



Logical Inference and Rule-based reasoning

AI & SW History (non-scholarly)

[Rule based reasoning](#) has a long history in AI and SW

- 1970s: [Prolog](#) programming language based on rules and used for many AI systems in Europe
- 1980s: the [expert systems](#) paradigm mostly used rule-based reasoning; Japan embraced Prolog for its [5th generation computing systems](#) push
- 2000s: semantic web promoted OWL, based on [description logic](#) KR rather than rules; [SWRL](#) rule system fails to become a W3C recommendation
- 2010s: various rule systems developed for RDF: [RIF](#), [RuleML](#), [SPIN](#), [Datalog](#)
- 2022: SWRL is still commonly supported and used

Automated inference for FOL

- Automated inference for FOL is harder than Propositional Logic
 - Variables can take on an infinite number of possible values from their domains
 - Hence there are potentially an infinite number of ways to apply the Universal Elimination rule
- Godel's Completeness Theorem says FOL entailment is only **semi-decidable**
 - If a sentence is **true** given a set of axioms, there is a procedure that will eventually determine this
 - If a sentence is **false**, there's no guarantee a procedure will ever discover this — it **may never halt**

Generalized Modus Ponens (GMP)

- Modus Ponens: $P, P \Rightarrow Q \models Q$
- Generalized Modus Ponens extends this to rules in FOL
- Combines And-Introduction, Universal-Elimination, and Modus Ponens, e.g.
 - given $P(c), Q(c), \forall x P(x) \wedge Q(x) \rightarrow R(x)$
 - derive $R(c)$
- Must deal with
 - more than one condition on rule's left side
 - variables

Often rules restricted to Horn clauses

- A Horn clause is a sentence of the form:

$$P_1(\dots) \wedge P_2(\dots) \wedge \dots \wedge P_n(\dots) \rightarrow Q(\dots)$$

where

- ≥ 0 P_i s and 0 or 1 Q
- P_i s and Q are positive (i.e., non-negated) literals
- The ... are a sequence of variables or literals
- Prolog and most rule-based systems are limited to Horn clauses
- Horn clauses are a subset of all FOL sentences

Horn clauses 2

- Special cases

- Typical rule: $P_1 \wedge P_2 \wedge \dots P_n \rightarrow Q$
- Constraint: $P_1 \wedge P_2 \wedge \dots P_n \rightarrow \text{false}$
- A fact: $\rightarrow Q$
- A goal: $Q \rightarrow$

- Examples

- $\text{parent}(P1,P2) \wedge \text{parent}(P2,P3) \rightarrow \text{grandparent}(P1,P3)$
- $\text{male}(X) \wedge \text{female}(X) \rightarrow \text{false}$
- $\rightarrow \text{male}(\text{john})$
- $\text{female}(\text{mary}) \rightarrow$

Horn clauses 3

- These are not Horn clauses:
 - $\text{married}(x, y) \rightarrow \text{loves}(x, y) \vee \text{hates}(x, y)$
 - $\neg \text{likes}(\text{john}, \text{mary})$
 - $\neg \text{likes}(x, y) \rightarrow \text{hates}(x, y)$
- Can't assert/conclude disjunctions (i.e., “or”)
- Can't have “true” negation
 - Though some systems, like Prolog, allow a negation operator that means “can't prove”
- No wonder Horn clause reasoning is easier

Horn clauses 3

- Where are the quantifiers?
 - Variables in conclusion **universally quantified**
 - Variables appearing only in premises **existentially quantified**
- Examples:
 - $\text{parentOf}(P,C) \rightarrow \text{childOf}(C,P)$
 $\forall P \forall C \text{parentOf}(P,C) \rightarrow \text{childOf}(C,P)$
 - $\text{parentOf}(P,X) \rightarrow \text{isParent}(P)$
 $\forall P \exists X \text{parent}(P,X) \rightarrow \text{isParent}(P)$
 - $\text{parent}(P1, X) \wedge \text{parent}(X, P2) \rightarrow \text{grandParent}(P1, P2)$
 $\forall P1,P2 \exists X \text{parent}(P1,X) \wedge \text{parent}(X, P2) \rightarrow \text{grandParent}(P1, P2)$

Definite Clauses

- A **definite clause** is a horn clause with a conclusion
- What's not allowed is a horn clause w/o a conclusion, e.g.
 - $\text{male}(x), \text{female}(x) \rightarrow$
 - i.e., $\sim \text{male}(x) \vee \sim \text{female}(x)$
- Most rule-based reasoning systems, like Prolog, allow only definite clauses in the KB

Limitations

- Most rule-based reasoning systems use only definite horn clauses
 - Limited ability to reason about **negation** and **disjunction**
 - Cannot have try negation in a rule
 - Cannot have rules whose conclusion is an or
- These are required for some problems (e.g., playing Clue)
- Benefit is **decidability** and **efficiency**
- Some limitations can be overcome by
 - Adding procedural components
 - Augmenting with other reasoners


Some non-limitations

- We can easily rewrite some non Horn clause rules to as several simpler Horn clauses
- Example, rewrite 1 as 2
 1. $\text{married}(x, y) \rightarrow \text{loves}(x, y) \wedge \text{knows}(x, y)$
 2. $\text{married}(x, y) \rightarrow \text{loves}(x, y)$
 $\text{married}(x, y) \rightarrow \text{knows}(x, y)$
- Example, rewrite 1 as 2
 1. $\text{loves}(x, y) \vee \text{likes}(x, y) \rightarrow \text{knows}(x, y)$
 2. $\text{loves}(x, y) \rightarrow \text{knows}(x, y)$
 $\text{likes}(x, y) \rightarrow \text{knows}(x, y)$
- Additional rewriting rules are available

Forward & Backward Reasoning

- We often talk about two reasoning strategies:
 - Forward chaining and
 - Backward chaining
- Both are equally powerful, but optimized for different use cases
- You can also have a mixed strategy
 - Each rule specifies how it's to be used

Forward chaining



- Proofs start with given axioms/premises in KB, deriving new sentences using GMP until the goal/query sentence is derived
 - Process follows a chain of rules and facts going from the KB to the conclusion
- This defines a **forward-chaining** inference procedure because it moves “forward” from the KB to the goal [eventually]
- Inference using GMP is **sound** and **complete** for KBs containing **only Horn clauses**

Forward chaining example

- KB:
 1. $\text{allergies}(X) \rightarrow \text{sneeze}(X)$
 2. $\text{cat}(Y) \wedge \text{allergicToCats}(X) \rightarrow \text{allergies}(X)$
 3. $\text{cat}(\text{felix})$
 4. $\text{allergicToCats}(\text{mary})$
- Goal:
 - $\text{sneeze}(\text{mary})$


Forward chaining example

- KB:
 1. $\text{allergies}(X) \rightarrow \text{sneeze}(X)$
 2. $\text{cat}(Y) \wedge \text{allergicToCats}(X) \rightarrow \text{allergies}(X)$
 3. $\text{cat}(\text{felix})$
 4. $\text{allergicToCats}(\text{mary})$
 5. #3 , #4 $\rightarrow \text{allergies}(\text{mary})$
 6. $\text{allergies}(\text{mary})$
 7. #6 , #1 $\rightarrow \text{sneeze}(\text{mary})$
 8. $\text{sneeze}(\text{mary})$
- Goal:
 - $\text{sneeze}(\text{mary})$

In practice...

- Most forward chaining systems compute and add to the KB **all** inferred facts
 - This is generally OK for Horn clause systems, as there will be a finite number of them
 - But could be a problem if the KB is large
- Another problem: suppose you want to remove a fact or rule?
 - You could delete all of the inferred facts and start over, or
 - Use a [truth maintenance system](#) to only remove inferred facts that depended on the deleted items

Backward chaining



- **Backward-chaining** deduction using GMP is also **complete** for KBs containing **only Horn clauses**
- Proofs start with the goal query, find rules with that conclusion, and then tries to prove each of the antecedents in the rule
- Keep going until you reach premises
- Avoid loops by checking if new subgoal is already on the goal stack
- Avoid repeated work: use a cache to check if new subgoal already proved true or failed

Backward chaining example

- KB:
 - $\text{allergies}(X) \rightarrow \text{sneeze}(X)$
 - $\text{cat}(Y) \wedge \text{allergicToCats}(X) \rightarrow \text{allergies}(X)$
 - $\text{cat}(\text{felix})$
 - $\text{allergicToCats}(\text{mary})$
- Goal:
 - $\text{sneeze}(\text{mary})$

Backward chaining example

- KB:

1. $\text{allergies}(X) \rightarrow \text{sneeze}(X)$
2. $\text{cat}(Y) \wedge \text{allergicToCats}(X) \rightarrow \text{allergies}(X)$
3. $\text{cat}(\text{felix})$
4. $\text{allergicToCats}(\text{mary})$

- Goal:

- $\text{sneeze}(\text{mary})$ if $\text{allergies}(\text{mary})$
- $\text{allergies}(\text{mary})$ if $\text{cat}(Y) \wedge \text{allergicToCats}(\text{mary})$
- $\text{allergies}(\text{mary})$ if $\text{cat}(Y) \wedge \text{cat}(\text{felix})$
- $\text{allergies}(\text{mary})$
- $\text{sneeze}(\text{mary})$

In practice...

- Backward chaining is very common and used in Prolog and many other rule-based systems
- It does not have a problem if the facts of rules change
- It can be less efficient if you want to prove the same conclusion multiple times

Forward vs. backward chaining

- Forward chaining is data-driven
 - Automatic, unconscious processing, e.g., object recognition, routine decisions
 - May do lots of work that is irrelevant to the goal
 - Efficient when you want to compute **all conclusions**
- Backward chaining is goal-driven, better for problem-solving and query answering
 - Where are my keys? How do I get to my next class?
 - Complexity can be much less than linear wrt KB size
 - Efficient when you want **one or a few conclusions**
 - Good where the underlying facts are changing

Mixed strategy

- Many practical reasoning systems do both forward and backward chaining
- How you encode rule determines how it's used:
 - spouse(X,Y) => spouse(Y,X) % forward chaining
 - father(X,Y) <= parent(X,Y), male(X) % backward chaining
- Forward chaining rules useful if you want to draw conclusions and **materialize them as facts**
 - This also easily avoid loops, e.g., don't trigger forward chaining if the fact already exists
- Given a model of your rules and the kind of reasoning needed, you can decide which to encode as FC and which as BC rules

Completeness of GMP

- GMP (using forward or backward chaining) is complete for KBs containing only Horn clauses
- not complete for simple KBs with non-Horn clauses
- What is entailed by the following sentences?
 1. $(\forall x) P(x) \rightarrow Q(x)$
 2. $(\forall x) \neg P(x) \rightarrow R(x)$
 3. $(\forall x) Q(x) \rightarrow S(x)$
 4. $(\forall x) R(x) \rightarrow S(x)$

Completeness of GMP

- The following entail that $S(A)$ is true:
 1. $(\forall x) P(x) \rightarrow Q(x)$
 2. $(\forall x) \neg P(x) \rightarrow R(x)$
 3. $(\forall x) Q(x) \rightarrow S(x)$
 4. $(\forall x) R(x) \rightarrow S(x)$
- If we want to conclude $S(A)$, with GMP we can't, since the second one isn't a Horn clause
- It is equivalent to $P(x) \vee R(x)$

How about Prolog?

- Prolog syntax is a bit different, putting the rule's conclusion first

hasMother(?x, ?m) :- hasParent(?x, ?m), female(?m) .

head = conclusion

body = conjunction of conditions

- A fact is a rule w/o a body (i.e., no conditions)

hasParent(john, tom).

hasParent(john, mary).

female(mary).

- Prolog 'proves' queries by matching a fact, or a rule's conclusion and then proving each condition in the rule's body

Can we do this in Prolog?

Try encoding this in Prolog

1. $q(X) :- p(X).$

2. $r(X) :- \text{neg}(p(X)).$

3. $s(X) :- q(X).$

4. $s(X) :- r(X).$

1. $(\forall x) P(x) \rightarrow Q(x)$

2. $(\forall x) \neg P(x) \rightarrow R(x)$

3. $(\forall x) Q(x) \rightarrow S(x)$

4. $(\forall x) R(x) \rightarrow S(x)$

- We should not use `\+` or **not** (in SWI) for negation since it means “**negation as failure**”
- i.e., selective use of the closed world assumption or “if I can’t prove it, it must be false”
- Prolog explores possible proofs independently
- It can’t take a larger view and realize that one branch must be true since $p(x) \vee \sim p(x)$ is always true

Fín