

Warm-up:

What is the relationship between number of constraints and number of possible solutions?

In other words, as the number of the constraints increases, does the number of possible solutions:

- A) Increase
- B) Decrease
- C) Stay the same

Announcements

Assignments:

- P2: Optimization
 - Due Thu 2/21, 10 pm

Midterm 1 Exam

- Mon 2/18, in class
- Recitation Fri is a review session
- See Piazza post for details

Alita Class Field Trip!

- Moved to Saturday, 2/23, afternoon

White card feedback

Warm-up:

What is the relationship between number of constraints and number of possible solutions?

In other words, as the number of the constraints increases, does the number of possible solutions:

- A) Increase
- ☒ B) Decrease
- C) Stay the same

Where is the knowledge in our CSPs?



AI: Representation and Problem Solving

Propositional Logic



Instructors: Pat Virtue & Stephanie Rosenthal

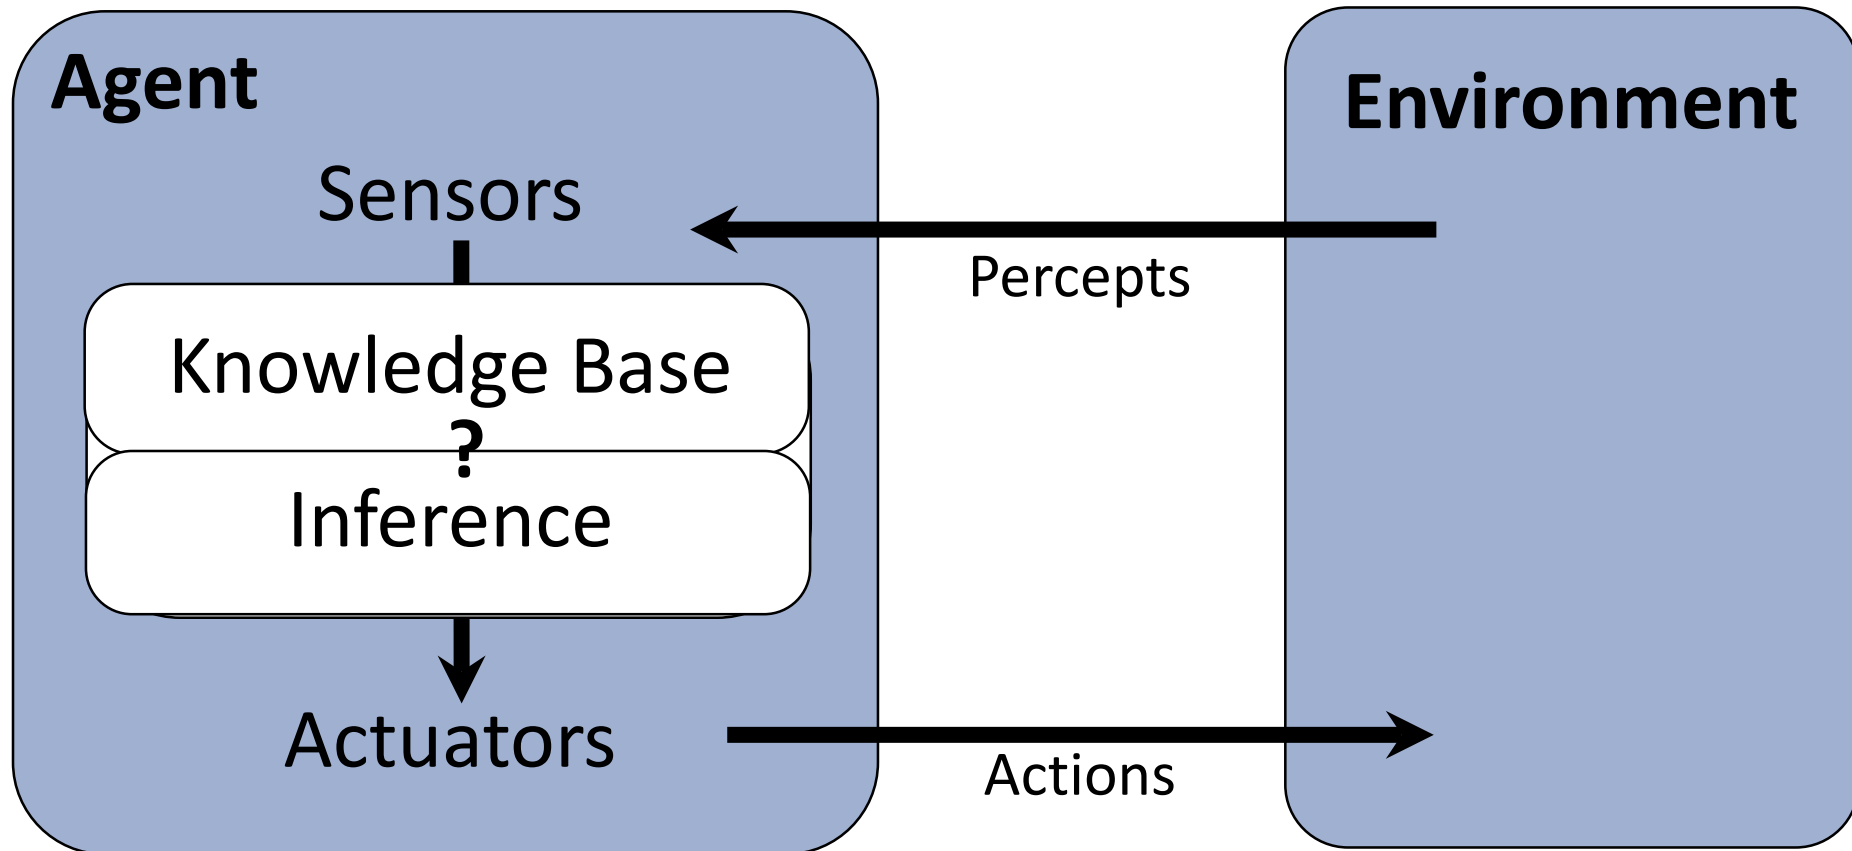
Slide credits: CMU AI, <http://ai.berkeley.edu>

Logic Representation and Problem Solving

To honk or not to honk

Logical Agents


Logical agents and environments



Wumpus World

Logical Reasoning as a CSP

- B_{ij} = breeze felt
- S_{ij} = stench smelt
- P_{ij} = pit here
- W_{ij} = wumpus here
- G = gold



A Knowledge-based Agent

function **KB-AGENT**(percept) returns an action

persistent: KB, a knowledge base

t, an integer, initially 0

TELL(KB, PROCESS-PERCEPT(percept, t))

action \leftarrow **ASK**(KB, **PROCESS-QUERY**(t))

TELL(KB, PROCESS-RESULT(action, t))

t \leftarrow t+1

return action

Logical Agents

So what do we TELL our knowledge base (KB)?

- Facts (sentences)
 - The grass is green
 - The sky is blue
- Rules (sentences)
 - Eating too much candy makes you sick
 - When you're sick you don't go to school
- Percepts and Actions (sentences)
 - Pat ate too much candy today

What happens when we ASK the agent?

- Inference – new sentences created from old
 - Pat is not going to school today

Logical Agents

Sherlock Agent

- Really good knowledge base
 - Evidence
 - Understanding of how the world works (physics, chemistry, sociology)
- Really good inference
 - Skills of deduction
 - “It’s elementary my dear Watson”




Dr. Strange?
Alan Turing?
Kahn?

Worlds

What are we trying to figure out?



- Who, what, when, where, why
- Time: past, present, future



- Actions, strategy
- Partially observable? Ghosts, Walls

Which world are we living in?

Models











How do we represent possible worlds with models and knowledge bases?
How do we then do inference with these representations?

Wumpus World

Possible Models


- $P_{1,2} P_{2,2} P_{3,1}$



Wumpus World

Possible Models


- $P_{1,2} P_{2,2} P_{3,1}$
- Knowledge base
 - Nothing in [1,1]
 - Breeze in [2,1]



Wumpus World

Possible Models


- $P_{1,2} P_{2,2} P_{3,1}$
- Knowledge base
 - Nothing in $[1,1]$
 - Breeze in $[2,1]$
- Query α_1 :
 - No pit in $[1,2]$



Wumpus World

Possible Models

- $P_{1,2} P_{2,2} P_{3,1}$
- Knowledge base
 - Nothing in $[1,1]$
 - Breeze in $[2,1]$
- Query α_2 :
 - No pit in $[2,2]$



Logic Language

Natural language?

Propositional logic

- Syntax: $P \vee (\neg Q \wedge R)$; $X_1 \Leftrightarrow (\text{Raining} \Rightarrow \text{Sunny})$
- Possible world: $\{P=\text{true}, Q=\text{true}, R=\text{false}, S=\text{true}\}$ or 1101
- Semantics: $\alpha \wedge \beta$ is true in a world iff α is true and β is true (etc.)

First-order logic

- Syntax: $\forall x \exists y P(x,y) \wedge \neg Q(\text{Joe}, f(x)) \Rightarrow f(x)=f(y)$
- Possible world: Objects o_1, o_2, o_3 ; P holds for $\langle o_1, o_2 \rangle$; Q holds for $\langle o_3 \rangle$; $f(o_1)=o_1$; $\text{Joe}=o_3$; etc.
- Semantics: $\phi(\sigma)$ is true in a world if $\sigma=o_j$ and ϕ holds for o_j ; etc.

Propositional Logic

Propositional Logic

Symbol:

- Variable that can be true or false
- We'll try to use capital letters, e.g. A , B , $P_{1,2}$
- Often include True and False

Operators:

- $\neg A$: not A
- $A \wedge B$: A and B (conjunction)
- $A \vee B$: A or B (disjunction) Note: this is not an “exclusive or”
- $A \Rightarrow B$: A implies B (implication). If A then B
- $A \Leftrightarrow B$: A if and only if B (biconditional)

Sentences

Propositional Logic Syntax

Given: a set of proposition symbols $\{X_1, X_2, \dots, X_n\}$

- (we often add **True** and **False** for convenience)

X_i is a sentence

If α is a sentence then $\neg\alpha$ is a sentence

If α and β are sentences then $\alpha \wedge \beta$ is a sentence

If α and β are sentences then $\alpha \vee \beta$ is a sentence

If α and β are sentences then $\alpha \Rightarrow \beta$ is a sentence

If α and β are sentences then $\alpha \Leftrightarrow \beta$ is a sentence

And p.s. there are no other sentences!

Notes on Operators

$\alpha \vee \beta$ is inclusive or, not exclusive

Truth Tables

$\alpha \vee \beta$ is inclusive or, not exclusive

α	β	$\alpha \wedge \beta$
F	F	F
F	T	F
T	F	F
T	T	T

α	β	$\alpha \vee \beta$
F	F	F
F	T	T
T	F	T
T	T	T

Notes on Operators

$\alpha \vee \beta$ is inclusive or, not exclusive

$\alpha \Rightarrow \beta$ is equivalent to $\neg\alpha \vee \beta$

- Says who?

Truth Tables

$\alpha \Rightarrow \beta$ is equivalent to $\neg\alpha \vee \beta$

α	β	$\alpha \Rightarrow \beta$	$\neg\alpha$	$\neg\alpha \vee \beta$
F	F	T	T	T
F	T	T	T	T
T	F	F	F	F
T	T	T	F	T

Notes on Operators

$\alpha \vee \beta$ is inclusive or, not exclusive

$\alpha \Rightarrow \beta$ is equivalent to $\neg\alpha \vee \beta$

- Says who?

$\alpha \Leftrightarrow \beta$ is equivalent to $(\alpha \Rightarrow \beta) \wedge (\beta \Rightarrow \alpha)$

- Prove it!

Truth Tables

$\alpha \Leftrightarrow \beta$ is equivalent to $(\alpha \Rightarrow \beta) \wedge (\beta \Rightarrow \alpha)$

α	β	$\alpha \Leftrightarrow \beta$	$\alpha \Rightarrow \beta$	$\beta \Rightarrow \alpha$	$(\alpha \Rightarrow \beta) \wedge (\beta \Rightarrow \alpha)$
F	F	T	T	T	T
F	T	F	T	F	F
T	F	F	F	T	F
T	T	T	T	T	T

Equivalence: it's true in all models. Expressed as a logical sentence:

$$\underline{(\alpha \Leftrightarrow \beta)} \Leftrightarrow \underline{[(\alpha \Rightarrow \beta) \wedge (\beta \Rightarrow \alpha)]}$$

Literals

A literal is an atomic sentence:

- True
- False
- Symbol
- \neg Symbol

Monty Python Inference

There are ways of telling whether she is a witch

Sentences as Constraints

Adding a sentence to our knowledge base constrains the number of possible models:

KB: Nothing

Possible
Models

P	Q	R
false	false	false
false	false	true
false	true	false
false	true	true
true	false	false
true	false	true
true	true	false
true	true	true

Sentences as Constraints

Adding a sentence to our knowledge base constrains the number of possible models:

KB: Nothing

KB: $[(P \wedge \neg Q) \vee (Q \wedge \neg P)] \Rightarrow R$

Possible
Models

P	Q	R
false	false	false
false	false	true
false	true	false
false	true	true
true	false	false
true	false	true
true	true	false
true	true	true

Sentences as Constraints

Adding a sentence to our knowledge base constrains the number of possible models:

KB: Nothing

KB: $[(P \wedge \neg Q) \vee (Q \wedge \neg P)] \Rightarrow R$

KB: R, $[(P \wedge \neg Q) \vee (Q \wedge \neg P)] \Rightarrow R$

Possible
Models

P	Q	R
false	false	false
false	false	true
false	true	false
false	true	true
true	false	false
true	false	true
true	true	false
true	true	true

Sherlock Entailment

“When you have eliminated the impossible, whatever remains, however improbable, must be the truth” – *Sherlock Holmes via Sir Arthur Conan Doyle*

- Knowledge base and inference allow us to remove impossible models, helping us to see what is true in all of the remaining models



Entailment

Entailment: $\alpha \models \beta$ (“ α entails β ” or “ β follows from α ”) iff in every world where α is true, β is also true

- I.e., the α -worlds are a subset of the β -worlds [$models(\alpha) \subseteq models(\beta)$]

Usually we want to know if $KB \models query$

- $models(KB) \subseteq models(query)$
- In other words
 - KB removes all impossible models (any model where KB is false)
 - If β is true in all of these remaining models, we conclude that β must be true


Entailment and implication are very much related

- However, entailment relates two sentences, while an implication is itself a sentence (usually derived via inference to show entailment)

Wumpus World

Possible Models


- $P_{1,2} P_{2,2} P_{3,1}$
- Knowledge base
 - Nothing in [1,1]
 - Breeze in [2,1]



Wumpus World

Possible Models

- $P_{1,2} P_{2,2} P_{3,1}$
- Knowledge base
 - Nothing in $[1,1]$
 - Breeze in $[2,1]$
- Query α_1 :
 - No pit in $[1,2]$



Possible Models

-

Propositional Logic Models

All Possible Models

Model Symbols

A	0	0	0	0	1	1	1	1
B	0	0	1	1	0	0	1	1
C	0	1	0	1	0	1	0	1

Piazza Poll 1

Does the KB entail query C?

Entailment: $\alpha \models \beta$
“ α entails β ” iff in every world where α is true, β is also true

All Possible Models

Model Symbols	A	0	0	0	0	1	1	1	1
	B	0	0	1	1	0	0	1	1
	C	0	1	0	1	0	1	0	1
Knowledge Base									
	A	0	0	0	0	1	1	1	1
	$B \Rightarrow C$	1	1	0	1	1	1	0	1
	$A \Rightarrow B \vee C$	1	1	1	1	0	1	1	1
Query									
	C	0	1	0	1	0	1	0	1


Piazza Poll 1

Does the KB entail query C?

Yes ☺

Entailment: $\alpha \models \beta$

“ α entails β ” iff in every world where α is true, β is also true



All Possible Models


Model Symbols

Knowledge Base

Query

A	0	0	0	0	1	1	1	1
B	0	0	1	1	0	0	1	1
C	0	1	0	1	0	1	0	1
A	0	0	0	0	1	1	1	1
$B \Rightarrow C$	1	1	0	1	1	1	0	1
$A \Rightarrow B \vee C$	1	1	1	1	0	1	1	1
$\neg B$	0	0	0	0	0	1	0	1
C	0	1	0	1	0	1	0	1

$KB \models C$



Entailment

How do we implement a logical agent that proves entailment?

- Logic language
 - Propositional logic
 - First order logic
- Inference algorithms
 - Theorem proving
 - Model checking

Propositional Logic

Check if sentence is true in given model

In other words, does the model *satisfy* the sentence?



function **PL-TRUE?**(α , model) returns true or false

→ if α is a symbol then return Lookup(α , model)

if Op(α) = \neg then return not(**PL-TRUE?**(Arg1(α), model))

if Op(α) = \wedge then return and(**PL-TRUE?**(Arg1(α), model),
PL-TRUE?(Arg2(α), model))

etc.

(Sometimes called “recursion over syntax”)

Simple Model Checking

$KB \models \alpha$

function **TT-ENTAILS?**(KB, α) returns true or false

return **TT-CHECK-ALL**(KB, α , symbols(KB) \cup symbols(α), {})

function **TT-CHECK-ALL**(KB, α , symbols, model) returns true or false

if empty?(symbols) then

if **PL-TRUE?**(KB, model) then return **PL-TRUE?**(α , model)

else return true

else

$P \leftarrow \text{first}(\text{symbols})$

$\text{rest} \leftarrow \text{rest}(\text{symbols})$

return **and** (**TT-CHECK-ALL**(KB, α , rest, model \cup { $P = \text{true}$ })

TT-CHECK-ALL(KB, α , rest, model \cup { $P = \text{false}$ })))

Simple Model Checking, contd.

Same recursion as backtracking
 $O(2^n)$ time, linear space
We can do much better!

