Implicitly Learning to Reason in First-Order Logic

Brendan Juba
Washington University in St. Louis
joint work with **Vaishak Belle**
University of Edinburgh & Alan Turing Institute

"*Implicit learning*": simulating learning without explicit representations

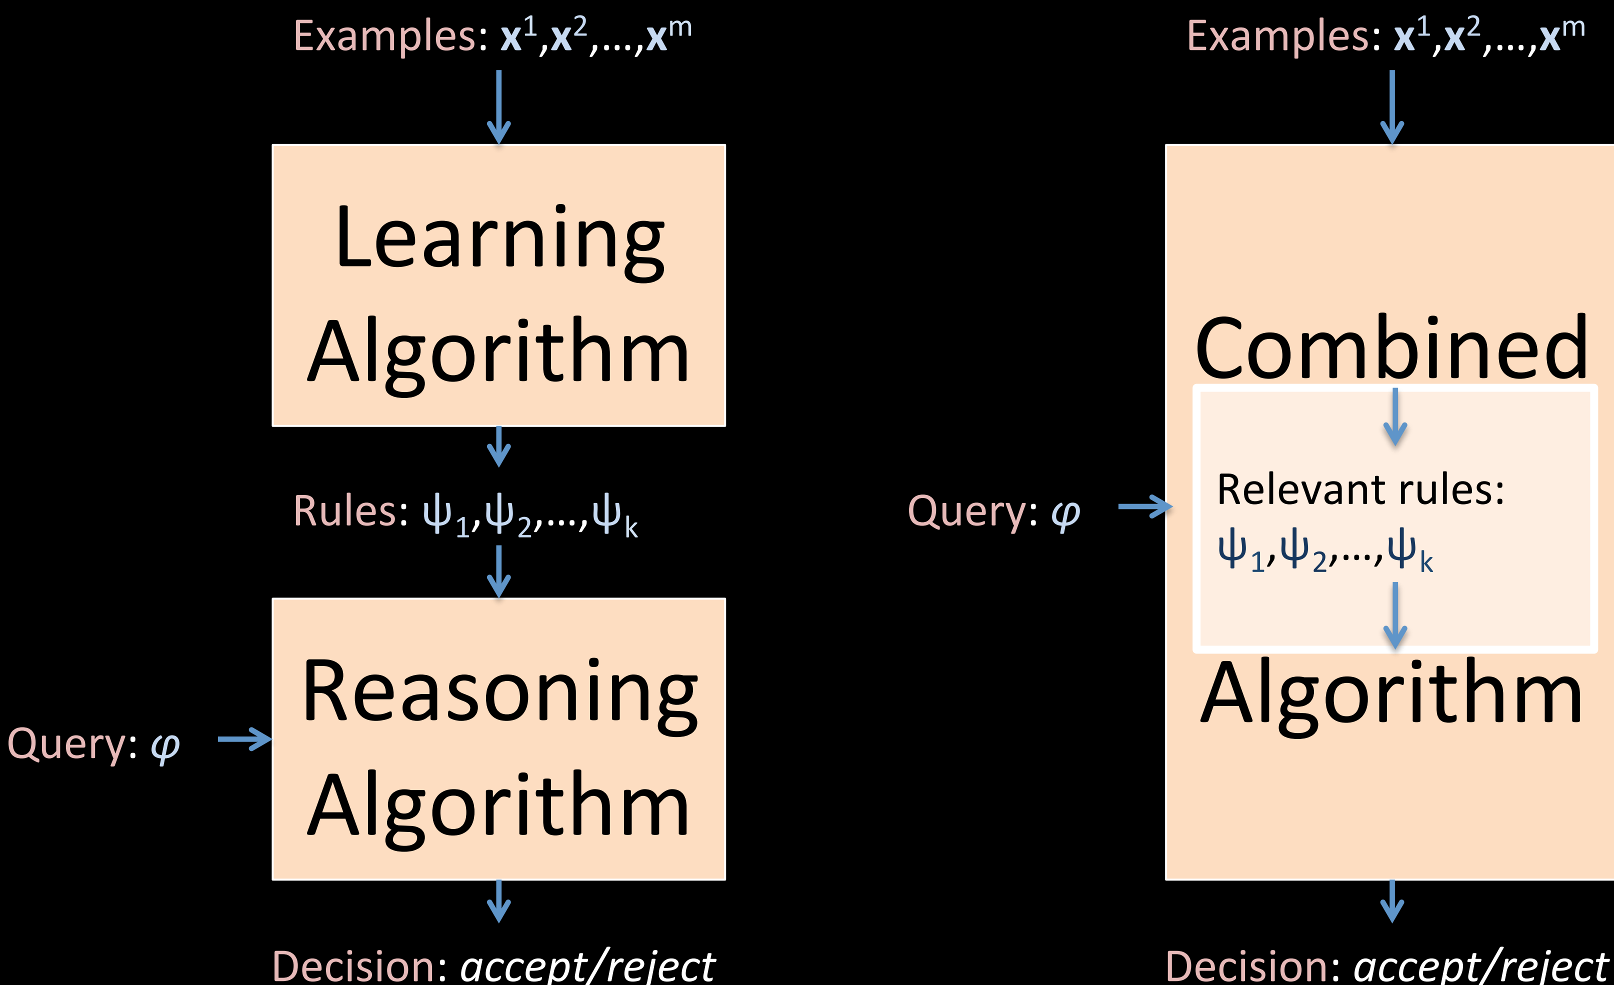

Why not use explicit representations?

- Often, **intractable** to guarantee that we discover all relevant rules
- *This work:* explicit representations are **impossible to learn.**

Language and *reasoning task*: universal clauses; ground clausal queries

- Language: First-Order Logic with equality, countably infinite domain of names (**N** wlog)
  - Variables x,y,z,…
  - Relation symbols P(x),…,Q(x_1,…,x_k),…
  - Usual connectives/quantifiers: $\wedge$, $\vee$, ¬, $\supset$, $\forall$, $\exists$
- Fragment: proper$^+$ KBs: finite set of $\forall$-clauses
  - Equality formulas: built over equality expressions of form "x = a" (*variable = name*) and $\wedge$, $\vee$, ¬
  - $\forall$-clause: $\forall$[e $\supset$ c] where e is an equality formula, c is a quantifier-free clause, $\forall$[] is universal closure
- Queries: ground clauses (OR of ground atoms)
  - Ground atoms: relations applied to names

Learning model: "Probably Approximately Correct"

- Suppose there exists an *arbitrary* probability distribution D on valuations of ground atoms
- Masking function θ: given valuation of ground atoms M, returns finite subset of the valuations N
- Suppose there exists an arbitrary masking process Θ: distribution on masking functions
- Given $N^1, N^2, \ldots, N^m$ drawn independently from Θ(D), ground clausal query φ, wish to certify φ is "1-ε valid" with high probability (over Θ(D) draw)
  - 1-ε valid: $\Pr_D[M \models \varphi] \geq 1-\varepsilon$ ($M \models \varphi$ means true on M)

Problem: nontrivial ∀-clauses **can't be learned**

- We'd like to *learn* a proper$^+$ KB – i.e., identify ∀-clauses that are 1-ε valid using partial valuations $N^1, N^2, \ldots, N^m$ from Θ(D).
- But, a ∀-clause ∀[e ⊃ c] either
  1. is equivalent to a ground clause (if c is trivial or e only permits a finite number of bindings)
  2. or else has an *infinite number of bindings* that must all be satisfied in the *full* valuation M

☞ In case 2, using (finite) partial N^i we *can't distinguish true* ∀-clauses from *false* ∀-clauses.

This work: *solution* using *implicit learning*

- We propose a new "testability" property that proper^{+} KBs may satisfy w.r.t. partial valuations.
- We describe a reduction of *learning and reasoning* to *classical reasoning*: using partial valuations, distinguish ground clausal queries φ
  - that are provable from a (implicit) testable proper^{+} KB
  - from those that are not 1-ε valid (thus: sound).
- Using, e.g., Liu et al. '04, obtain *polynomial-time learning and reasoning* for limited belief system

What's new?
Relationship to other work

- Sound learning and reasoning for proper^{+} KBs in infinite domains, with arbitrary distributions
  - Prior work on *learning to reason/reasoning in PAC-semantics* was **essentially propositional**, resorted to *propositionalization* for first-order
  - Work in *statistical relational learning* generally relies on *independence structure* in distributions (but produces more *explicit representations*)
  - *Inductive Logic Programming* treats input as *defining* a correct solution, rather than analyzing *predictive power* against unknown "ground truth"

OUR APPROACH

Key observation: the *"grounding trick"* (Levesque'98, Belle'07)

- Observation: Names not appearing in the KB or query behave identically

Theorem *(Belle'07)*: For a proper^{+} KB Δ and a ground clause φ, $\Delta \models \varphi$ iff $GND^{-}(\Delta \wedge \neg\varphi)$ is unsatisfiable.

- *Formally:* for a proper^{+} KB Δ, $GND^{-}(\Delta)$ is the set of all $c\theta$ for $\forall[e \supset c]$ in Δ such that $e\theta$ is valid and θ ranges over all variables in a set Z containing the names appearing in Δ plus rank(Δ) (arbitrary) additional names
  - rank(Δ): max # of quantified variables over clauses in Δ.

The *grounding trick* enables *evaluation* from partial valuations

- *Grounding trick*: As long as a partial valuation N^i gives values to a suitable set of names, we can check that a KB Δ entails a query φ.
- Witnessing: recursive evaluation on partial valuation N^i.
  - Propositional formulas: substitute partial valuation for atoms, $\varphi \vee \psi$ is witnessed true if either φ or ψ is; witnessed false if both are. (Other connectives similar.)
  - $\forall$-clause : $\forall x\, \varphi(x)$ is witnessed true for the set of names C if for all bindings of x to c from C, the propositional formula $\varphi(c)$ is witnessed true.

The *grounding trick* enables *evaluation* from partial valuations

- *Grounding trick*: As long as a partial valuation N^i gives values to a suitable set of names, we can check that a KB Δ entails a query *φ*.
- Witnessing: recursive evaluation on partial valuation N^i.
- Implicit KB I is *witnessed true* in N^i for a query *φ* and explicit KB Δ if for a set of names C containing all of the names appearing in I, Δ, and *φ* plus rank(Δ∧I) additional ones, every ∀-clause in I is witnessed true.
- Implicit KB I is *1-ε testable* for a query *φ* and explicit KB Δ if it is witnessed true with probability at least 1-ε on partial valuations from Θ(D).

Main Theorem

Theorem. For confidence δ, accuracy γ, and rank bound k, there is an algorithm that given a KB Δ, query *φ*, and $m \geq {}^{1}/_{2\gamma^2} \ln {}^{2}/_{\delta}$ partial valuations from Θ(D), returns an estimate of validity ṽ such that with probability at least 1-δ,

1. (*sound*) If $\Delta \supset \varphi$ is v–valid (w.r.t. D), $\tilde{v} \leq v+\gamma$
2. (*complete*) If there is an (implicit) KB I such that
   - $\Delta \wedge I \models \varphi$
   - Both I and Δ have rank at most k, and
   - I is v-testable for *φ* and Δ

   then $\tilde{v} \geq v-\gamma$.

Can compare ṽ to 1-ε to decide *"accept"*/*"reject"*

THE ALGORITHM & SKETCH OF ITS ANALYSIS

The algorithm (reduction to classical reasoning)

1. Initialize count = 0
2. Loop over partial valuations $N^1, N^2, \ldots, N^m$
   – Loop over k-tuples of names $c_1, \ldots, c_k$ from N^i not appearing in φ or Δ
     a. Construct Γ from $GND^-(\Delta \wedge \neg\varphi)$ using $\{c_1, \ldots, c_k\}$ as the additional names by recursively substituting truth values for subformulas witnessed in N^i
     b. If Γ is (detected) unsatisfiable, increment count and skip to next partial valuation N^{i+1}.
3. Return $\tilde{v}$ = count/m.

Sketch of analysis, condition 1 (*"soundness"*)

- When $\Delta \supset \varphi$ is falsified on complete valuation M drawn from D, $\Delta \wedge \neg \varphi$ is satisfied by M
- Therefore, it must also be satisfiable on any partial valuation N obtained from M, and in particular, satisfiable for any grounding.
- Therefore, the fraction of times Γ could be refuted is at most the fraction of times $\Delta \supset \varphi$ was satisfied on the actual valuations M.
- *Chernoff bound*: the observed fraction is greater than the true probability by at most γ.

Sketch of analysis, condition 2 (*"completeness"*)

- *Grounding trick*: $\Delta \wedge I \models \varphi$ iff $GND^-(\Delta \wedge I \wedge \neg\varphi)$ is unsatisfiable for any suitable choice of names.
- Substituting truth values for witnessed subformulas for any partial valuation N still yields an unsatisfiable formula.
- When I is witnessed true for the set of additional names $\{c_1,\ldots,c_k\}$, we'd substitute T for every clause of I in $GND^-(\Delta \wedge I \wedge \neg\varphi)$; the result is identical to the Γ we obtain.
- Therefore, when I is witnessed true, $\Delta \wedge I \models \varphi$ iff Γ is unsatisfiable.
- *Chernoff bound, again*: the observed fraction is less than the true probability (testability of I) by at most γ.

Recap: *learning and reasoning* for proper$^+$ KBs using *implicit learning*

- Obtain sound learning and reasoning for proper$^+$ KBs in infinite domains, with arbitrary distributions
- We proposed a new "testability" property that proper$^+$ KBs may satisfy w.r.t. partial valuations.
- We described a reduction of *learning and reasoning* to *classical reasoning*: using partial valuations, we distinguish ground clausal queries φ
  - that are provable from a (implicit) testable proper$^+$ KB
  - from those that are not 1-ε valid (thus: sound).
- Using, e.g., Liu et al. '04, obtain *polynomial-time learning and reasoning* for limited belief system

Future directions

- Queries on atoms with names that are rarely/ never observed
- Queries with quantifiers

Both require new assumptions for learning from partial valuations – perhaps *"bounded concealment"* (Michael '10).