Notes for Meeting 6
Abductive Inference

Review of Deductive Reasoning

Much AI work on rule-based reasoning focuses on deductive inference:

- Given: A set of inference rules and a set of facts

- Find: Proofs that derive beliefs which follow deductively.

- Most approaches to this task reason backward from a query.

Deductive inference is important for many tasks, and it underlies logic programming languages like Prolog.

But such reasoning is not the only type that arises in humans or the only form we need for AI systems.

A Motivating Example

You come into class and realize that student X has been choked with the projector cord. You know that X always comes to class early. Moreover, this instructor Y prides herself on the difficulty of her exams, grades on a curve, and X always scores the highest on them. You also know that student Z has applied for a scholarship that requires near perfect grades, but that Z has not been doing especially well in this one class.

(a) Did X commit suicide? Why is this implausible?

(b) Did Z murder X? What motive might he have had?

(c) Did Y murder X? What motive might she have had?

Did you arrive at this conclusion by deduction? Does your reasoning constitute a proof of the conclusion you reached?

Some Other Examples

We encounter non-deductive reasoning frequently in both life and art:

- When you see a doctor about a medical problem, how does the doctor reach a conclusion about what ails you?

- In Arthur Conan Doyle's stories about Sherlock Holmes, does the detective solve crimes through deduction (as the author claims)?

- In the television show CSI, how does the investigative unit decide who committed the crime?

- When astronomers discovered pulsars, how did they explain their unusual behavior? Did they consider different explanations?

Each of these examples involve ABDUCTION (Peirce, 1878), a form of reasoning with important differences from deduction.

The Task of Abductive Inference

We can specify the generic task of abductive reasoning as:

- Given: A set of inference rules (or other knowledge elements) and a set of observed facts

- Find: One or more EXPLANATIONS of the observed facts in terms of the knowledge and observations.

An explanation is always stated in terms of something already known.

Some define abduction as inference to the best explanation, but this comes from a misguided focus on optimality.

Even if we could define "best", finding it may be intractable.

Applications of Abductive Inference

The abductive reasoning task arises in many different AI settings:

- Medical and mechanical diagnosis

- Natural language understanding

- Plan understanding (e.g., in games)

- Scientific reasoning and discovery

However, abduction has received relatively little attention with the AI community, especially in recent years.

Issues in Abduction

Consider again the classroom scenario with the asphyxiated student:

- Why do we feel the need to explain the cause of death? Why not explain what the student was wearing that day?

- What knowledge is needed to construct an explanation? What form might such knowledge take?

- Do explanations take the same form as deductive proofs? Do they constitute valid deductive proofs?

- Why does one explanation seems better than another? What criteria come into play?

- What mechanisms can one use to generate candidates? Does it make sense to carry out exhaustive search?

As Leake (1995) notes, a complete account approach of abduction must address each of these questions.

Heuristics for Plausible Reasoning

Abductive inference relies on making assumptions that do not follow deductively from the givens; it concerns PLAUSIBLE reasoning.

But this requires some way to determine which assumptions are more plausible than others; some candidates are:

- preferring simpler explanations

- preferring more probable accounts

- preferring more coherent explanations

Coherence can produce explanations that are similar to those humans prefer (Ng & Mooney, 1990), although probabilities may modulate it.

Backward-Chaining Approaches to Abduction

A common approach to abduction involves adapting backward-chaining methods for deduction; such methods:

- start from a fact that needs to be explained

- chain through rule consequents that unify with this fact

- unify rule antecedents with facts when possible

  - chain off unmatched antecedents when not OR

  - make default assumptions for unmatched antecedents

- continue until producing a "proof" of the original fact
  that terminates in other facts or assumptions

This process involves a query-driven, AND-OR search through the space of explanations.

Academic Knowledge, Facts, and Inferences

Background knowledge:

```
(happy ?x) <= (optimist ?x)

(happy ?x) <= (succeed ?x ?y)

(succeed ?x ?y) <= (exam ?y) (easy ?y) (study ?x ?y) (take ?x ?y)
```

Observed facts:

```
(name j john) (happy j) (exam e) (easy e)
```

Plausible inferences:

```
(succeed j e) (study j e) (take j e)
```

As Ng and Mooney (1990) note, this explanation is more coherent than a simpler one that assumes (optimist j).

Case-Based Approaches to Abduction

Leake (1985) reviews a different approach to abduction associated with the paradigm of case-based reasoning that:

- encodes knowledge as "cases" rather than as general rules;

- generates accounts that do not correspond to proof trees;

- attempts to find explanations only when anomalies occur; and

- produces explanations by adapting cases that fail to handle
  these anomalies.

This framework is legitimate, but his analysis confounds issues like plausible reasoning with case-based representations.

Another Approach to Abduction

Bridewell and Langley (2011) report another approach to abduction.

Like other work, it uses a form of plausible reasoning to generate explanations and guides search by a coherence metric.

But their framework differs from earlier techniques in that it:

- accepts new facts as they arrive, forming explanations in an
  incremental fashion;

- operates in a flexible manner, chaining both backward and forward
  over its rules; and

- calculates coherence locally, ensuring that inference will scale.

They claim this approach provides a reasonable account of everyday inference in humans.

Some Medical Knowledge

Background knowledge:

```
(has-flu ?person) <= (has-symptom ?person ?s1) (fever ?s1)
  (has-symptom ?person ?s2) (cough ?s2)

(has-food-poisoning ?person) <= (has-symptom ?person ?s1)
  (fever ?s1) (has-symptom ?person ?s2) (vomiting ?s2)

(has-lung-cancer ?person) <= (has-symptom ?person ?s1)
  (cough ?s1) (has-symptom ?person ?s2) (yellow-teeth ?s2)

(caught-flu ?person1 ?person2) <= (at-meetings ?person1 ?project)
  (has-flu ?person1) (at-meetings ?person2 ?project) (has-flu ?person2)

(project ?project) <= (member-of ?person1 ?project)
  (paid-from ?person1 ?project) (at-meetings ?person1 ?project)
```

Some Medical Facts and Inferences

Observed facts:

```
(member-of Ann p1) (member-of Bob p1)
(has-symptom Ann s1) (fever s1)
(has-symptom Bob s2) (cough s2)
```

Plausible inferences:

```
(has-flu Ann) (has-flu Bob)
(at-meetings Ann p1) (at-meetings Bob p1)
(has-symptom Bob s3) (fever s3)
(has-symptom Ann s4) (fever s4)
(caught-flu Ann Bob)
```

This set of inferences is more coherent than ones that assume Ann has food poisoning and Bob has lung cancer.

Note that this explanation does NOT take the form of a proof tree.

The Challenge of Belief Revision

Humans operate over time, so that facts they learn later may cause inferences based on earlier ones to become less plausible.

- Consider the sentences "John needed money. He got his gun." These might lead one to think John plans to rob a store or bank.

- Now consider the sentence "He drove to the shooting contest." This changes the picture considerably.

A robust abductive system must be able to retract its assumptions through a process of BELIEF REVISION.

- Work on truth maintenance systems addresses this issue by storing dependencies among beliefs.

- But detecting and responding to inconsistencies both incrementally and efficiently is still an open problem.

Assignments for Meeting 7
Analogical Reasoning

Read the articles:

- Falkenhainer, B., Forbus, K., & Gentner, D. (1986). The Structure-Mapping Engine. Proceedings of the Fifth National Conference on Artificial Intelligence. Philadelphia: Morgan Kaufmann. [required]

- Sowa, J. F., & Majumdar, A. K. (2003). Analogical reasoning. In A. Aldo, W. Lex, & B. Ganter, (Eds)., Conceptual structures for knowledge creation and communication. Berlin: Springer-Verlag. [optional]

- Gentner, D., & Forbus, K. (1991). MAC/FAC: A model of similarity-based retrieval. Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society (pp. 504-509). [optional]

- Complete the second exercise (due 11:59 PM on 2/9/2011).