CSE 513T Theory of Artificial Intelligence & Machine Learning March 19, 2015

Lecture 18

Lecturer: Brendan Juba *Scribe: Dingwen Li*

1 Today's Topics

1. Explicit vs. Implicit representations
2. Testable rules
3. Restriction-closed proofs
4. Reasoning with partial examples

2 Common Sense Reasoning

1. Reasoning with examples nicely captured reasoning about "typical" objects, scenes, etc.
   -but- required complete information.
2. Declarative representations seem necessary to cope with partial information, but may suffer from some pitfalls (most can be addressed ... next time).

3 Common Sense is "Big"

estimated size: $0.5 \times 10^9 - 3.4 \times 10^9$ bits.
or $10^7 - 10^8$ elements of knowledge.

The most prominent attempt to write down the contents of common sense knowledge:
CYC (1983 - present)
contains 7×10^6 assertions (about 5×10^6 terms). Still far from complete.
Other attempts crowd sourcing. OMCS.
from 4.5×10^5 facts in 2002 to 1.1×10^6 facts in 2014 (at this rate, might take 100 years).
K NEXT (scraping from blogs, wikipedia, ...).
contains 10^8 "factorials".
probably missing common sense facts, contains trivia.

Reasoning with examples can be a much more compact representation, doesn't require explicitly considering 10^7+ assertions and placing them in a KB. Seems, we need both.

4 Testable Rules

Declarative representations that are "implicitly represented" in partial examples.
More specifically assertions that are "demonstrably" $(1-\epsilon)$ valid on examples from $M(D)$, where M is the masking process, D is an arbitrary distribution.
Note, some rules may be tautologies (true on every partial example and this may be hard to verify (3 DNFs are NP-complete to check).

Recall: for a CNF φ, $\varphi|_\rho$ simplifies to 1 if for every clause C, $C|_\rho$ simplifies to 1; this happens if some literal in C is satisfied by ρ.

Notice: can compute $\varphi|_\rho$ in linear time.

Defn: we say that the representation φ is $(1-\epsilon)$-testable with respect to $M(D)$ if for $\rho \in M(D)$, $\varphi|_\rho = 1$ with probability $\geq 1-\epsilon$.

Ex: if we sample from $M(D|bird)$, and 99% of the examples, the partial examples have fly=1, then 'fly' is .99-testable (even if many other attributes are missing).
We'll use these testable rules as the "implicitly specified" knowledge, together with declarative representations.

Need: proofs that behave well under partial information.

5 Restriction-Closed Proofs

Defn: We say that a set of proofs P for a proof systems is closed under restriction if, whenever there is a proof $\pi = \pi_1, ..., \pi_l$ in D of a representation φ from premises $\psi_1, ..., \psi_k$ then, for every partial assignment ρ, there is also a proof of $\varphi|_\rho$ from $\psi_1|_\rho, ..., \psi_n|_\rho$, whenever none of these simplify to 'O'(for a CNF, when some clause has all literals falsified by ρ).

Claim from last week: clause space S treelike resolution proofs are closed under restrictions (H-W4).

Also, length-bounded [treelike] resolution is closed under restrictions.
Most, "natural" sets of proofs are closed under restrictions.

Claim: Chaining proofs are closed under restrictions.

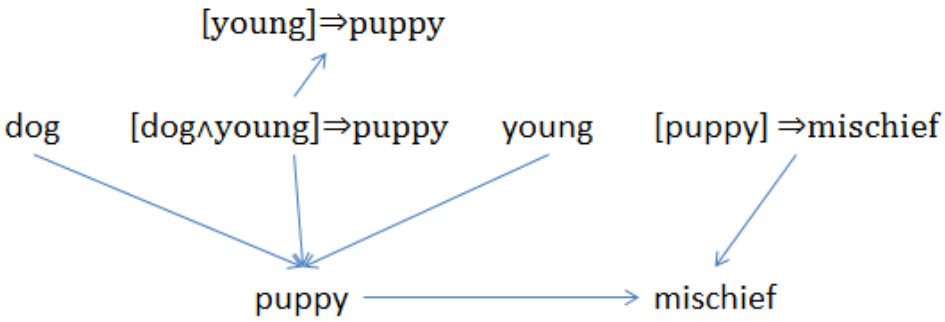
Pf: Consider a chaining proof $\pi = \pi_1, ..., \pi_l$ of φ from $\psi_1, ..., \psi_k$ and any partial assignment ρ. Consider the sequence $\pi|_\rho = \pi_1|_\rho, ..., \pi_l|_\rho$, we next omit any portion of the graph of π leading to some π_i where $\pi_i|\rho$ is a constant. Call this modified sequence π'.

Recall: Chaining proofs consist of "facts" (literals) & "rules" of the form $[l_1 \wedge ... \wedge l_r] \Rightarrow l$ (which corresponds to $\neg l_1 \vee ... \vee \neg l_r \vee l$). So under ρ, when an l_1 is set to 1, it gets dropped from the body. If l_1 is set to 0, then the rule simplifies to 1. Either we obtain a Boolean constant, another rule, or a single literal, ie, a "fact" under ρ. So π' also is a sequence of "facts" & "rules". We argue by induction on the number of steps of π' that we have a chaining derivation of all steps up to i.
Base case i=0 is done.

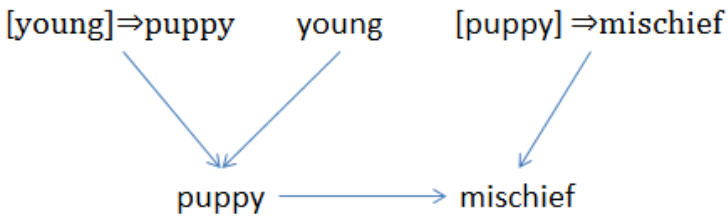
Induction Step: Consider whether the i^{th} representation was originally either among $\psi_1, ...\psi_k$ or was obtained by chaining. In the first case, as long as the $\psi_i|_\rho$ is not a constant it can still appear, and we have the premise $\psi_i|_\rho$ available. In the second case, we know that there was a rule $[l_1 \wedge ... \wedge l_r] \Rightarrow l$ and facts $l_1, ..., l_r$ derived in some previous steps. We know, since this step survived in π' that l was not set to a Boolean constant by ρ. So either the rule $[l_1 \wedge ... \wedge l_r] \Rightarrow l$ simplified itself to l in which case we can apply the same inference to derive l again, or for some subset of $l_1, ..., l_r$, that are not fixed to constants under ρ, which were derived earlier in π, and hence are derived in π'. So the chaining inference rule applied to $([l_1 \wedge ... \wedge l_r] \Rightarrow l)|_\rho$ and these surviving facts, we can derive π_i. This completes the IS.

Ex: We had a chaining proof of mischief.

Suppose first that a partial assignment sets. Dog and puppy to 1. Then the KB becomes young, mischief, so there is a trivial proof of mischief.



If only dog is set to 1, we still have the chaining proof.



6 Reasoning with Partial Examples

Suppose that we have an efficient algorithm for deciding whether or not a proof of a query φ from premise. $\psi_1, ..., \psi_k$ is contained in a set of restriction -closed proofs P. (e.g. our algorithm for treelike, space bounded resolution proofs or an algorithm for chaining or -).

Algorithm for reasoning with partial examples.
Input: $\varphi, \psi_1, ..., \psi_k$, partial examples $\rho^{(1)}, ..., \rho^{(n)}$

Initialize counter $s \leftarrow 0$
For i=1,...,n
    If $P(\varphi|_{\rho(i)}, \psi_1|_{\rho(i)}, ..., \psi_k|_{\rho(i)}) = 1$
    Then increment S
Return s/m

Theorem: For $m \geq \frac{1}{2r^2} \ln(\frac{2}{\delta})$ examples $(r, \delta \in (0,1))$ independently drawn from $M(D)$, then with probability $1-\delta$, the value $\hat{p}$ returned by the algorithm satisfies

1. $[\psi_1 \wedge ... \wedge \psi_k] \Rightarrow \varphi$ is $(\hat{p} - \delta)$-valid with respect to D.

2. If (a) there are "implicit" representation $\psi_{k+1}, ..., \psi_s$ s.t. $\psi_{k+1} \wedge ... \wedge \psi_s$ is (p+r)-testable (for some $p \in [0,1]$)

(b) p is closed under restrictions.
(c) There is a proof of φ from $\psi_1, ..., \psi_s$ in P, then $\hat{p} \geq p$.

7 Key Observation for Proof

The condition $\psi_{k+1} \wedge ... \wedge \psi_s$ is $(p + r)$-testable guarantees that with probability $\geq p + \delta$ over draws of partial examples ρ from $M(D)$, all of $\psi_{k+1}, ..., \psi_s$ simplify to 1 under ρ simultaneously. In this case, if we apply ρ to the proof of φ from $\psi_1, ..., \psi_s$,$\psi_{k+1}, ..., \psi_s$ drop out.