15-780: Graduate AI Lecture 1. Intro & Logic

Geoff Gordon (this lecture)
Tuomas Sandholm
TAs Byron Boots, Sam Ganzfried



15-780 Graduate Artificial Intelligence Spring 2009

Geoff Gordon and Tuomas Sandholm School of Computer Science, Carnegie Mellon University



About | People | Lectures | Recitations | Homework | Projects

Mailing lists

Textbook

Grading

Homework policy

Collaboration policy

Late policy

Regrade policy

Final project

Note to people outside CMU Class Lectures: Tuesdays and Thursdays 10:30-11:50am in 4623 Wean Hall

Recitations: TBA

This course is targeted at graduate students who need to learn about and perform current-day research in artificial intelligence---the discipline of designing intelligent decision-making machines. Techniques from probability, statistics, game theory, algorithms, operations research and optimal control are increasingly important tools for improving the intelligence and autonomy of machines, whether those machines are robots surveying Antarctica, schedulers moving billions of dollars of inventory, spacecraft deciding which experiments to perform, or vehicles negotiating for lanes on the freeway. This Al course is a review of a selected set of these tools. The course will cover the ideas underlying these tools, their implementation, and how to use them or extend them in your research. Students entering the class should have a pre-existing working knowledge of

http://www.cs.cmu.edu/~ggordon/780/

http://www.cs.cmu.edu/~sandholm/cs15-780S09/

Website highlights

- Book: Russell and Norvig. Artificial
 Intelligence: A Modern Approach, 2nd ed.
- o Grading: 4–5 HWs, "mid" term, project
- Project: proposal, 2 interim reports, final report, poster
- Office hours

Website highlights

- Authoritative source for readings, HWs
- Please check the website regularly for readings (for Lec. 1–3, Russell & Norvig Chapters 7–9)

Background

- No prerequisites
- But, suggest familiarity with at least some of the following:
 - Linear algebra
 - Calculus
 - Algorithms & data structures
 - Complexity theory

Waitlist, Audits

 If you need us to approve something, send us email

Course email list

- o 15780students@...
- o domain mailman.srv.cs.cmu.edu
- To subscribe/unsubscribe:
 - o email 15780students-request@...
 - word "help" in subject or body

Matlab

- Should all have access to Matlab via school computers
 - Those with access to CS license servers, please use if possible
 - Limited number of Andrew licenses
- Tutorial TBA soon
- HWs: please use C, C++, Java, or Matlab



• Easy part: A

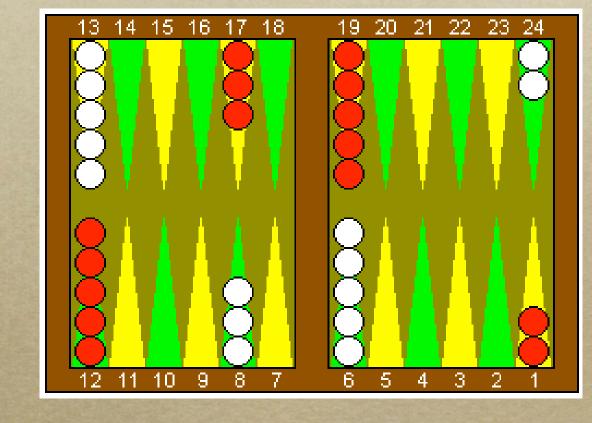
- Easy part: A
- Hard part: I

- Easy part: A
- Hard part: I
 - Anything we don't know how to make a computer do yet

- Easy part: A
- Hard part: I
 - Anything we don't know how to make a computer do yet
 - Corollary: once we do it, it isn't AI
 anymore :-)

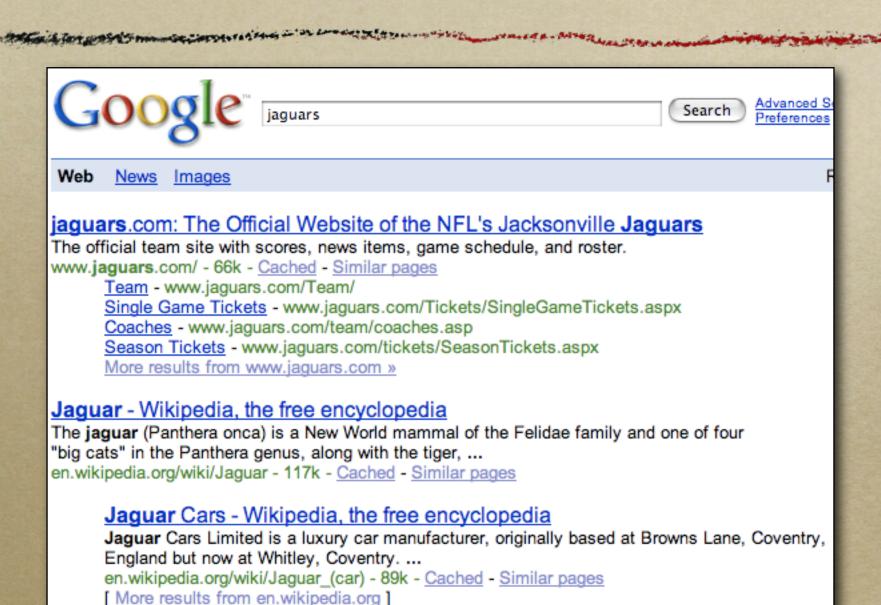
Definition by examples

- Card games
 - o Poker
 - Bridge
- Board games
 - o Deep Blue
 - TD-Gammon

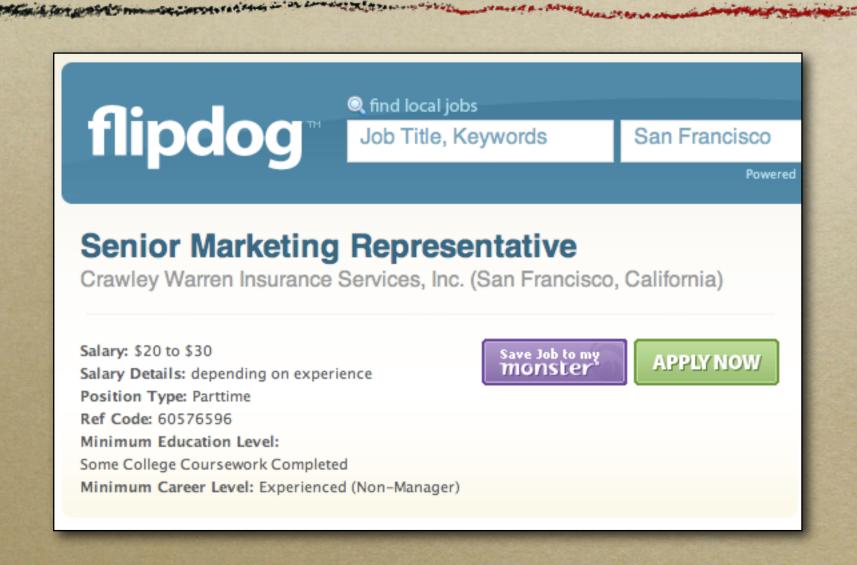


Samuels's checkers player

Web search

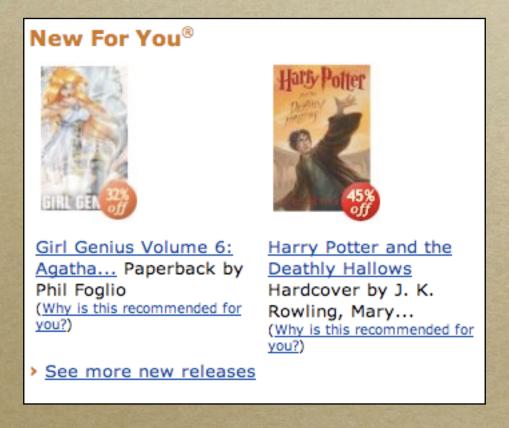


Web search, cont'd

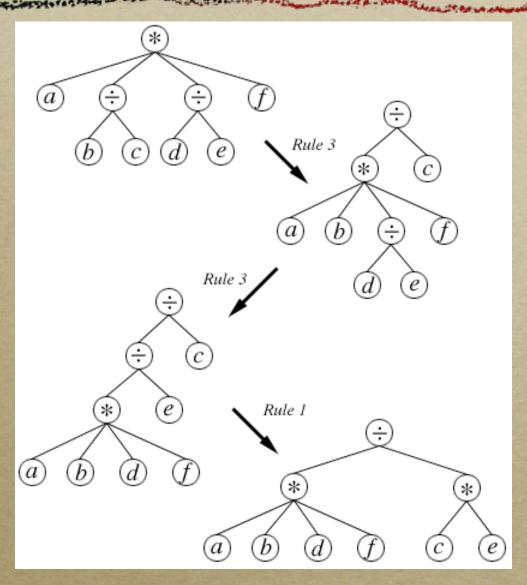


Recommender systems





Computer algebra systems



from http://www.math.wpi.edu/IQP/BVCalcHist/calctoc.html

Grand Challenge road race



Getting from A to B

Round Trip	One Way Multi-Segment								
nound mp	<u>Mail Segment</u>								
	an any simant within (0, 1)								
from	or any airport within 0 miles 💠								
to	or any airport within 0 miles 💠								
outbound date	Oct † 1 † on this day only † departing † anytime †								
return date	Oct \$ 8 \$ on this day only \$ departing \$ anytime \$								
travelers	adults seniors youths children infants in seat infants on lap (18 to 61) (62 plus) (12 to 17) (2 to 11) (under 2) (under 2)								
stops	○ nonstops only ○ up to 1 stop ○ up to 2 stops • no limit								
sales city	sales city BOS (change only for trips originating outside the United States: learn more)								
more options (cabin, airport changes, seat availability, etc)									
	Go!								

ITA software (<u>http://beta.itasoftware.com</u>)

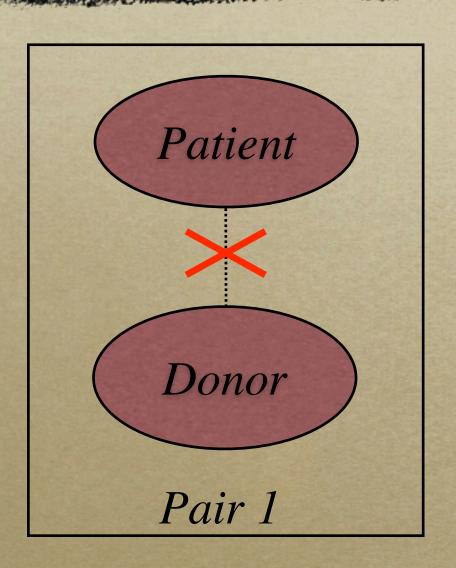
Robocup

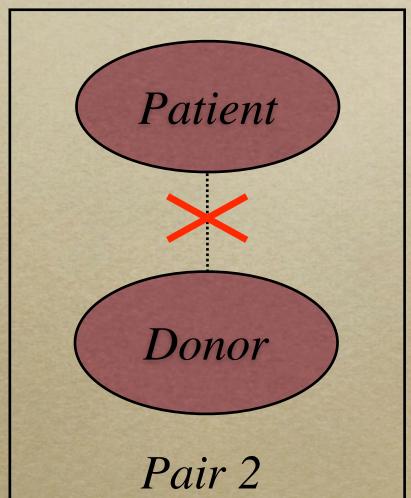


Kidney exchange

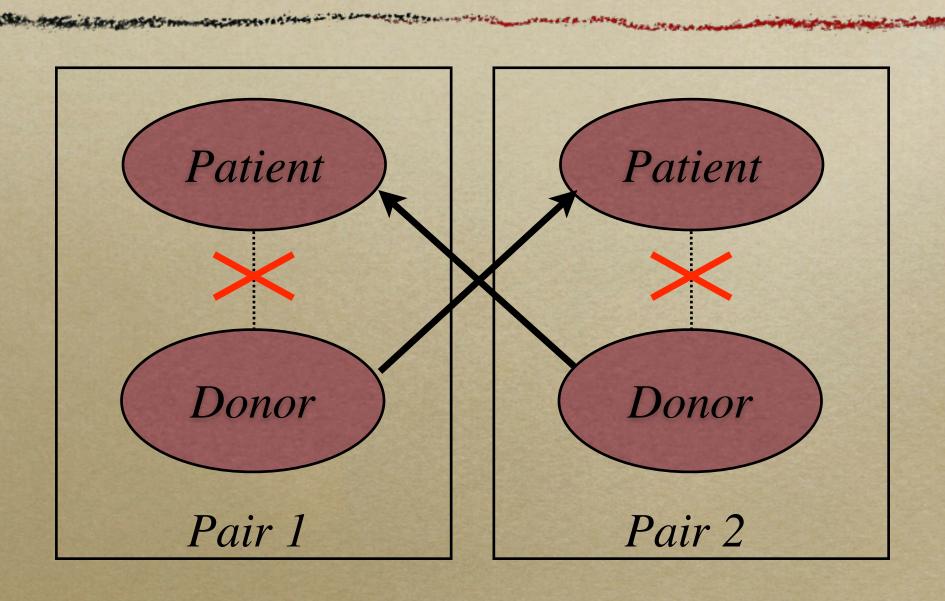
- In US, ≥ 50,000/yr get lethal kidney disease
- Cure = transplant, but donor must be compatible (blood type, tissue type, etc.)
- Wait list for cadaver kidneys: 2–5 years
- Live donors: have 2 kidneys, can survive w/ 1
- Illegal to buy/sell, but altruists/friends/family donate

Kidney Exchange

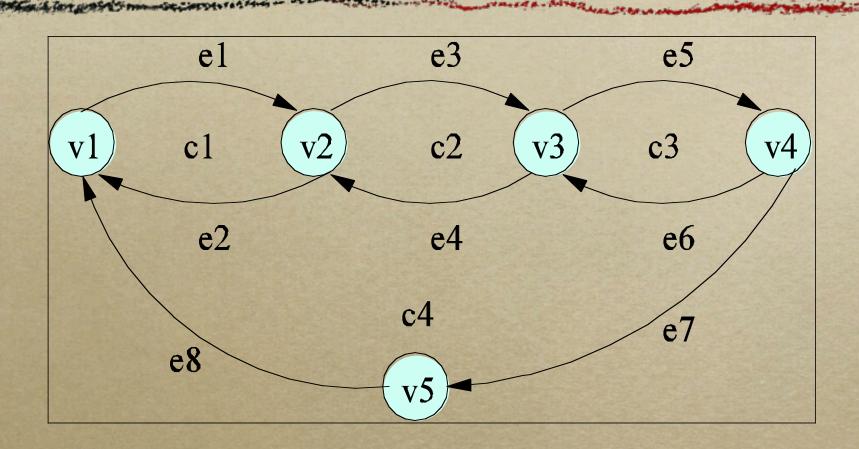




Kidney Exchange

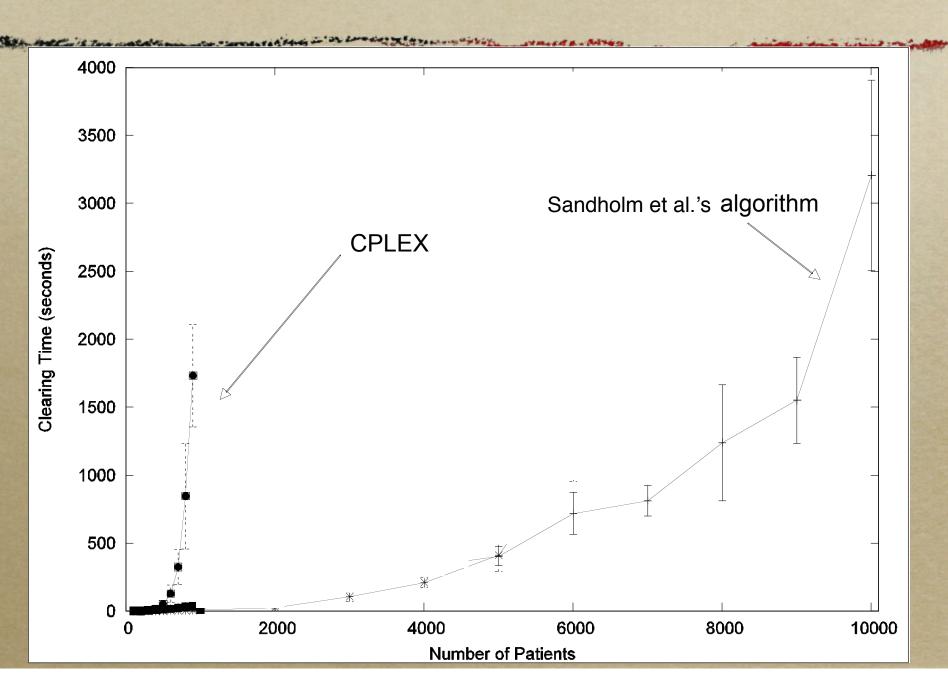


Optimization: cycle cover



Cycle length constraint => extremely hard (NP-complete)
combinatorial optimization problem
National market predicted to have 10,000 patients at any one time

Optimization performance



More examples

- Motor skills: riding a bicycle, learning to walk, playing pool, ...
- Vision

More examples

- Valerie and Tank, the Roboceptionists
- Social skills: attending a party, giving directions, ...



More examples

- Natural language
- Speech recognition

Common threads

- Finding the needle in the haystack
 - Search
 - Optimization
 - Summation / integration
- Set the problem up well (so that we can apply a standard algorithm)

Common threads

- Managing uncertainty
 - o chance outcomes (e.g., dice)
 - sensor uncertainty ("hidden state")
 - o opponents
- The more different types of uncertainty, the harder the problem (and the slower the solution)

Classic AI

- No uncertainty, pure search
 - Mathematica
 - o deterministic planning
 - Sudoku
- This is the topic of Part I of the course

http://www.cs.ualberta.ca/~aixplore/search/IDA/Applet/SearchApplet.html

Classic AI

- o No uncertainty, pure search
 - Mathematica
 - o deterministic planning
 - Sudoku

			6	3			4	7	
			5	8		7			
	1							2	3
		6		1	9				
	4	9							
							1	9	8
	6					3	5		
			8		5				2
		7	4			6		8	
-	and the last	minutes of		Name and		2000000	-		

• This is the topic of Part I of the course

http://www.cs.ualberta.ca/~aixplore/search/IDA/Applet/SearchApplet.html

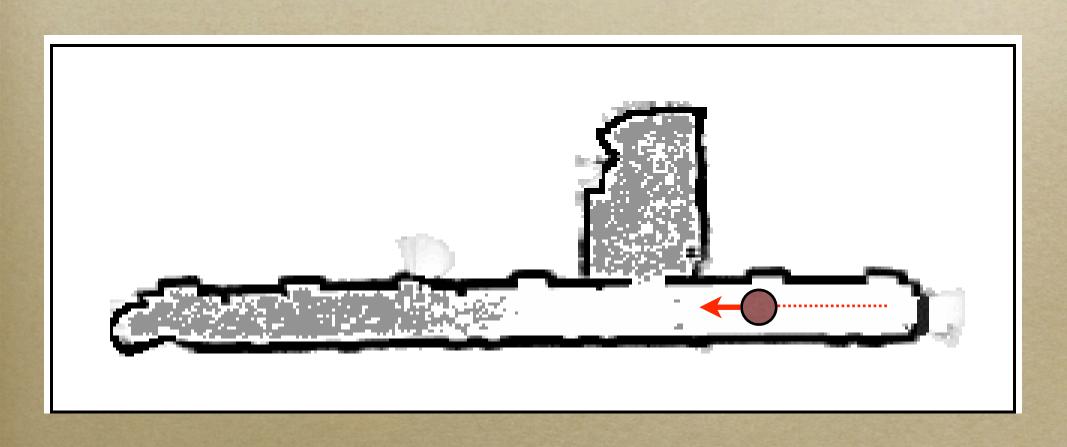
Outcome uncertainty

- In backgammon, don't know ahead of time what the dice will show
- When driving down a corridor, wheel slippage causes unexpected deviations
- o Open a door, find out what's behind it
- MDPs (later)

Sensor uncertainty

- For given set of immediate measurements, multiple world states may be possible
- Image of a handwritten digit $\rightarrow 0, 1, ..., 9$
- Image of room → person locations, identities
- Laser rangefinder scan of a corridor → map, robot location

Sensor uncertainty example

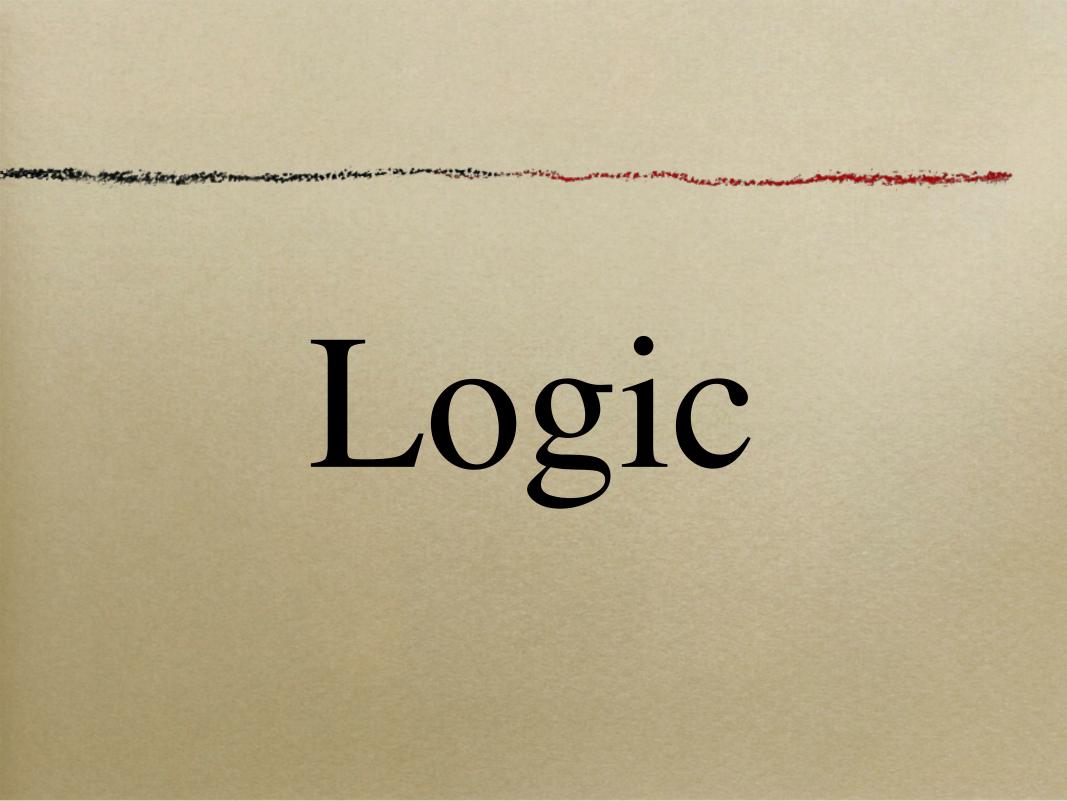


Opponents cause uncertainty

- In chess, must guess what opponent will do; cannot directly control him/her
- Alternating moves (game trees)
 - not really uncertain (Part I)
- Simultaneous or hidden moves: game theory (later)

Other agents cause uncertainty

- In many AI problems, there are other agents who aren't (necessarily) opponents
 - Ignore them & pretend part of Nature
 - Assume they're opponents (pessimistic)
 - Learn to cope with what they do
 - Try to convince them to cooperate (paradoxically, this is the hardest case)
- More later



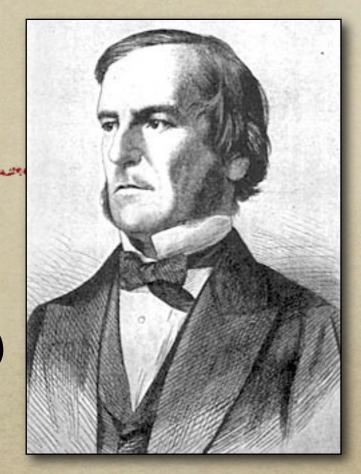
Why logic?

- Search: for problems like Sudoku, can write compact description of rules
- Reasoning: figure out consequences of the knowledge we've given our agent
- ...and, logical inference is a special case of probabilistic inference

Propositional logic

- Constants: Tor F
- Variables: x, y (values T or F)
- ∘ Connectives: ∧, ∨, ¬
 - Can get by w/ just NAND
 - Sometimes also add others:

$$\oplus$$
, \Rightarrow , \Leftrightarrow , ...



George Boole 1815–1864

Propositional logic

- Build up expressions like $\neg x \Rightarrow y$
- \circ Precedence: \neg , \land , \lor , \Rightarrow
- Terminology: variable or constant with or w/o negation = literal
- Whole thing = formula or sentence

Expressive variable names

- Rather than variable names like x, y, may use names like "rains" or "happy(John)"
- For now, "happy(John)" is just a string with no internal structure
 - o there is no "John"
 - $happy(John) \Rightarrow \neg happy(Jack)$ means the same as $x \Rightarrow \neg y$

But what does it mean?

- A formula defines a mapping (assignment to variables) $\mapsto \{T, F\}$
- Assignment to variables = model
- \circ For example, formula $\neg x$ yields mapping:

X	$\neg x$
$\mid T \mid$	F
$oxed{F}$	T

Questions about models and sentences

- How many models make a sentence true?
 - A sentence is satisfiable if it is True in some model
 - If not satisfiable, it is a contradiction (False in every model)
 - A sentence is valid if it is True in every model (called a tautology)

Questions about models and sentences

- How is the variable X set in {some, all} true models?
- This is the most frequent question an agent would ask: given my assumptions, can I conclude X? Can I rule X out?

More truth tables

X	y	$x \wedge y$
T	$\mid T \mid$	T
T	F	F
F	T	$oxed{F}$
F	F	F

\mathcal{X}	y	$x \vee y$
T	$\mid T \mid$	T
T	F	T
$oxed{F}$	T	T
$oxed{F}$	F	F

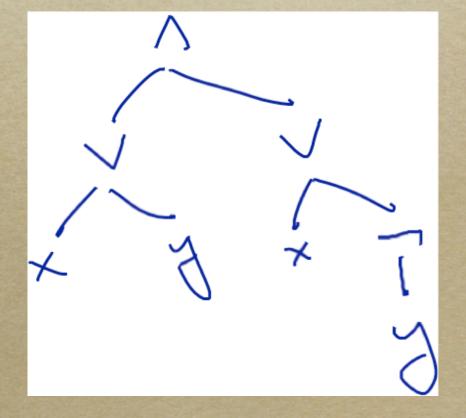
Truth table for implication

- $(a \Rightarrow b)$ is logically equivalent to $(\neg a \lor b)$
- o If a is True, b must be True too
- If a False, no requirement on b
- E.g., "if I go to the movie I will have popcorn": if no movie, may or may not have popcorn

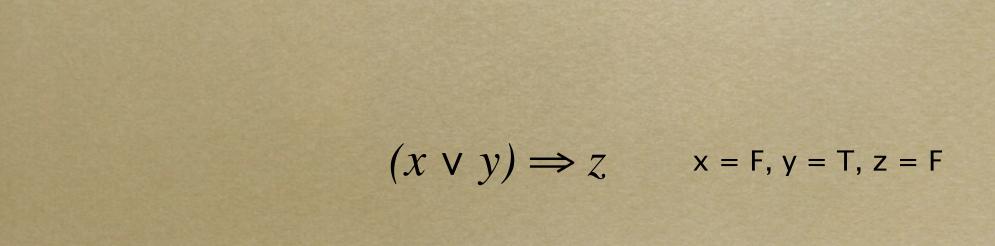
a	b	$a \Rightarrow b$
T	T	T
T	F	F
$oxed{F}$	T	T
$oxed{F}$	F	T

Complex formulas

- o To evaluate a bigger formula
 - \circ $(x \lor y) \land (x \lor \neg y)$ when x = F, y = F
- o Build a parse tree
- Fill in variables at leaves using model
- Work upwards using truth tables for connectives

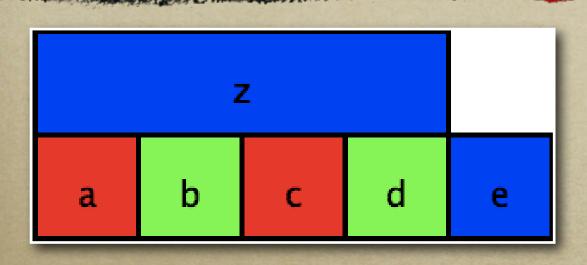


Another example





3-coloring

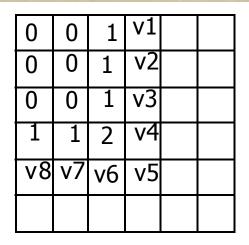


Sudoku

SuDoku Puzzle								
		6	3			4	7	
		5	8		7			
1							2	3
	6		1	9				
4	9							
						1	9	8
6					3	5		
		8		5				2
	7	4			6		8	

http://www.cs.qub.ac.uk/~I.Spence/SuDoku/SuDoku.html

Minesweeper



 $V = \{ v1, v2, v3, v4, v5, v6, v7, v8 \}, D = \{ B (bomb), S (space) \}$ $C = \{ (v1,v2) : \{ (B,S), (S,B) \}, (v1,v2,v3) : \{ (B,S,S), (S,B,S), (S,S,B) \}, ... \}$ $v1 \subseteq \{ (v1,v2) : \{ (B,S,S), (S,B,S), (S,S,B) \}, ... \}$

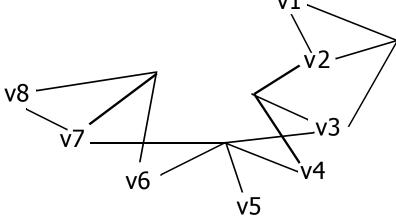


image courtesy Andrew Moore

Propositional planning

```
init: have(cake)
goal: have(cake), eaten(cake)
eat(cake):
 pre: have(cake)
 eff: -have(cake), eaten(cake)
bake(cake):
 pre: -have(cake)
 eff: have(cake)
```

Other important logic problems

- Scheduling (e.g., of factory production)
- Facility location
- Circuit layout
- Multi-robot planning

Important issue: handling uncertainty

Working with formulas

Truth tables get big fast

x	y	Z	$(x \lor y) \Rightarrow z$
T	T	T	
T	T	F	
T	F	T	
T	F	F	
$oxed{F}$	T	T	
$oxed{F}$	T	F	
$oxed{F}$	F	T	
$oxed{F}$	F	F	

Truth tables get big fast

x	y	z	a	$(x \vee y \vee a) \Rightarrow z$
T	T	T	T	
T	T	F	T	
T	F	T	T	
T	F	F	T	
F	T	T	T	
F	T	F	T	
F	F	T	T	
F	F	F	T	
T	T	T	F	
T	T	F	F	
T	F	T	F	
T	F	F	F	
F	T	T	F	
F	T	F	F	
F	F	T	F	
F	F	F	F	

Definitions

- Two sentences are equivalent, $A \equiv B$, if they have same truth value in every model
 - \circ (rains \Rightarrow pours) \equiv (\neg rains \lor pours)
 - o reflexive, transitive, commutative
- Simplifying = transforming a formula into a shorter*, equivalent formula

Transformation rules

```
(\alpha \wedge \beta) \equiv (\beta \wedge \alpha) \quad \text{commutativity of } \wedge \\ (\alpha \vee \beta) \equiv (\beta \vee \alpha) \quad \text{commutativity of } \vee \\ ((\alpha \wedge \beta) \wedge \gamma) \equiv (\alpha \wedge (\beta \wedge \gamma)) \quad \text{associativity of } \wedge \\ ((\alpha \vee \beta) \vee \gamma) \equiv (\alpha \vee (\beta \vee \gamma)) \quad \text{associativity of } \vee \\ \neg(\neg \alpha) \equiv \alpha \quad \text{double-negation elimination}
```

$$(\alpha \land (\beta \lor \gamma)) \equiv ((\alpha \land \beta) \lor (\alpha \land \gamma))$$
 distributivity of \land over \lor $(\alpha \lor (\beta \land \gamma)) \equiv ((\alpha \lor \beta) \land (\alpha \lor \gamma))$ distributivity of \lor over \land

 α, β, γ are arbitrary formulas

More rules

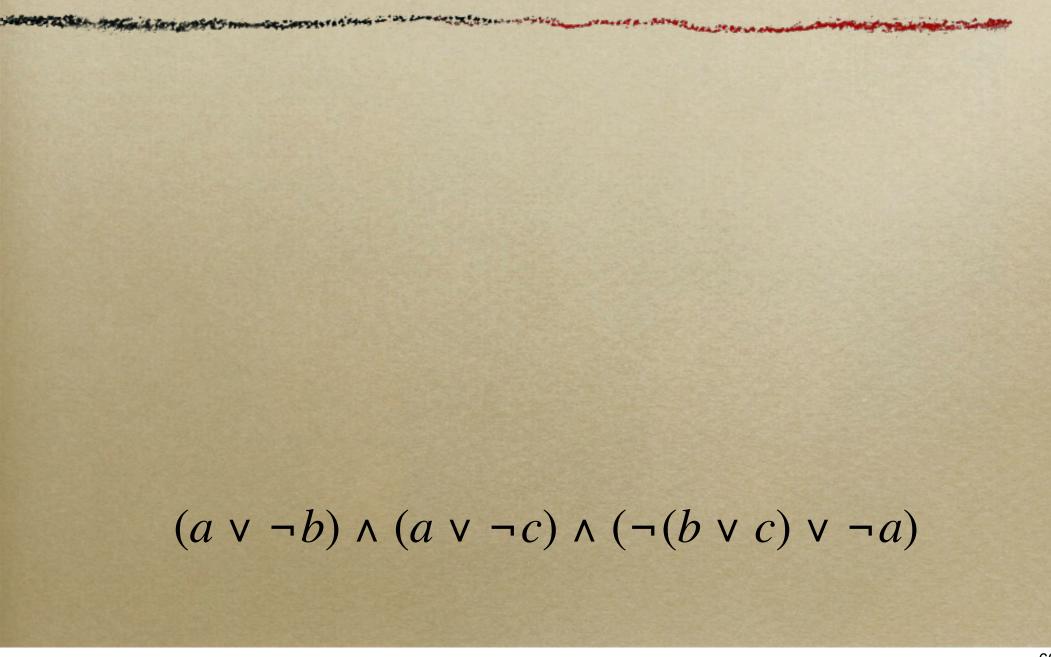
$$(\alpha \Rightarrow \beta) \equiv (\neg \beta \Rightarrow \neg \alpha)$$
 contraposition $(\alpha \Rightarrow \beta) \equiv (\neg \alpha \lor \beta)$ implication elimination $(\alpha \Leftrightarrow \beta) \equiv ((\alpha \Rightarrow \beta) \land (\beta \Rightarrow \alpha))$ biconditional elimination $\neg(\alpha \land \beta) \equiv (\neg \alpha \lor \neg \beta)$ de Morgan $\neg(\alpha \lor \beta) \equiv (\neg \alpha \land \neg \beta)$ de Morgan

 α , β are arbitrary formulas

Still more rules...

- o ... can be derived from truth tables
- For example:
 - \circ $(a \lor \neg a) \equiv True$
 - \circ (True \vee a) \equiv True
 - \circ (False \land a) \equiv False

Example



Normal Forms

Normal forms

- A normal form is a standard way of writing a formula
- E.g., conjunctive normal form (CNF)
 - conjunction of disjunctions of literals
 - $\circ (x \vee y \vee \neg z) \wedge (x \vee \neg y) \wedge (z)$
 - Each disjunct called a clause
- Any formula can be transformed into CNF w/o changing meaning

CNF cont'd

```
happy(John) ∧
(¬happy(Bill) ∨ happy(Sue)) ∧
man(Socrates) ∧
(¬man(Socrates) ∨ mortal(Socrates))
```

- Often used for storage of knowledge database
 - called knowledge base or KB
- Can add new clauses as we find them out
- Each clause in KB is separately true (if KB is)

Another normal form: DNF

- DNF = disjunctive normal form = disjunction of conjunctions of literals
- Doesn't compose the way CNF does: can't just add new conjuncts w/o changing meaning of KB
- Example:

```
(rains ∨ ¬pours) ∧ fishing

≡

(rains ∧ fishing) ∨ (¬pours ∧ fishing)
```

Transforming to CNF or DNF

- Naive algorithm:
 - ∘ replace all connectives with ∧∨¬
 - move negations inward using De Morgan's laws and double-negation
 - repeatedly distribute over ∧ over ∨ for DNF (∨ over ∧ for CNF)

Example

Put the following formula in CNF

 $(a \lor b \lor \neg c) \land \neg (d \lor (e \land f)) \land (c \lor d \lor e)$

 $(a \lor b \lor \neg c) \land \neg (d \lor (e \land f)) \land (c \lor d \lor e)$

Example

Now try DNF

 $(a \lor b \lor \neg c) \land \neg (d \lor (e \land f)) \land (c \lor d \lor e)$

Discussion

- Problem with naive algorithm: it's exponential! (Space, time, size of result.)
- Each use of distributivity can almost double the size of a subformula

A smarter transformation

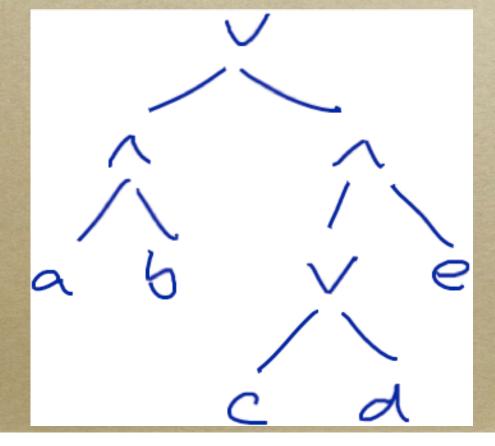
- Can we avoid exponential blowup in CNF?
- Yes, if we're willing to introduce new variables
- G. Tseitin. On the complexity of derivation in propositional calculus. Studies in Constrained Mathematics and Mathematical Logic, 1968.

Tseitin example

• Put the following formula in CNF:

 $(a \wedge b) \vee ((c \vee d) \wedge e)$

o Parse tree:

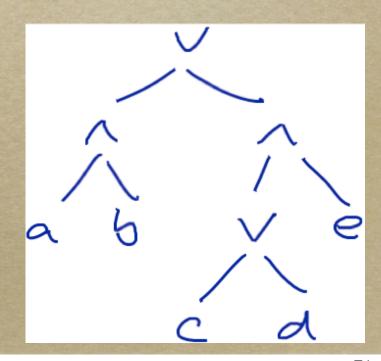


Introduce temporary variables

$$\circ \ x = (a \land b)$$

$$\circ \ y = (c \lor d)$$

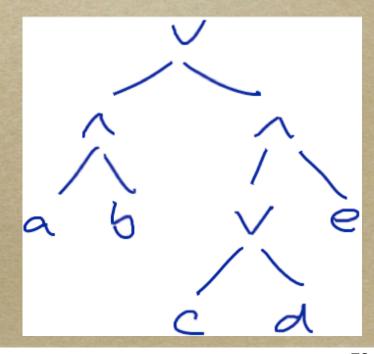
$$\circ \ z = (y \land e)$$



• To ensure $x = (a \land b)$, want

$$\circ x \Rightarrow (a \land b)$$

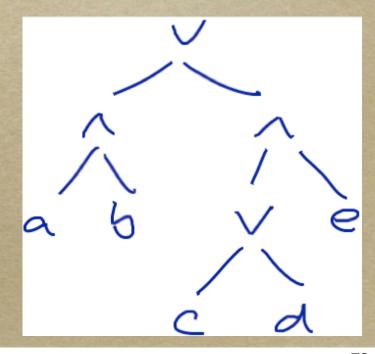
$$\circ (a \land b) \Longrightarrow x$$



$$\circ x \Rightarrow (a \land b)$$

$$\circ (\neg x \lor (a \land b))$$

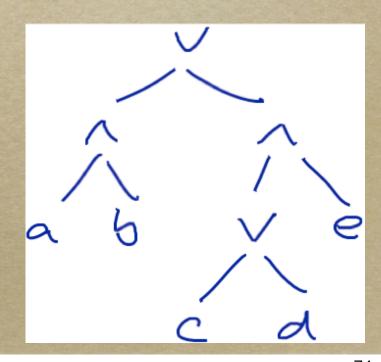
$$\circ (\neg x \lor a) \land (\neg x \lor b)$$



$$\circ (a \land b) \Rightarrow x$$

$$\circ (\neg (a \land b) \lor x)$$

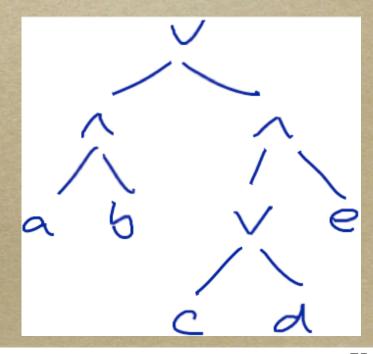
$$\circ (\neg a \lor \neg b \lor x)$$



• To ensure $y = (c \lor d)$, want

$$\circ y \Rightarrow (c \lor d)$$

$$\circ (c \lor d) \Rightarrow y$$



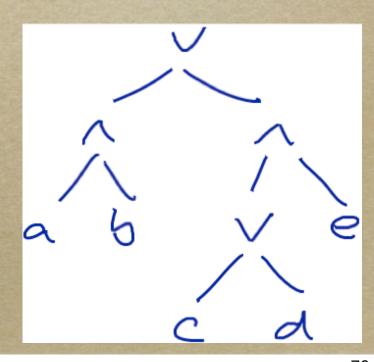
$$\circ y \Rightarrow (c \lor d)$$

$$\circ (\neg y \lor c \lor d)$$

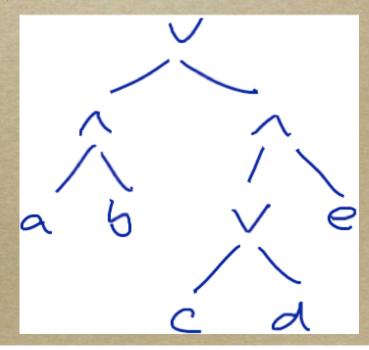
$$\circ (c \lor d) \Rightarrow y$$

$$\circ ((\neg c \land \neg d) \lor y)$$

$$\circ (\neg c \lor y) \land (\neg d \lor y)$$



- \circ Finally, $z = (y \land e)$
- $\circ z \Rightarrow (y \land e) \equiv (\neg z \lor y) \land (\neg z \lor e)$
- $\circ (y \land e) \Rightarrow z \equiv (\neg y \lor \neg e \lor z)$



Tseitin end result

$$(a \wedge b) \vee ((c \vee d) \wedge e) \equiv$$

$$(\neg x \lor a) \land (\neg x \lor b) \land (\neg a \lor \neg b \lor x) \land$$

 $(\neg y \lor c \lor d) \land (\neg c \lor y) \land (\neg d \lor y) \land$
 $(\neg z \lor y) \land (\neg z \lor e) \land (\neg y \lor \neg e \lor z) \land$
 $(x \lor z)$



Entailment

- Sentence A entails sentence B, $A \models B$, if B is True in every model where A is
 - same as saying that $(A \Rightarrow B)$ is valid

Proof tree

- o A tree with a formula at each node
- At each internal node, children ⊨ parent
- Leaves: assumptions or premises
- Root: consequence
- If we believe assumptions, we should also believe consequence

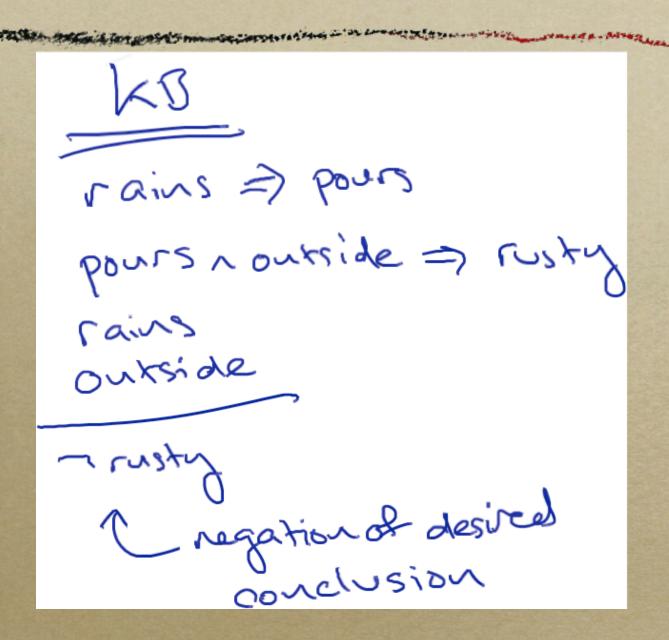
Proof tree example

rains => pours
pours noutside => rusty
rains
outside

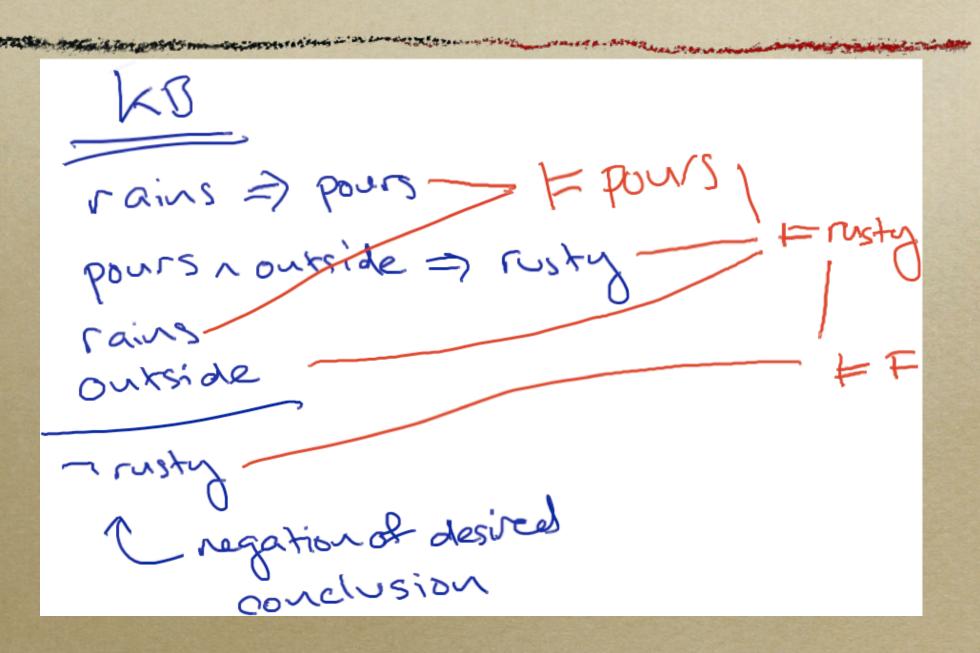
Proof by contradiction

- Assume opposite of what we want to prove, show it leads to a contradiction
- Suppose we want to show $KB \models S$
- Write KB' for $(KB \land \neg S)$
- Build a proof tree with
 - o assumptions drawn from clauses of KB'
 - \circ conclusion = F
 - \circ so, $(KB \land \neg S) \models F$ (contradiction)

Proof by contradiction



Proof by contradiction





Inference rule

- To make a proof tree, we need to be able to figure out new formulas entailed by KB
- Method for finding entailed formulas = inference rule
- We've implicitly been using one already

Modus ponens

$$(a \land b \land c \Rightarrow d) \ a \ b \ c$$

$$d$$

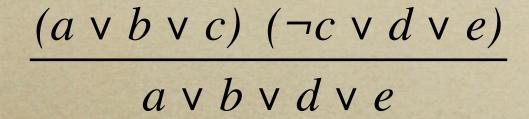
- Probably most famous inference rule: all men are mortal, Socrates is a man, therefore Socrates is mortal
- Quantifier-free version:
 man(Socrates) ∧
 (man(Socrates) ⇒ mortal(Socrates))

Another inference rule

$$\frac{(a \Rightarrow b) \ \neg b}{\neg a}$$

- Modus tollens
- If it's raining the grass is wet; the grass is not wet, so it's not raining

One more...



- o Resolution
- Combines two sentences that contain a literal and its negation
- Not as commonly known as modus ponens / tollens

Resolution example

- Modus ponens / tollens are special cases
- Modus tollens:

```
(\neg raining \lor grass-wet) \land \neg grass-wet \models \neg raining
```

$$\frac{(a \lor b \lor c) (\neg c \lor d \lor e)}{a \lor b \lor d \lor e}$$

- Simple proof by case analysis
- Consider separately cases where we assign c = True and c = False

 $(a \lor b \lor c) \land (\neg c \lor d \lor e)$

• Case
$$c = True$$

$$(a \lor b \lor T) \land (F \lor d \lor e)$$

$$= (T) \land (d \lor e)$$

$$= (d \lor e)$$

 $(a \lor b \lor c) \land (\neg c \lor d \lor e)$

• Case
$$c = False$$

$$(a \lor b \lor F) \land (T \lor d \lor e)$$

$$= (a \lor b) \land (T)$$

$$= (a \lor b)$$

 $(a \lor b \lor c) \land (\neg c \lor d \lor e)$

Since c must be True or False, conclude
 (d v e) v (a v b)

as desired

Soundness and completeness

- An inference procedure is sound if it can only conclude things entailed by KB
 - common sense; haven't discussed anything unsound
- A procedure is **complete** if it can conclude everything entailed by KB

Completeness

- o Modus ponens by itself is incomplete
- Resolution is complete for propositional logic
 - not obvious—famous theorem

Variations

- Horn clause inference (faster)
- Ways of handling uncertainty (slower)
- CSPs (sometimes more convenient)
- o Quantifiers / first-order logic

Horn clauses

- Horn clause: $(a \land b \land c \Rightarrow d)$
- \circ Equivalently, $(\neg a \lor \neg b \lor \neg c \lor d)$
- Disjunction of literals, at most one of which is positive
- Positive literal = head, rest = body

Use of Horn clauses

People find it easy to write Horn clauses
 (listing out conditions under which we can conclude head)

 $happy(John) \land happy(Mary) \Rightarrow happy(Sue)$

No negative literals in above formula;
 again, easier to think about

Why are Horn clauses important

- Modus ponens alone is complete
- So is modus tollens alone
- Inference in a KB of propositional Horn clauses is linear

Forward chaining

- Look for a clause with all body literals satisfied
- Add its head to KB (modus ponens)
- Repeat
- See RN for more details

Handling uncertainty

- Fuzzy logic / certainty factors
 - o simple, but don't scale
- Nonmonotonic logic
 - o also doesn't scale
- Probabilities
 - may or may not scale—more later
 - Dempster-Shafer (interval probability)

Certainty factors

- KB assigns a certainty factor in [0, 1] to each proposition
- o Interpret as "degree of belief"
- When applying an inference rule, certainty factor for consequent is a function of certainty factors for antecedents (e.g., minimum)

Problems w/ certainty factors

- Hard to separate a large KB into mostlyindependent chunks that interact only through a well-defined interface
- Certainty factors are not probabilities
 (i.e., do not obey Bayes' Rule)

- Suppose we believe all birds can fly
- Might add a set of sentences to KB

$$bird(Polly) \Rightarrow flies(Polly)$$

 $bird(Tweety) \Rightarrow flies(Tweety)$

 $bird(Tux) \Rightarrow flies(Tux)$

 $bird(John) \Rightarrow flies(John)$

. . .

- Fails if there are penguins in the KB
- Fix: instead, add

 $bird(Polly) \land \neg ab(Polly) \Rightarrow flies(Polly)$ $bird(Tux) \land \neg ab(Tux) \Rightarrow flies(Tux)$

...

- ab(Tux) is an "abnormality predicate"
- Need separate $ab_i(x)$ for each type of rule

- Now set as few abnormality predicates as possible
- Can prove flies(Polly) or flies(Tux) with no ab(x) assumptions
- If we assert ¬flies(Tux), must now assume ab(Tux) to maintain consistency
- Can't prove flies(Tux) any more, but can still prove flies(Polly)

- Works well as long as we don't have to choose between big sets of abnormalities
 - is it better to have 3 flightless birds or 5 professors that don't wear jackets with elbow-patches?
 - even worse with nested abnormalities:
 birds fly, but penguins don't, but
 superhero penguins do, but ...