



TECHNISCHE
UNIVERSITÄT
DRESDEN

DRESDEN
concept



Faculty of Computer Science Chair of Automata Theory

INTRODUCTION TO NONMONOTONIC REASONING

Anni-Yasmin Turhan

Dresden, WS 2019/20

Section 3

Default Logic

Subsection 3.1

Introducing defaults and default logics

Introducing defaults and default logics: an example

Suppose you are asked how you get to the university in the morning.

By bike (usually)! $\frac{\textit{gotoWork} : \textit{byBike}}{\textit{byBike}}$

new information: It is snowing heavily and your bike's tire is flat.
(You cannot assume that you go by bike.)

The default is no longer applicable \rightsquigarrow Revise previous conclusion

Why can **classical logic** not model this?

$\textit{goToUni} \wedge \neg \textit{snow} \wedge \neg \textit{FlatTire} \longrightarrow \textit{useBike}$

- There are more reasons not to use the bike, e.g. broke brakes, demonstration in the city, ... The formula would need to **list all possible obstacles!**
- **All preconditions** would need to be established to be true, so that the rule applies!

Default reasoning for kinds of common sense reasoning

Defaults can be used to model several forms of (common sense) reasoning.

- **Prototypical reasoning:** most instances of a category have a property.

“Typically, children have parents.” $\frac{child(X) : hasParents(X)}{hasParents(X)}$

- **No-risk reasoning:** concerns situations where a conclusion is drawn even if it is not the most probable, because another conclusion could lead to disaster.

“In absence of evidence to the contrary assume the accused is innocent.” $\frac{accused(X) : innocent(X)}{innocent(X)}$

- **Best-guess reasoning:** for instance, we know that there are some shopping centers in this city and some are open on Sundays, but we don't know which one. We would try the closest first, although we don't have evidence of it.

$$\frac{closest(X) : openSundays(X)}{openSundays(X)}$$

Default reasoning appears in many application domains: legal reasoning, diagnosis, reasoning about actions, etc.

Introducing defaults and default logics

- Default logics were introduced by Ray Reiter in 1980
- Default reasoning appears when reasoning is done under the closed world assumption and using inference rules that admit exceptions (rules that hold under the normality assumption)
“... in absence of any information to the contrary, assume ...”
- Classical inference rules sanction the derivation of a formula whenever some other formulae are derived.
- Default rules require an additional consistency condition to hold.

E.g.: the rule “normally birds fly” is represented as $\frac{bird(x) : flies(x)}{flies(x)}$

This states that:

“if $bird(J)$ is given or derived for a particular ground term J and $flies(J)$ is consistent (there is no information that $\neg flies(J)$ holds), then $flies(J)$ can be derived “by default”.

Consistent with what? Set of formulae that can “reasonably” be accepted based on the available information.

Syntax of Default Logic

Definition 3.1 (Default theory)

A default theory is a pair (W, D) consisting of

- W : a set of FOL formulae (called facts or axioms) and
- D : a countable set of defaults

A default δ has the form

$$\frac{\varphi : \psi_1, \dots, \psi_n}{\chi},$$

where $\varphi, \psi_1, \dots, \psi_n$, and χ are closed FOL formulae and $n > 0$.

The formula

- φ is called the prerequisite (denoted by $pre(\delta)$),
- ψ_1, \dots, ψ_n the justifications (denoted by $just(\delta)$), and
- χ the consequent of δ (denoted by $cons(\delta)$).

Why closed formulae?

What happens when open formulae are used?

Actually,

$$\frac{\textit{bird}(x) : \textit{flies}(x)}{\textit{flies}(x)}$$

is not a default according to Definition 3.1, since the formulae are not closed. We call such “defaults” **open defaults**. An open default is interpreted as a **default schema** representing a (possibly infinite) **set of defaults**.

A **default schema** differs from a default in that $\varphi, \psi_1, \dots, \psi_n, \chi$ are arbitrary FOL formulae (may contain free variables). A **default schema** defines a set of defaults

$$\frac{\varphi\sigma : \psi_1\sigma, \dots, \psi_n\sigma}{\chi\sigma}$$

for *all* ground substitutions σ that assign values to all free variables occurring in the schema.

↪ Free variables : interpreted as being **universally quantified** over the **whole default schema**

Why closed formulae?

Trying to use quantifiers

The open default

$$\frac{bird(x) : flies(x)}{flies(x)}$$

would read under

- universally quantified variables as
"If all X are birds, and if for all X we may assume that they fly, then we conclude that *all* X fly."
 - Does not match the intuition. All birds are treated uniformly.
 - Would only be applicable, if every object in the domain is a bird.
- existentially quantified variables as
"If there is a bird and if there is an X that flies, then conclude that there is some flying object."
 - Would not allow to conclude from $bird(tweety)$ that $flies(tweety)$ holds.
 - instead we would conclude: $\exists X flies(X)$

↪ Using one of the quantifiers is not adequate!

Towards the semantics of defaults

The informal meaning of a default $\frac{\varphi : \psi_1, \dots, \psi_n}{\chi}$ is:

“If φ is known, and if it is consistent to assume that ψ_1, \dots, ψ_n , then conclude χ .”

In the formal semantics we must say

1. where φ should be included
2. with what should ψ_1, \dots, ψ_n be consistent

With what should ψ_1, \dots, ψ_n be consistent? A first attempt: the facts.

Consider the default

$$\frac{\text{friend}(X, Y) \wedge \text{friend}(Y, Z) : \text{friend}(X, Z)}{\text{friend}(X, Z)}$$

Given the facts: $\text{friend}(\text{tom}, \text{bob}), \text{friend}(\text{bob}, \text{sally}), \text{friend}(\text{sally}, \text{tina})$.

Wanted conclusion: $\text{friend}(\text{tom}, \text{tina})$

This is only possible if we apply the appropriate instance of the default schema to $\text{friend}(\text{sally}, \text{tina})$ and $\text{friend}(\text{tom}, \text{sally})$. But $\text{friend}(\text{tom}, \text{sally})$ is derived by a previous application of the default schema!

Without this intermediate result and from the facts alone, we could not derive this.

Towards the semantics of defaults

Example 3.2

Let's consider $T = (W, D)$ with $W = \{green, ADACmember\}$ and $D = \{\delta_1, \delta_2\}$, where

$$\delta_1 = \frac{green : \neg likesCars}{\neg likesCars} \quad \text{and} \quad \delta_2 = \frac{ADACmember : likesCars}{likesCars}$$

If consistency is tested against W , both defaults can be applied. Deriving $\neg likesCars$ and $likesCars$, which is a contradiction!

Alternative:

apply the first default δ_1 , check for consistency with the knowledge derived so far. Would block the application of the second default δ_2 .

Informal semantics of defaults

A general formulation:

If φ is currently known, and if all ψ_i are consistent with the current knowledge base, then conclude χ .

The current knowledge base E is obtained from

- the facts and
- the consequents of some defaults that have been applied previously.

A more formal version:

Default $\delta = \frac{\varphi : \psi_1, \dots, \psi_n}{\chi}$ is applicable to a deductively closed set of formulae E iff $\varphi \in E$ and $\neg\psi_1 \notin E, \dots, \neg\psi_n \notin E$.

Towards extensions

Example 3.2 suggests that there can be several competing knowledge bases which maybe inconsistent with each other.

Extensions

represent possible world views which are based on the given default theories. They seek to extend the set of known facts with “reasonable” conjectures based on the available defaults.

Desirable properties of extensions:

- an extension should include the set W of facts—the certain information
- an extension should be deductively closed
(Keep classical reasoning! Derive more from the defaults.)
- an extension should be closed under the application of the defaults. Apply defaults exhaustively.

Formally: if $\frac{\varphi : \psi_1, \dots, \psi_n}{\chi} \in D, \varphi \in E$ and $\neg\psi_1 \notin E, \dots, \neg\psi_n \notin E$ then $\chi \in E$.

↪ Extensions are maximal possible world views.

Towards extensions — unwanted effects

1. "Ungrounded" beliefs

An extension must not contain "ungrounded" beliefs, i.e. every formula in the extension must be derivable from W and the consequents of applied defaults. We require extensions to be minimal w.r.t. to these properties.

Consider: $T = (W, D)$ with $W = \{german\}$ and $D = \left\{ \frac{german : drinksBeer}{drinksBeer} \right\}$

Now, $E = Th(\{german, \neg drinksBeer\})$ is minimal w.r.t. to the properties, but unintuitive.

2. Applications of defaults can contradict the application of an earlier default.

Consider:

$$\frac{true : creditworthy}{approveCredit}, \quad \frac{true : \neg creditworthy}{\neg creditworthy}$$

We apply the first default, since nothing contradicts the assumption $creditworthy$. We then apply the second, since $\neg creditworthy$ is consistent with the knowledge, $\neg creditworthy$ is derived.

Inclusion of $\neg creditworthy$ shows *a-posteriori* that, we should not have assumed $creditworthy$.

Subsection 3.2

Operational semantics of Default Logic

- based on the process in which inferences are drawn
- gives a procedure that can be applied

Idea:

- apply defaults as long as possible
- If a default should not have been applied, backtrack and try an alternative

Operational Semantics

Given a default theory $T = (W, D)$ let $\Pi = (\delta_0, \delta_1 \dots)$ be (a finite or infinite) sequence of defaults from D without multiple occurrences.
(Possible order in which some defaults from D are applied.)

$\Pi[k]$ denotes the initial segment of sequence Π of length k .¹

Each sequence Π is associated with two sets: $In(\Pi)$ and $Out(\Pi)$

- $In(\Pi) = Th(W \cup \{cons(\delta) \mid \delta \text{ occurs in } \Pi\})$.
- $Out(\Pi) = \{\neg\psi \mid \psi \in just(\delta) \text{ for some } \delta \text{ in } \Pi\}$.

Intuition:

- $In(\Pi)$ represents the current knowledge base after the defaults in Π have been applied
- $Out(\Pi)$ represents the formulae that should not become true even after subsequent application of other defaults.

¹We assume (from now on) that the length of Π is at least k .

Example: default sequences

Example 3.3

Consider $T = (W, D)$ with $W = \{a\}$ and the defaults from D :

$$\delta_1 = \frac{a : \neg b}{\neg b}, \quad \delta_2 = \frac{b : c}{c}$$

For $\Pi_a = (\delta_1)$ we have $In(\Pi_a) = Th(\{a, \neg b\})$ and $Out(\Pi_a) = \{b\}$.

For $\Pi_b = (\delta_2, \delta_1)$ we have $In(\Pi_b) = Th(\{a, c, \neg b\})$ and $Out(\Pi_b) = \{\neg c, b\}$

We have not assured that the defaults can be applied in the order given.
 (δ_2, δ_1) cannot be applied in this order, since $b \notin In(\emptyset) = Th(W) = Th(a)$.

"Applicable sequences" are formalized by the notion of a *process*.

Processes

Definition 3.4 (Process, successful, closed)

Π is a **process** of \mathcal{T} iff δ_k is applicable to $In(\Pi[k])$ for every k s.t.² δ_k occurs in Π .

Let Π be a process. We define:

- Π is **successful** iff $In(\Pi) \cap Out(\Pi) = \emptyset$. Otherwise, it is **failed**.
- Π is **closed** iff every $\delta \in D$ that is applicable to $In(\Pi)$ already occurs in Π .

Intuition:

Success of a process captures that it was "okay" to have assumed the justifications of the applied defaults; no formula $\neg\psi \in Out(\Pi)$ is part of the current knowledge base, so it was consistent to assume ψ .

Closed processes correspond to the extension being closed under application of the defaults.

²"such that"

Example: properties of processes

Consider the default theory $T = (W, D)$ with $W = \{a\}$ and D containing

$$\delta_1 = \frac{a : \neg b}{d}, \quad \delta_2 = \frac{true : c}{b}$$

$\Pi_1 = (\delta_1)$

is successful, but not closed, since δ_2 may be applied to $In(\Pi_1) = Th(\{a, d\})$.

$\Pi_2 = (\delta_2, \delta_1)$

is closed, but not successful. Since both $In(\Pi_2) = Th(a, b, d)$ and $Out(\Pi_2) = \{b, \neg c\}$ contain b .

$\Pi_3 = (\delta_2)$

is a closed and successful process of T .

Extension and closure—operational semantics

Definition 3.5 (Extension)

Let T be a default theory. A set of formulae E is an **extension of T** iff there is some closed and successful process Π s.t. $E = In(\Pi)$.

This definition may be applied directly to concrete examples.

To find a successful process, it suffices to generate a process Π , test whether $In(\Pi) \cap Out(\Pi) = \emptyset$ holds. If not, then backtrack.

A **(in)finite default theory** is a default theory, where D has (in)finitely many elements.

For finite default theories ensuring closure is conceptually easy: apply an applicable default that has not been applied yet, until no more a left.

How about closure of *infinite* default theories?

Closure of infinite theories

Lemma 3.6

An *infinite process* Π of a default theory $T = (W, D)$ is *closed* iff each default in D that is applicable to $In(\Pi[k])$, for infinitely many numbers k , is already contained in process Π .

The proof of Lemma 3.6 needs:

Compactness Theorem

The *compactness theorem* says that a set of M of formulae is satisfiable iff every finite subset of M is satisfiable.

“A set of first-order sentences has a model iff every finite subset of it has a model.”
↪ method for constructing models of any set of sentences that is finitely consistent.

Proof of Lemma 3.6 : blackboard

A strategy to guarantee the closure of an infinite process Π must ensure that any default which from k on, demands application, will eventually be applied.
(A.k.a. *fairness* condition from concurrent programming)

Finding closed and successful processes

The process of finding a closed and successful process can be represented by a kind of (search) tree.

Definition 3.7 (Process tree)

Let $T = (W, D)$ be a default theory. A **process tree** is tree $G = (V, E)$, s.t. all nodes $v \in V$ are labeled with two sets of formulae:

- an **In-set** $In(v)$ and
- an **Out-set** $Out(v)$.

The root of G is labeled with $Th(W)$ as In-set and \emptyset as Out-set.

The edges in E correspond to default applications and are labelled with the default applied.

The **paths of a process** are the paths in G starting at the root.

A node v is **expanded** if $In(v) \cap Out(v) = \emptyset$, otherwise it is marked "failed" and is a leaf of the process tree.

Process Trees (cont.)

Definition 3.7: Process tree (cont.)

If $v \in V$ is expanded it possesses for each default $\delta = \frac{\varphi : \psi_1, \dots, \psi_n}{\chi}$ one successor node w that

- does not appear on the path from the root node to v ,
- is applicable to $In(v)$,
- is connected to v by a edge labeled with δ , and
- is labeled with $Th(In(v) \cup \{\chi\})$ and $Out(v) \cup \{\neg\psi_1, \dots, \neg\psi_n\}$.

Subsection 3.3

Original semantics of default logics

- original definition by Ray Reiter
- fixed-point based, not constructive

Consistency w.r.t. to what?

When applying defaults we need to ensure consistency. But consistency w.r.t. to which theory?

We consider again Example 3.2:

$T = (W, D)$ with $W = \{green, ADACmember\}$ and $D = \{\delta_1, \delta_2\}$, where

$$\delta_1 = \frac{green : \neg likesCars}{\neg likesCars} \quad \text{and} \quad \delta_2 = \frac{ADACmember : likesCars}{likesCars}$$

Consistency w.r.t. to W alone is not enough.

Solution by Reiter: Use a theory that plays the role of a context or belief set. Check consistency against this context.

A formalization of this idea:

A default $\delta = \frac{\varphi : \psi_1, \dots, \psi_n}{\chi}$ is applicable to a deductively closed set of formulae F w.r.t. belief set E (the context) iff $\varphi \in F$ and $\neg\psi_1 \notin E, \dots, \neg\psi_n \notin E$ (each ψ_i is consistent with E).

Note that the concept “ δ is applicable to E ” is so far a special case where $E = F$.

Which context to use?

Observation:

If a default has been applied to a belief set E , a formula has been derived and is part of the knowledge base. Therefore it should be believed, i.e. become an element of belief set E .

On the other hand, E should contain only formulae that can be derived from the axioms in W by default application.

Definition 3.8 (Closure under a set of defaults w.r.t. a belief set)

Let D be a set of defaults and F a deductively closed set of formulae.

F is closed under D w.r.t. belief set E iff, for every default $\delta \in D$ that is applicable to F w.r.t. belief set E , its consequent χ is also contained in E .

Lemma 3.9

Let $E' \subseteq E$ and F be a set of formulae closed under some set of defaults D w.r.t. E' . Then F is closed under D w.r.t. E .

Proof: exercise

Extension — original semantics

Definition 3.10 ($\Lambda_T(E)$, extension)

Given $T = (W, D)$ and a set of formulae E . Let $\Lambda_T(E)$ be the smallest³ set of formulae that

- contains W ,
- is closed under deduction, i.e. contains all conclusions, and
- is closed under D w.r.t. E .

E is an extension of T , iff $E = \Lambda_T(E)$.

Intuition:

- $\Lambda_T(E)$ contains all formulae that are sanctioned by T w.r.t. E .
- E is an extension, iff by the use of E as a belief set, exactly the formulae in E will be obtained from default application.

Observe: one first needs to guess E and then check whether the fixed-point equation is fulfilled.

³. i.e. has smallest number of elements

Are the two definitions of extension equivalent?
Are operational semantics and the original one the same?

Theorem 3.11

Let $T = (W, D)$ be a default theory.

E is an extension of T (according Definition 3.5) iff $E = \Lambda_T(E)$.

Proof: blackboard

Minimality of Reiter's extensions

Reiter's characterization fulfills the desirable properties of an extension:

- includes the set W of facts: E includes W .
- is deductively closed: E is deductively closed
- is closed under the application of the defaults: E is closed under D w.r.t. E

Claim: E is minimal w.r.t. these properties.

Corollary 3.12 (Minimality of extension)

Let E and E' be two extensions of a default theory T . If $E' \subseteq E$, then $E' = E$.

Proof: If E' is an extension and $E' \subseteq E$, then E' is closed under D w.r.t. E (by Lemma 3.9).

By definition, $E = \Lambda_T(E) \subseteq E'$ and thus $E' = E$.

Properties of extensions

Theorem 3.13 (Consistency preservation)

A default theory $T = (W, D)$ has an inconsistent extension iff W is inconsistent.

Proof: Exercise

Corollary 3.14

If a default has an inconsistent extension E , then it is its only extension.

Theorem 3.15

Let $T = (W, D)$ be a default theory s.t. the set

$M = W \cup \{\psi_1 \wedge \dots \wedge \psi_n \wedge \chi \mid \frac{\varphi : \psi_1, \dots, \psi_n}{\chi} \text{ is a default in } D\}$ is consistent. Then T has exactly one extension.

Proof: blackboard

Nonmonotonic nature of Default logic

Nonmonotonic behaviour may appear when the default theory is changed!

Example 3.16 (Changing the defaults)

Let $T_{ex}(W, D)$ be a default theory with $W = \emptyset$ and $D = \{\frac{true:a}{a}\}$.
 T_{ex} has exactly one extension: $E = Th(\{a\})$.

- $\delta_1 = \frac{true:b}{\neg b}$. Then $(W, D \cup \{\delta_1\})$ has no extensions.
- $\delta_2 = \frac{b:c}{c}$. Then $(W, D \cup \{\delta_2\})$ has E as only extension.
- $\delta_3 = \frac{true:\neg a}{\neg a}$. Then $(W, D \cup \{\delta_3\})$ has two extensions: E and $Th(\{\neg a\})$.
- $\delta_4 = \frac{a:b}{b}$. Then $(W, D \cup \{\delta_4\})$ has the two extensions: $Th(\{a, b\})$.

Nonmonotonic nature of Default logic

Nonmonotonic behaviour may appear when the default theory is changed!

Example 3.17 (Changing the facts)

Let $T_{ex}(W, D)$ be a default theory with $W = \emptyset$ and $D = \{\delta_1, \delta_2, \delta_3, \delta_4, \delta_5\}$ with

$$\delta_1 = \frac{true : a, \neg c}{a}, \delta_2 = \frac{a : b, \neg c}{b}, \delta_3 = \frac{true : \neg a, c}{c}, \delta_4 = \frac{d : e}{e}, \delta_5 = \frac{f : g}{\neg g}$$

T_{ex} has two extensions: $E_1 = Th(\{a, b\})$ and $E_2 = Th(\{c\})$. Consider

- $W_1 = \{f\}$. (W_1, D) has no extensions.
- $W_2 = \{\neg a\}$. (W_2, D) has the only extension $Th(\{\neg a, c\})$.
- $W_3 = \{\neg a, \neg c, d\}$. (W_3, D) the only new extension: $Th(\{\neg a, \neg c, d, e\})$.
- $W_4 = \{d\}$. (W_4, D) the two extension: $Th(E_1 \cup \{d, e\})$ and $Th(E_2 \cup \{d, e\})$.

Subsection 3.4

Normal default logics

Normal defaults and normal default theories

Definition 3.18 (Normal defaults, normal default theory)

A default is *normal* iff its consequent is its only justification. They have the form

$$\frac{\varphi : \psi}{\psi}.$$

A default theory $T = (W, D)$ is *normal* iff all defaults in D are normal.

Normal defaults are well-behaved

- have always extensions
- rule out "pathological cases" such as: $\frac{true : a}{\neg a}$
- have limited expressivity: no interactions among defaults

A normal default theory

draws the conclusion ψ when φ is known and it is consistent to conclude ψ .

Processes of normal default theories

Lemma 3.19

Each process of a normal default theory is successful.

Proof: blackboard

Expanding the sequence Π in a fair way (starting from a consistent set of facts W), a closed and successful process is obtained. Thus we have established:

Theorem 3.20 (Existence of extensions)

Normal default theories (with a consistent set of facts W) always possess extensions. Every finite process Π may be expanded to a closed process Π' .

By Lemma 3.19 every such process is also successful.

Default theories are not semi-monotonic and need not have orthogonal extensions

Semi-monotonicity

A class of default theories, is **semi-monotonic** iff the addition of default rules never eliminates, but extends or adds new extensions.

We saw in Example 3.16 that this property need not hold for default theories in general.

Orthogonality of extensions

Two extensions of the same default theory are **orthogonal** to each other, if they are inconsistent.

Consider the general default theory $T = (\emptyset, D)$ with

$$D = \left\{ \frac{true : \neg p}{q}, \frac{true : \neg q}{p} \right\}$$

Due to conflicting defaults T has two extensions $E_1 = Th(\{q\})$ and $E_2 = Th(\{p\})$. But clearly, $E_1 \cup E_2$ is consistent. Default theories in general need not have orthogonal extensions.

Properties of normal default theories

Theorem 3.21 (Semi-monotonicity)

*Let $T = (W, D)$ and $T' = (W, D')$ be normal default theories s.t. $D \subseteq D'$.
Then each extension of T is contained in an extension of T' .*

Proof: blackboard

Theorem 3.22 (Orthogonality of extensions)

*Let E and F be different extensions of a normal default theory T .⁴
Then $E \cup F$ is inconsistent.*

Proof: blackboard

⁴Note that it follows by Consistency preservation (Theo. 3.13 and Cor. 3.14) that W must be consistent.

Limitations of normal default theories

Normal default theories are well-behaved. But are normal default theories **expressive enough** to model common sense reasoning?

Statements such as:

- “Typically birds fly”
- “Assume the accused is innocent unless you know otherwise”

can be captured by normal defaults:

$$\frac{\textit{bird}(x) : \textit{flies}(x)}{\textit{flies}(x)} \qquad \frac{\textit{accused}(x) : \textit{innocent}(x)}{\textit{innocent}(x)}$$

Often a default rule on its own is normal, but problems arise when **several defaults have to interact in a theory**.

Limitations of normal default theories—example

Example 3.23

Consider the example of a normal default theory:

$$T = \left(\{ \text{dropout}(\text{bill}) \}, \left\{ \frac{\text{dropout}(x) : \text{adult}(x)}{\text{adult}(x)}, \frac{\text{adult}(x) : \text{employed}(x)}{\text{employed}(x)} \right\} \right)$$

T has the single extension $Th(\{\text{dropout}(\text{bill}), \text{adult}(\text{bill}), \text{employed}(\text{bill})\})$. But it is counter-intuitive to assume that Bill is employed!

How to prevent the application of the 2. default, if X is a dropout?

$$\frac{\text{adult}(x) : \text{employed}(x) \wedge \neg \text{dropout}(x)}{\text{employed}(x)}$$

But this is no longer a normal default!

↪ Normal defaults are too restrictive for modeling!

Semi-normal Defaults

Definition 3.24 (Semi-normal default, semi-normal default theory)

A default is a **semi-normal default**, if it has the form

$$\frac{\phi : \psi \wedge \chi}{\psi}.$$

A default theory $T = (W, D)$ is **semi-normal**, if all defaults in D are semi-normal.

Example 3.23 used a semi-normal default.

Semi-normal defaults can “implement” priorities between defaults.

Consider $T = (\{bird(X) \rightarrow animal(X), bird(tweety)\}, \{\delta_1, \delta_2\})$ with

$$\delta_1 = \frac{animal(X) : \neg flies(X)}{\neg flies(X)}, \quad \delta_2 = \frac{bird(X) : flies(X)}{flies(X)}$$

T has two extensions: one includes $\neg flies(tweety)$ and the other $flies(tweety)$.

The second extension is preferable, since more specific information was used.

This can be achieved by:

$$\delta_1' = \frac{animal(X) : \neg bird(X) \wedge \neg flies(X)}{\neg flies(X)}$$

Are semi-normal defaults well-behaved?

Do semi-normal default theories always have extensions?

No.

Consider the example: $T = (W, D)$, with $W = \emptyset$ and

$$D = \left\{ \frac{\text{true} : \neg q \wedge p}{p}, \frac{\text{true} : \neg r \wedge q}{q}, \frac{\text{true} : \neg p \wedge r}{r} \right\}$$

Only some classes of restricted semi-normal theories do always have extensions.

Semi-normal default theories are not well-behaved

Semi-normal default theories do not have ...

- Semi-monotonicity

Consider:

$T = (W, D)$ with $W = \emptyset$ and $D = \left\{ \frac{\text{true} : \neg q \wedge p}{p} \right\}$ and

$T' = (W, D')$ with $W = \emptyset$ and $D' = \left\{ \frac{\text{true} : \neg q \wedge p}{p}, \frac{\text{true} : \neg r \wedge q}{q} \right\}$.

We have $E = Th(\{p\})$ and $E' = Th(\{q\})$ as the extensions of the theories, but $E \not\subseteq E'$.

- Success of all processes

T' has a failed process.

- Orthogonality of extensions

Consider:

$T'' = (W, D)$ with $W = \emptyset$ and $D = \left\{ \frac{\text{true} : p \wedge q}{q}, \frac{\text{true} : \neg q \wedge \neg p}{\neg p} \right\}$.

T'' has two extensions: $E_1 = Th(\{q\})$ and $E_2 = Th(\{\neg p\})$, but $E_1 \cup E_2$ is consistent.

Reasoning in Default Logics

Classical reasoning problems of default theories are:

- deciding whether a default theory has an extension
- deciding whether a given formula is **element of all extensions**.
("cautious reasoning" or "skeptical reasoning")
- deciding whether a given formula is **element of one extension**.
("brave reasoning" or "credulous reasoning")
choosing a different extension, may yield different consequences.

Example: reasoning in Default Logics

Example 3.25 (Reprise of Example 3.17)

Let $T_{ex} = (W, D)$ be a default theory with $W = \{d\}$ and $D = \{\delta_1, \delta_2, \delta_3, \delta_4, \delta_5\}$ with

$$\delta_1 = \frac{true : a, \neg c}{a}, \delta_2 = \frac{a : b, \neg c}{b}, \delta_3 = \frac{true : \neg a, c}{c}, \delta_4 = \frac{d : e}{e}, \delta_5 = \frac{f : g}{\neg g}$$

T_{ex} has two extensions: $E_1 = Th(\{d, e, a, b\})$ and $E_2 = Th(\{d, e, c\})$.

Formula $(d \wedge e)$ is a consequence for T_{ex} under cautious reasoning.

Formula c is a consequence for T_{ex} under brave reasoning.

Recap on complexity classes

Complexity class with **oracle** admits the use of a subroutine “at no cost”

Polynomial hierarchy: the classes Π_k^P , Σ_k^P and Δ_k^P are defined as follows:

$$P = \Sigma_0^P = \Pi_0^P = \Delta_0^P$$

and for all $k \geq 0$:

$$\Sigma_{k+1}^P = NP^{\Sigma_k^P} \quad \Pi_{k+1}^P = \text{co-}\Sigma_{k+1}^P \quad \Delta_{k+1}^P = P^{\Sigma_k^P}$$

Note: $\Sigma_1^P = NP$, $\Pi_1^P = \text{co-}NP$ and $\Delta_{k+1}^P = P$.

Complexity of reasoning in default theories

Typically, default reasoning is harder than classical reasoning.

Computability / complexity results for different classes of default theories:

- FOL default theories: **undecidable**
(since classical reasoning is already undecidable)
Reasoning in default theories is **not even semi-decidable**, since computing an extension requires FOL consistency tests which are semi-decidable.
- normal default theories: **undecidable**
- propositional default theories:
 - deciding existence of an extension: Σ_2^P -complete
 - brave reasoning is Σ_2^P -complete
 - cautious reasoning is Π_2^P -complete