Tools for Formal Epistemology:
Doxastic Logic, Probability, and Default Logic
– ESSLLI 2023 –

Lecture 4 – Default Logic c'd, Starting Applications

Aleks Knoks, University of Luxembourg
Eric Pacuit, University of Maryland

1 Quick recap

1.1 Default logic

- Default logic is one kind of defeasible logic.
  (The context here is symoblic AI.)
- The basic idea behind default logic: supplement
  CL with a set of *default rules*.

1.2 Core question

- What can one conclude from a default theory?
  Alternatively: What is the *extension* of a given default
  theory $\langle \mathcal{W}, \mathcal{D}, < \rangle$?
- We answer the question in a roundabout way..

$$\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$$

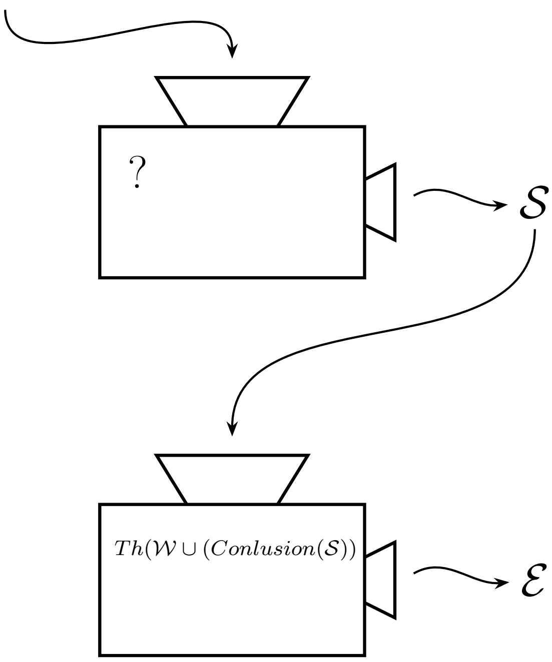

- We still need to specify how to select proper scenarios..

2 Binding defaults

2.1 Intuition

- *Binding defaults* are, intuitively, the *good* or *correct* ones to base the conclusions on in the context. Proper scenarios will contain all and only such defaults.

- The concept emerges as a combination of three preliminary notions: (i) *triggering*, (ii) *conflict*, and (iii) *defeat*.

2.2 Triggered defaults

- $Triggered_{\langle \mathcal{W},\mathcal{D},<\rangle}(\mathcal{S}) = \{\delta \in \mathcal{D} : \mathcal{W} \cup Conclusion(\mathcal{S}) \vdash Premise(\delta)\}$

2.3 Example (Tweety)

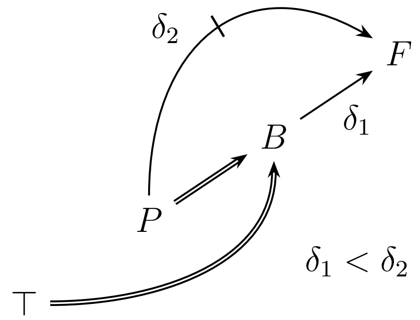

- $\Delta_1' = \langle \mathcal{W}, \mathcal{D}, <\rangle$ where
  $\mathcal{W} = \{B, P \supset B\}$
  $\mathcal{D} = \{\delta_1, \delta_2\}$
  - $\delta_1 = B \to F$
  - $\delta_2 = P \to \neg F$

  $\delta_1 < \delta_2$
- Here $Triggered_{\Delta_1}(\emptyset) = \{\delta_1\}$

2.4 Conflicting defaults

- $Conflicted_{\langle \mathcal{W},\mathcal{D},<\rangle}(\mathcal{S}) = \{\delta \in \mathcal{D} : \mathcal{W} \cup Conclusion(\mathcal{S}) \vdash \neg Conclusion(\delta)\}$

2.5 Example (Nixon Diamond)

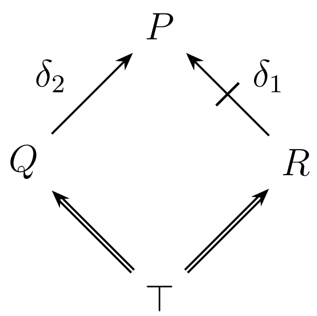

- $\Delta_2 = \langle \mathcal{W}, \mathcal{D}, <\rangle$ where
  $\mathcal{W} = \{Q, R\}$
  $\mathcal{D} = \{\delta_1, \delta_2\}$

- $\delta_1 = R \to \neg P$
- $\delta_2 = Q \to P$

$<= \emptyset$

(Q = Quaker, R = Republican, P = Pacifist)

- Here
  $Triggered_{\Delta_2}(\emptyset) = \{\delta_1, \delta_2\}$
  $Conflicted_{\Delta_2}(\emptyset) = \emptyset$
  Yet
  $Conflicted_{\Delta_2}(\{\delta_1\}) = \{\delta_2\}$
  $Conflicted_{\Delta_2}(\{\delta_2\}) = \{\delta_1\}$

2.6 Defeated defaults

- Basic idea: a default is defeated if there's a stronger default
  supporting a contrary conclusion
  Formally:

$$Defeated_{\langle \mathcal{W}, \mathcal{D}, < \rangle}(\mathcal{S}) = \{\delta \in \mathcal{D} : \exists \delta' \in Triggered_{\langle \mathcal{W}, \mathcal{D}, < \rangle}(\mathcal{S}) \text{ with } (1)\ \delta < \delta' \text{ and } (2)\ Conclusion(\delta') \vdash \neg Conclusion(\delta)\}.$$

2.7 Example (Tweety)

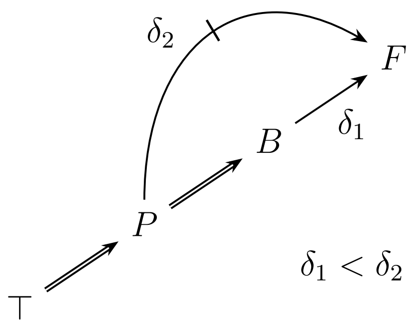

- $\Delta_1 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ where
  $\mathcal{W} = \{P, P \supset B\}$
  $\mathcal{D} = \{\delta_1, \delta_2\}$
  - $\delta_1 = B \to F$
  - $\delta_2 = P \to \neg F$

  $\delta_1 < \delta_2$

- Here δ_1 is defeated:
  $Defeated_{\Delta_1}(\emptyset) = \{\delta_1\}$.

2.8 Binding defaults

- Now we can define binding defaults:

$$Binding_{\langle \mathcal{W}, \mathcal{D}, < \rangle}(\mathcal{S}) = \{\delta \in \mathcal{D} : \delta \in Triggered_{\langle \mathcal{W}, \mathcal{D}, < \rangle}(\mathcal{S}), \delta \notin Conflicted_{\langle \mathcal{W}, \mathcal{D}, < \rangle}(\mathcal{S}), \delta \notin Defeated_{\langle \mathcal{W}, \mathcal{D}, < \rangle}(\mathcal{S})\}.$$

2.9 Stable scenarios

- Stable scenarios:
  $\mathcal{S}$ is *stable* just in case $\mathcal{S} = Binding_{\langle \mathcal{W}, \mathcal{D}, < \rangle}(\mathcal{S})$.

2.10 Example (Tweety, again)

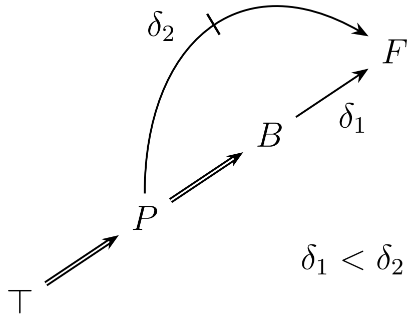

$\delta_1 < \delta_2$

- There are four scenarios based on Δ_1:

  $\mathcal{S}_1 = \emptyset$

  $\mathcal{S}_2 = \{\delta_1\}$

  $\mathcal{S}_3 = \{\delta_2\}$

  $\mathcal{S}_4 = \{\delta_1, \delta_2\}$

- Only $\mathcal{S}_3$ stable, since $\mathcal{S}_3 = Binding_{\langle \mathcal{W}, \mathcal{D}, < \rangle}(\mathcal{S}_3)$

3 Complication #1: Self-triggering chains

- There are three complications
- Complication #1: Can we identify the proper scenarios with the stable ones?

3.1 Problem

- The problem is with "groundedness" or the possibility of *self-triggering chains*
- Consider $\Delta_3 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ where
  $\mathcal{W} = \emptyset$
  $\mathcal{D} = \{\delta_1\}$
  - $\delta_1 = A \rightarrow A$

  $<= \emptyset$
- $\mathcal{S}_1 = \{\delta_1\}$ is a stable scenario, but it shouldn't be proper!

  Intuitively, we wouldn't conclude A on the basis of Δ_3

3.2 Solution (Approximating sequences)

- Let $\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$ be a default theory and $\mathcal{S} \subseteq \mathcal{D}$.
- Then $\mathcal{S}_0, \mathcal{S}_1, \mathcal{S}_2, \ldots$ is an *approximating sequence* based on Δ and constrained by $\mathcal{S}$ just in case:

$$\begin{array}{lcll} \mathcal{S}_0 & = & \emptyset & \\ \mathcal{S}_{i+1} & = & \{\delta \in \mathcal{D}: & \delta \in Triggered_{\langle \mathcal{W}, \mathcal{D}, < \rangle}(\mathcal{S}_i), \\ & & & \delta \notin Conflicted_{\langle \mathcal{W}, \mathcal{D}, < \rangle}(\mathcal{S}), \\ & & & \delta \notin Defeated_{\langle \mathcal{W}, \mathcal{D}, < \rangle}(\mathcal{S})\}. \end{array}$$

  (This is called *quasi-induction*)

  The limit of an approximating sequence is $\bigcup_{i \geq 0} \mathcal{S}_i$.

- Approximating sequence constrained by their own limits give us the *proper scenarios*:

  Given a $\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$ and an approximating sequence $\mathcal{S}_0, \mathcal{S}_1, \mathcal{S}_2, \ldots$ constrained by some scenario $\mathcal{S}$ based on Δ, the scenario $\mathcal{S}$ is proper just in case $\mathcal{S} = \bigcup_{i \geq 0} \mathcal{S}_i$.

- Theorem: if $\mathcal{S}$ is proper, then $\mathcal{S}$ is also stable.

3.3 Example

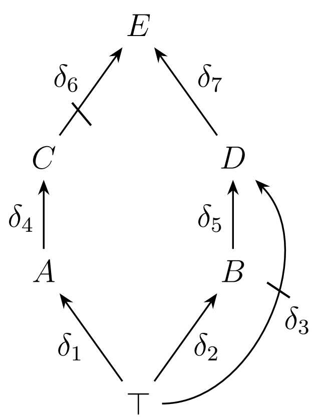

4 Complication #2: Multiple extensions

4.1 Problem

- The problem is that some theories have **multiple** proper scenarios, and so multiple extensions

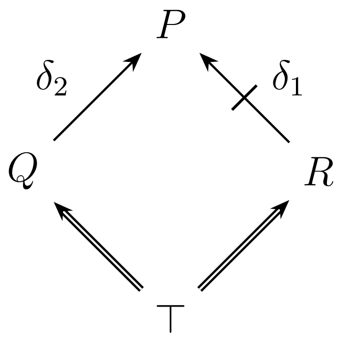

- Consider the Nixon Diamond theory $\Delta_2 = \langle \mathcal{W}, \mathcal{D}, < \rangle$
  where
  $\mathcal{W} = \{Q, R\}$
  $\mathcal{D} = \{\delta_1, \delta_2\}$
  - $\delta_1 = R \rightarrow \neg P$
  - $\delta_2 = Q \rightarrow P$

  $< = \emptyset$
- There are two proper scenarios:
  $\mathcal{S}_1 = \{\delta_1\}$
  $\mathcal{S}_2 = \{\delta_2\}$
  and two extensions
  $\mathcal{E}_1 = Th(\{Q, R, P\})$
  $\mathcal{E}_2 = Th(\{Q, R, \neg P\})$
  But what should we conclude?

4.2 Options

- #1: "Credulous": give some weight to any conclusion X contained in **some** extension
  In the Nixon case:
  either P and $\neg P$ follow, or $\mathcal{B}(P)$ and $\mathcal{B}(\neg P)$ do, where $\mathcal{B}(X)$ says that X is believable
- #2: "Skeptical": endorse X as conclusion whenever X is contained in **every** extension
  In the Nixon case:
  neither P, nor $\neg P$ follow
- The skeptical option is the most popular one
- My strategy: go with skeptical, since it's less committal

5 Complication #3: No extensions

5.1 Problem

- Here the problem is that some default theories have **no** proper scenarios

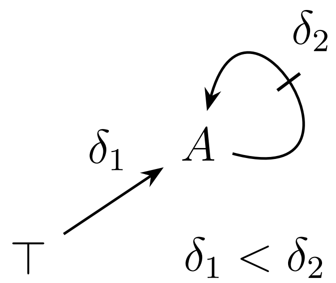

- Consider $\Delta_4 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ with
  $\mathcal{W} = \emptyset$
  $\mathcal{D} = \{\delta_1, \delta_2\}$
  - $\delta_1 = \top \to A$
  - $\delta_2 = A \to \neg A$

  $\delta_1 < \delta_2$

5.2 Options

- Syntactic restriction that rules out "vicious cycles" or *self-defeating chains*
- Live with it!
  (might seem like a benign choice if you like the skeptical option, but it doesn't give you a universal free pass)
- Move to formal argumentation theory

6 Some results

1. Let $\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$ be a (fixed priority) default theory. If $\mathcal{S}$ is a proper scenario based on Δ, it is also a stable scenario based on Δ.
2. A default theory $\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$ has an inconsistent extension if and only if $\mathcal{W}$ is inconsistent.
3. If a default theory Δ has an inconsistent extension, it's the only extension.
4. Let $\mathcal{S}$ and $\mathcal{S}'$ be proper scenarios based on the same default theory Δ such that $\mathcal{S} \subseteq \mathcal{S}'$. Then $\mathcal{S} = \mathcal{S}'$.
5. Let $\mathcal{E}$ be an extension of the default theory $\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$, and let $\mathcal{A} \subseteq \mathcal{E}$. Then $\mathcal{E}$ is also an extension of $\Delta' = \langle \mathcal{W} \cup \mathcal{A}, \mathcal{D}, < \rangle$.

6. If Δ has distinct extensions $\mathcal{E}$ and $\mathcal{E}'$, then $\mathcal{E} \cup \mathcal{E}'$ is inconsistent.

..

- Which of Reiter's (1980) results hold in *prioritized* default logic? (There's work left to do here..)

7 Some application in AI

7.1 Example 1: Closed-world assumption (CWA)

- CWA is, roughly, a system's assumption that all relevant information is at its disposal.

  Context: database theory

  But see also Daniel Kahneman's (2011) WYSIATI

- Formalization in default logic:
  Set $\mathcal{D}^{cwa} = \{\top \to \neg A : A$ is a propositional atom$\}$
  Now, instead of applying classical logic to $\mathcal{W}$, apply default logic to $\langle \mathcal{W}, \mathcal{D}^{cwa}, < \rangle$ with an empty $<$

7.2 Example 2: Planning and the Frame Problem

- The frame problem revolves around formalizing what's been called "causal inertia".

- Formalization in default logic:
  $P(x, s) \to P(x, do(a, s))$

8 Caveats

- Defaults rules we have been working with look quite different from what Reiter (1980) calls *default rules*:

  $$\delta = \frac{X\colon MZ_1, \ldots, MZ_n}{Y}$$

  (M is to be read as "it is consistent to assume"
  $MZ_1, \ldots, MZ_n$ is called the *justification*)

  Normal defaults:

  $$\delta = \frac{X\colon MY}{Y} \qquad \equiv \qquad \delta = X \to Y$$

  "In fact I know of no naturally occurring default which cannot be represented in this form" (Reiter, 1980, p. 95)

- Our definition of defeat is very simple; Horty presents it as *preliminary* (getting defeat right is an **open problem**)

- Horty (2012) presents the default logic as the "logic of reasons", but one has to be careful..

9 Exclusionary default theories

- An exclusionary default theory is a tuple $\langle \mathcal{W}, \mathcal{D}, < \rangle$ where:
  $\mathcal{W}$ is a set of propositional formulas, and
  $\mathcal{D}$ is a set of default rules.
- Default rules can be of two types:
  - Ordinary defaults of the form $X \to Y$
  - *Exclusionary defaults* of the form $X \to Out(d)$

  $<$ is a preorder on $\mathcal{D}$
- We extend the background language with a new predicate $Out(\cdot)$ and rule names $d_1, d_2, \ldots$

  Intuitively, $Out(d_1)$ says that default δ_1 is excluded or taken out of consideration

9.1 Example (Blue Lights)

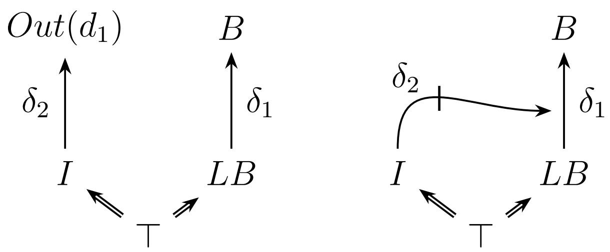

- $\Delta_1 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ where
  $\mathcal{W} = \{LB\}$
  $\mathcal{D} = \{\delta_1, \delta_2\}$
  - $\delta_1 = LB \to B$
  - $\delta_2 = I \to Out(d_1)$

  $< = \emptyset$
- $\Delta_2 = \langle \mathcal{W}', \mathcal{D} \rangle$ where
  $\mathcal{W}' = \mathcal{W} \cup \{I\}$
  $< = \emptyset$
  (LB = Looks blue; B = Is blue;
  I = Illuminated by blue lights)
- Assumption: a weaker rule can exclude a stronger one

  (There's an **open problem** having to do with this assumption. We leave it aside..)

9.2 Example (Blue Lights c'd)

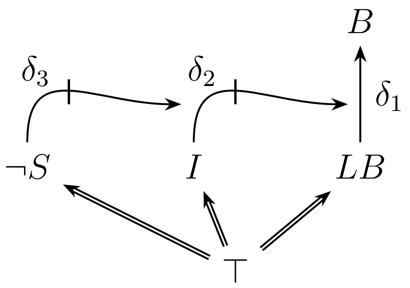

- $\Delta_3 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ where
  $\mathcal{W} = \{LB, I, \neg S\}$
  $\mathcal{D} = \{\delta_1, \delta_2, \delta_3\}$
  - $\delta_1 = LB \to B$
  - $\delta_2 = I \to Out(d_1)$
  - $\delta_3 = \neg S \to Out(d_2)$

  $< = \emptyset$
- So there can be exclusion of exclusion. Intuitively, it should be safe to rely on δ_1..

9.3 Exclusion

- $Excluded_{\langle \mathcal{W}, \mathcal{D}, < \rangle}(\mathcal{S}) = \{\delta \in \mathcal{D} : \mathcal{W} \cup Conclusion(\mathcal{S}) \vdash Out(d)\}$

9.4 Stable* scenarios

- A scenario $\mathcal{S}$ based on $\langle \mathcal{W}, \mathcal{D}, < \rangle$ is stable* just in case:

$$\begin{aligned} \mathcal{S} \quad = \quad \{\delta \in \mathcal{D} : \quad & \delta \in Triggered_{\langle \mathcal{W}, \mathcal{D}, < \rangle}(\mathcal{S}), \\ & \delta \notin Conflicted_{\langle \mathcal{W}, \mathcal{D}, < \rangle}(\mathcal{S}), \\ & \delta \notin Defeated_{\langle \mathcal{W}, \mathcal{D}, < \rangle}(\mathcal{S}), \\ & \delta \notin Excluded_{\langle \mathcal{W}, \mathcal{D}, < \rangle}(\mathcal{S})\}. \end{aligned}$$

9.5 Optional problem set

10 Peer disagreement and conciliatory views

10.1 Peer disagreement debate

- Core problem: You've thought hard about some nontrivial question and arrived at the belief that X. Then you find out that someone you consider an *epistemic peer* is convinced that $\neg X$. How should you adjust your beliefs?
  - The core question is normative
  - Who makes for an epistemic peer?

- Main views:
  - Conciliatory views: Back off from your belief that X!
  - Steadfast views: You can stick to your belief
- This is in terms of all-or-nothing attitudes; in the literature, these views are typically formulated in terms of *credences* / *degrees of belief* / *degrees of confidence* /
- I think of the core question from a first-personal perspective

10.2 Sample case

- **Mental Math**. My friend Megan and I have been going out to dinner for many years. We always tip 20% and divide the bill equally, and we always do the math in our heads. We're quite accurate, but on those occasions where we've disagreed in the past, we've been right equally often. This evening seems typical, in that I don't feel unusually tired or alert, and neither my friend nor I have had more wine or coffee than usual. I get $43 in my mental calculation, and become quite confident of this answer. But then Megan says she got $45. I dramatically reduce my confidence that $43 is the right answer. (Christensen, 2010)

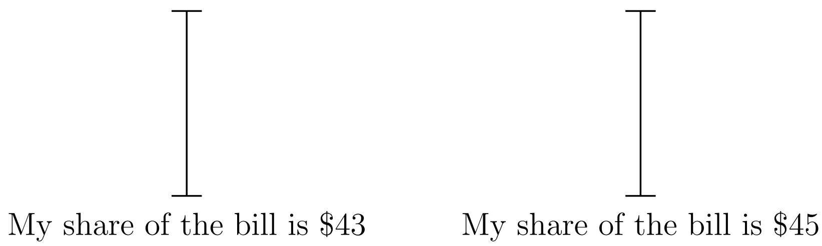

- Many other examples in the literature..
- Why we might (should) worry about the epistemic effects of peer disagreement?
- Interview with D. Christensen from Brown University

10.3 Higher-order disagreements

- Conciliatory views seem intuitive, but
- run into trouble in scenarios involving *higher-order disagreements*:
  (1) disagreements over the peerhood status of an apparent epistemic peer (Mulligan 2015)
  (2) disagreements over conciliatory views themselves (Elga 2010)

..

Tools for Formal Epistemology: Doxastic Logic, Probability, and Default Logic

– ESSLLI 2023 –

Problem Set (Default logic)

Task

Complete as many of the exercises as you like. (Exercise 7 will require more work than others.)

Exercise 1 (Inference graphs)

Draw inference graphs depicting the following (fixed priority) default theories:

(a) $\Delta_1 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ where:
$\mathcal{W} = \{B, C \supset D\}$;
$\mathcal{D} = \{\delta_1, \delta_2, \delta_3\}$ with $\delta_1 = B \rightarrow C$, $\delta_2 = D \rightarrow E$, and $\delta_3 = E \rightarrow \neg B$;
$< = \emptyset$

(b) $\Delta_2 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ where:
$\mathcal{W} = \{A\}$;
$\mathcal{D} = \{\delta_1, \delta_2, \delta_3, \delta_4, \delta_5\}$ with $\delta_1 = A \rightarrow B$, $\delta_2 = B \rightarrow C$, $\delta_3 = C \rightarrow D$, $\delta_4 = B \rightarrow \neg D$, and $\delta_5 = A \rightarrow D$;
$\delta_3 < \delta_4 < \delta_5$

(c) $\Delta_3 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ where:
$\mathcal{W} = \{A, B\}$;
$\mathcal{D} = \{\delta_1, \delta_2, \delta_3, \delta_4\}$ with $\delta_1 = A \rightarrow C$, $\delta_2 = B \rightarrow \neg C$, $\delta_3 = C \rightarrow D$, $\delta_4 = \top \rightarrow \neg D$;
$< = \emptyset$

(d) $\Delta_4 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ where:
$\mathcal{W} = \{A, B, (D \wedge B) \supset E\}$;
$\mathcal{D} = \{\delta_1, \delta_2, \delta_3\}$ with $\delta_1 = A \rightarrow C$, $\delta_2 = C \rightarrow D$, and $\delta_3 = E \rightarrow \neg C$;
$< = \emptyset$

(e) $\Delta_5 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ where:
$\mathcal{W} = \emptyset$
$\mathcal{D} = \{\delta_1, \delta_2\}$ with $\delta_1 = A \rightarrow A$, $\delta_2 = \top \rightarrow B$;
$< = \emptyset$

(f) $\Delta_6 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ where:
$\mathcal{W} = \emptyset$;
$\mathcal{D} = \{\delta_1, \delta_2, \delta_3\}$ with $\delta_1 = \top \rightarrow A$, $\delta_2 = A \rightarrow \neg A$, and $\delta_3 = \top \rightarrow B$;
$\delta_1 < \delta_2$

Recall that an atomic formula A from the hard information $\mathcal{W}$ is to be treated as $\top \supset A$. Also, recall that there can be only one node representing the true proposition $\top$.

Exercise 2 (Interpretation)

Come up with a plausible interpretation for the default theories Δ_2 and Δ_3. (Hint: You might want to think about the Tweety Triangle and the Nixon Diamond.)

Exercise 3 (Stable scenarios)

Determine the stable scenarios of the default theories Δ_1–Δ_6 from Exercise 1.

Exercise 4 (Proper scenarios)

Find all proper scenarios of the default theories from Exercise 1, using approximating sequences.

Exercise 5 (Extensions)

Determine the extensions of default theories from Exercise 1.

Example: The extension of the Tweety Triangle (Δ_1 in the slides) is

$$\mathcal{E} = Th(\mathcal{W} \cup Conclusion(\{\delta_2\})) = Th(\mathcal{W} \cup \{\neg F\}) = Th(\{P, P \supset B\} \cup \{\neg F\}) = Th(\{P, P \supset B, \neg F\})$$

Exercise 6 (Exclusion)

(a) How would you represent the information from the Tweety Triangle in a default theory with an empty priority ordering? What's lost (or gained)?

(b) The addition of exclusionary default rules and the corresponding notion of exclusion prompted us to modify the notion of stable scenarios. One important adjustment is still missing from the slides. What is it?

Exercise 7 (Proofs)

Prove some of the facts listed in Section 6 of Lecture 4.

Tools for Formal Epistemology:
Doxastic Logic, Probability, and Default Logic

– ESSLLI 2023 –

Problem Set (Default logic)

Task

Complete as many of the exercises as you like. (Exercise 7 will require more work than others.)

Exercise 1 (Inference graphs)

Draw inference graphs depicting the following (fixed priority) default theories:

(a) $\Delta_1 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ where:
$\mathcal{W} = \{B, C \supset D\}$;
$\mathcal{D} = \{\delta_1, \delta_2, \delta_3\}$ with $\delta_1 = B \to C$, $\delta_2 = D \to E$, and $\delta_3 = E \to \neg B$;
$< = \emptyset$

(b) $\Delta_2 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ where:
$\mathcal{W} = \{A\}$;
$\mathcal{D} = \{\delta_1, \delta_2, \delta_3, \delta_4, \delta_5\}$ with $\delta_1 = A \to B$, $\delta_2 = B \to C$, $\delta_3 = C \to D$, $\delta_4 = B \to \neg D$, and $\delta_5 = A \to D$;
$\delta_3 < \delta_4 < \delta_5$

(c) $\Delta_3 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ where:
$\mathcal{W} = \{A, B\}$;
$\mathcal{D} = \{\delta_1, \delta_2, \delta_3, \delta_4\}$ with $\delta_1 = A \to C$, $\delta_2 = B \to \neg C$, $\delta_3 = C \to D$, $\delta_4 = \top \to \neg D$;
$< = \emptyset$

(d) $\Delta_4 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ where:
$\mathcal{W} = \{A, B, (D \wedge B) \supset E\}$;
$\mathcal{D} = \{\delta_1, \delta_2, \delta_3\}$ with $\delta_1 = A \to C$, $\delta_2 = C \to D$, and $\delta_3 = E \to \neg C$;
$< = \emptyset$

(e) $\Delta_5 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ where:
$\mathcal{W} = \emptyset$
$\mathcal{D} = \{\delta_1, \delta_2\}$ with $\delta_1 = A \to A$, $\delta_2 = \top \to B$;
$< = \emptyset$

(f) $\Delta_6 = \langle \mathcal{W}, \mathcal{D}, < \rangle$ where:
$\mathcal{W} = \emptyset$;
$\mathcal{D} = \{\delta_1, \delta_2, \delta_3\}$ with $\delta_1 = \top \to A$, $\delta_2 = A \to \neg A$, and $\delta_3 = \top \to B$;
$\delta_1 < \delta_2$

Recall that an atomic formula A from the hard information $\mathcal{W}$ is to be treated as $\top \supset A$. Also, recall that there can be only one node representing the true proposition $\top$.

Exercise 2 (Interpretation)

Come up with a plausible interpretation for the default theories Δ_2 and Δ_3. (Hint: You might want to think about the Tweety Triangle and the Nixon Diamond.)

Exercise 3 (Stable scenarios)

Determine the stable scenarios of the default theories Δ_1–Δ_6 from Exercise 1.

Exercise 4 (Proper scenarios)

Find all proper scenarios of the default theories from Exercise 1, using approximating sequences.

Exercise 5 (Extensions)

Determine the extensions of default theories from Exercise 1.

Example: The extension of the Tweety Triangle (Δ_1 in the slides) is

$$\mathcal{E} = Th(\mathcal{W} \cup Conclusion(\{\delta_2\})) = Th(\mathcal{W} \cup \{\neg F\}) = Th(\{P, P \supset B\} \cup \{\neg F\}) = Th(\{P, P \supset B, \neg F\})$$

Exercise 6 (Exclusion)

(a) How would you represent the information from the Tweety Triangle in a default theory with an empty priority ordering? What's lost (or gained)?

(b) The addition of exclusionary default rules and the corresponding notion of exclusion prompted us to modify the notion of stable scenarios. One important adjustment is still missing from the slides. What is it?

Exercise 7 (Proofs)

Prove some of the facts listed in Section 6 of Lecture 4.