CSE 513T Theory of Artificial Intelligence & Machine Learning March 8, 2018

Lecture 16

Lecturer: Brendan Juba *Scribe: Han Liu*

Contents

1 "Common sense" knowledge & reasoning **1**

2 Nonmonotonic reasoning with filtered examples **2**

3 Issues with representation of common sense knowledge **3**
3.1 Qualification problem 3
3.2 Elaboration Tolerance 4
3.3 Ramification Problem 4
3.4 Size of common sense 4

4 Shortcomings of reasoning with examples **4**

1 "Common sense" knowledge & reasoning

Distinguish two kinds of knowledge: (Shaded area for common sense)

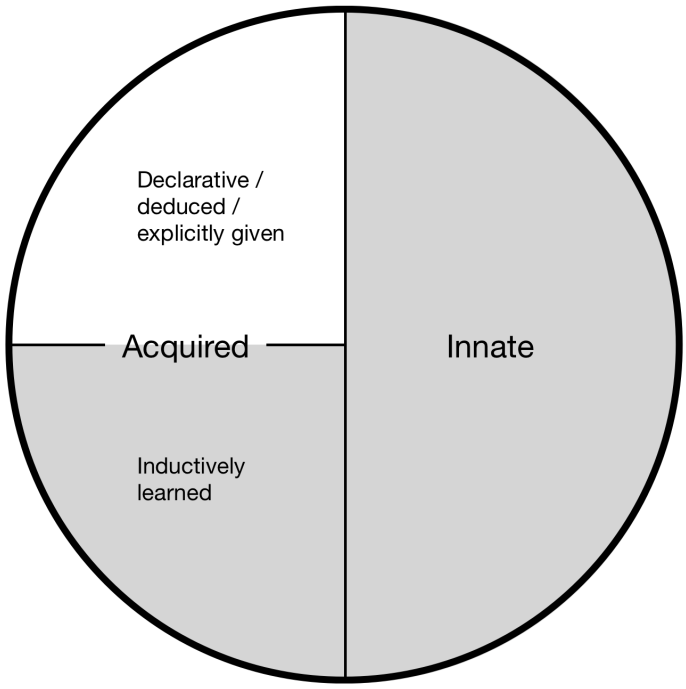


Figure 1: Categories of knowledge

Knowledge

- Acquired
    - Declarative/deduced/explicitly given
    - Inductively learned (common sense)
- Innate (common sense)

Common sense is a varitey of tacitly shared knowledge. We rely on in action, communication, etc. How much of "common sense" falls in each category?

Cognitive Psychology identifies several domains as innate including (mostly not addressed by reasoning with examples):

- Space and objects (persistence, identity)
- Naive physics ← e.g. based on extra-logical "simulations"?
- Naive psychology (intentions, emotions) ← e.g. may not be learnable at all - rich?
- Quantitative reasoning ("more" or "less", small #s)

These domains are not usually present at birth, but developed according to some standard schedule across individuals. We're currently focusing on inductively learned knowledge. e.g. natural kinds (birds, games, etc.) ← reasoning with (filtered) examples

2 Nonmonotonic reasoning with filtered examples

Suppose we wish to reason about "birds" (in general), we keep the subset of our examples with bird(x)=1, and filtering out the rest gives such a set of ex's.

	bird	fly	has beak	red	purple	penguin
$x^{(1)}$	1	1	1	0	0	0
$x^{(2)}$	1	1	1	1	0	0
$x^{(3)}$	1	0	1	0	0	1
$x^{(4)}$	1	0	1	0	0	1
$x^{(5)}$	1	1	1	0	0	0
$x^{(6)}$	1	1	1	1	0	0
$x^{(7)}$	1	1	1	0	0	0

Table 1: Truth table of examples

This filtering operation effectively simulates a set of ex's drawn from the conditional distribution:

$$D|bird = \Pr_{x\in D}[y|bird(y)] = \frac{\Pr_{y\in D}[y \wedge bird]}{\Pr_{x\in D}[bird(y)]}$$

Probabilities proportional to prob under D (for bird=1) scaled up by $1/\Pr[bird]$.

Suppose "error tolerance" $\epsilon = 1/3$. Try reasoning about examples:

$$\hat{\Pr}[fly(x)] = \frac{1}{7}\sum_{i=1}^{7} fly(x^{(i)}) = 5/7 \geq 2/3 = 1-\epsilon \Rightarrow \text{accept}$$

$$\hat{\Pr}[has\ beak] = \frac{1}{7}\sum_{i=1}^{7} has\ beak(x^{(i)}) = 1 \geq 2/3 = 1 - \epsilon \Rightarrow \text{accept}$$

$$\hat{\Pr}[red] = \frac{1}{7}\sum_{i=1}^{7} red(x^{(i)}) = 2/7 \leq 2/3 = 1 - \epsilon \Rightarrow \text{reject}$$

Similarly, $\hat{\Pr}[purple] = 0 \Rightarrow$ reject, and $\hat{\Pr}[penguin] = 2/7 \Rightarrow$ reject. These might possibly characterize our intuitive beliefs about birds.

Now suggest we are interest in red birds. Then what can we reason about $bird \wedge red$? (Focus on $x^{(2)}$ and $x^{(6)}$)

Fly?	$\hat{\Pr}[bird \wedge red \wedge fly] = 1 \geq 2/3 = 1 - \epsilon \Rightarrow$ accept
Has beak?	$\hat{\Pr}[bird \wedge red \wedge has\ beak] = 1 \geq 2/3 = 1 - \epsilon \Rightarrow$ accept
Red?	Conditioned on red. $\Rightarrow$ accept.
Purple?	$\hat{\Pr}[bird \wedge red \wedge purple] = 0 \Rightarrow$ reject
Penguin?	$\hat{\Pr}[bird \wedge red \wedge penguin] = 0 \Rightarrow$ reject

Now how about $bird \wedge penguin$? (Focus on $x^{(3)}$ and $x^{(4)}$)

Fly?	$\hat{\Pr}[bird \wedge penguin \wedge fly] = 0 \Rightarrow$ reject
Has beak?	$\hat{\Pr}[bird \wedge penguin \wedge has\ beak] = 1 \Rightarrow$ accept
Red?	$\hat{\Pr}[bird \wedge penguin \wedge red] = 0 \Rightarrow$ reject
Purple?	$\hat{\Pr}[bird \wedge penguin \wedge purple] = 0 \Rightarrow$ reject
Penguin?	Conditioned on penguin. $\Rightarrow$ accept.

Further restrict to $bird \wedge penguin \wedge has\ beak$? Same examples. Adding conditions in this case doesn't affect our result.

Finally, reason about $bird \wedge penguin \wedge purple$: No examples, only accept tautologies. Otherwise refrain from drawing conclusions.

As long as $\epsilon > 0$ is allowed, we can get "leaping to conclusions" - birds fly, etc. As a consequence, also we get nonmonotonicity of additional knowledge.

3 Issues with representation of common sense knowledge

3.1 Qualification problem

Originally from reasoning about actions: essentially impossible to list out all of the conditions under which some rule holds.

e.g. bird $\Rightarrow$ fly ... except not penguins/emus/ostriches

- [bird $\wedge$ $\neg$penguin $\wedge$ $\neg$emu $\wedge$ $\neg$ostrich] $\Rightarrow$ fly
  Other conditions: healthy (has spine, no broken wing, alive)
- [bird $\wedge$ $\neg$penguin $\wedge$ $\neg$emu $\wedge$ $\neg$ostrich $\wedge$ has spine $\wedge$ $\neg$broken wing $\wedge$ alive] $\Rightarrow$ fly
  Other conditions: $\neg$baby, $\neg$chained down, $\neg$wings clipped, air pressure, reasonable gravity

Must tolerate exceptions due to unanticipated and unmodeled effects.

3.2 Elaboration Tolerance

Classical representations often require editing an entire KB to incorporate knowledge of new exceptional circumstances. By contrast, we'd like to handle exceptions just by making note of them (coping/responding to exceptions shouldn't require $\Omega(|KB|)$ time...)

E.g. in reasoning with examples, we just add examples to our set (first penguin), conditioning on property $(D|penguin)$ of exceptional example immediately draws from appropriate conditional (obtains ¬fly) distribution. (Didn't amend "bird⇒fly")

3.3 Ramification Problem

Again, arises in reasoning about actions. One small additional piece of knowledge can have many consequences.

E.g. load a box on to a truck, truck drives to Milwaukee ⇒ box is also in Milwaukee etc. Must be able to infer all such consequences efficiently.

Comes "for free" in reasoning with examples since the entailed properties are true on the unfiltered examples, we can just read them off.

3.4 Size of common sense

Various estimates of the size of human common sense: 0.5×10^9 bits - 3.4×10^9 bits or $10^7 - 10^8$ "elements of knowledge". Curating such a KB by hand is daunting...

- Systems (KNEXT) proposed for mining knowledge from Wikipedia, blogs, etc. - mine 10^8 "factoids" in reasonable time. We're looking for mundane facts. "eat with your mouth", "need a container to hold water", etc.
- Systems like CYC attempted this, initiated in 1983 at last count: 7×10^6 assertions about 5×10^6 items.
- Crowdsourcing: "Open Mind Common Sense" (OMCS) grown from 4.5×10^5 facts in 2002 to about 11×10^6 in 2014. - not going to reach desired size in our lifetime.

May wonder how is it possible for everyone to know so much?? Some is innate, but also explicitly stated rules may not be the best representation. Reasoning with examples gives an exponentially large set of query answers from a polynomial size set of examples.

4 Shortcomings of reasoning with examples

1 Assumes complete knowledge of examples.
If important attributes are not recorded in examples, can't generally evaluate queries.

Next couple lectures:

- extend to reasoning from partial examples
- we'll need to use deductive reasoinging to cope with missing information

2 Some common sense such as persistence or "laws of inertia" explicitly refer to unobserved things, therefore can't be learned from observations. After reasoning with partial examples, we'll see non-monotonic logics that explicitly represent such "leaping to conclusions" can represent "law of inertia" compactly.