existing in the "real world".
  - Partial exception: conditional logics.
  - But no approach is fully adequate for reasoning with defaults.
  - (See [Delgrande, 2011] for more.)
- Other types of defaults, such as deontics, counterfactuals, etc.?

First-Order Defaults

- Default logic and circumscription are most appropriate for reasoning about individuals.
  - For the first order case they are either lacking (DL) or inadequate (circ) for dealing with quantification.

☞ Basically, we don't have a good theory of first-order defaults.

First-Order Defaults

- Default logic and circumscription are most appropriate for reasoning about individuals.
  - For the first order case they are either lacking (DL) or inadequate (circ) for dealing with quantification.

☞ Basically, we don't have a good theory of first-order defaults.

- Example problem (with suggestive notation):

$$\begin{aligned} \forall x, y \; Elephant(x) \land Keeper(y) &\rightarrow \; Likes(x, y) \\ \forall x \; Elephant(x) \land Keeper(Fred) &\rightarrow \neg Likes(x, Fred) \\ Elephant(Clyde) \land Keeper(Fred) &\rightarrow \; Likes(Clyde, Fred) \end{aligned}$$

NMR and Belief Revision

- The area of belief change is an important subarea and, in many aspects, largely unexplored.
- [Gärdenfors and Makinson, 1994] shows a strong connection between preferential reasoning and belief revision.
- As well, the *Ramsey Test* gives an appealing connection between BR and NMR:

  An agent accepts a default $\alpha \to \beta$ just if, in revising its beliefs by α it comes to believe β.

- General issue: What is the connection between NMR and BR?

References

James Delgrande.
What's in a default? thoughts on the nature and role of defaults in nonmonotonic reasoning.
In Gerhard Brewka, Victor W. Marek, and Miroslaw Truszczynski, editors, *Nonmonotonic Reasoning: Essays Celebrating its 30th Anniversary*. College Publications, 2011.

Marc Denecker, Victor W. Marek, and Miroslaw Truszczyński.
Uniform semantic treatment of default and autoepistemic logics.
Artificial Intelligence, 143(1):79–122, 2003.

T. Eiter and G. Gottlob.
Propositional circumscription and extended closed world reasoning are Π_2^p-complete.
Theoretical Computer Science, 114:231–245, 1993.

Peter Gärdenfors and David Makinson.
Nonmonotonic inference based on expectations.
Artificial Intelligence, 65(2):197–245, 1994.

M. Gelfond and V. Lifschitz.
Classical negation in logic programs and deductive databases.
New Generation Computing, 9:365–385, 1991.

M. Gelfond.
Answer sets.
In F. van Harmelen, V. Lifschitz, and B. Porter, editors, *Handbook of Knowledge Representation*, pages 285–316. Elsevier Science, San Diego, USA, 2008.

Georg Gottlob.
Complexity results for nonmonotonic logics.
Journal of Logic and Computation, 2(3):397–425, 1992.

S. Kraus, D. Lehmann, and M. Magidor.
Nonmonotonic reasoning, preferential models and cumulative logics.