AI: Representation and Problem Solving

Logical Agent Algorithms

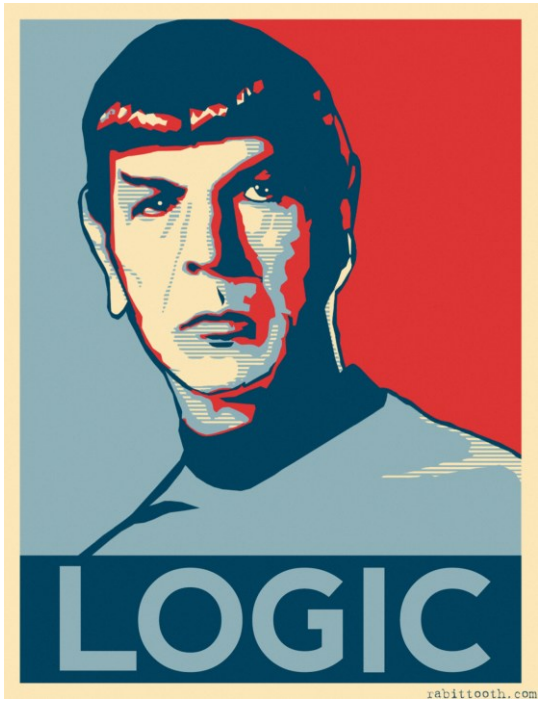

Instructor: Pat Virtue

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

Plan

Last Time:

- Propositional logic
- Models and Knowledge Bases
- Satisifiability and Entailment

Today: Logical Agent Algorithms

- Entailment
  - Model checking: Truth table entailment
  - Theorem proving: (Forward chaining), resolution
- Satisfiability: DPLL
- (Next time) Planning with logic

time

Propositional Logic Vocab

Literal

- Atomic sentence: True, False, Symbol, ¬Symbol

Clause

- Disjunction (OR) of literals: $A \vee B \vee \neg C$

Definite clause

- Disjunction (OR) of literals, *exactly one* is positive
- $\neg A \vee B \vee \neg C$

Horn clause

- Disjunction of literals, *at most one* is positive
- All definite clauses are Horn clauses

PL: Conjunctive Normal Form (CNF)

Every sentence can be expressed as a conjunction (AND) of clauses

Each clause is a disjunction (OR) of literals

Each literal is a symbol or a negated symbol

- Example: $(\neg A \vee \neg C \vee B) \wedge (\neg A \vee \neg B \vee C)$

PL: Conjunctive Normal Form (CNF)

Every sentence can be expressed as a conjunction of clauses

Each clause is a disjunction of literals

Each literal is a symbol or a negated symbol

Conversion to CNF by a sequence of standard transformations:

- $At\_1,1\_0 \Rightarrow (Wall\_0,1 \Leftrightarrow Blocked\_W\_0)$
- $At\_1,1\_0 \Rightarrow ((Wall\_0,1 \Rightarrow Blocked\_W\_0) \wedge (Blocked\_W\_0 \Rightarrow Wall\_0,1))$
- $\neg At\_1,1\_0 \vee ((\neg Wall\_0,1 \vee Blocked\_W\_0) \wedge (\neg Blocked\_W\_0 \vee Wall\_0,1))$
- $(\neg At\_1,1\_0 \vee \neg Wall\_0,1 \vee Blocked\_W\_0) \wedge (\neg At\_1,1\_0 \vee \neg Blocked\_W\_0 \vee Wall\_0,1)$

PL: Conjunctive Normal Form (CNF)

Every sentence can be expressed

Each clause is a disjunction of literal

Each literal is a symbol or a neg sym

Conversion to CNF by a se quence standard transform s:

- At_1,1_0 $\Rightarrow$ (Wall_0,1 $\Leftrightarrow$ Block W_0)
- At_1,1_0 $\Rightarrow$ ((Wall_0,1 $\Rightarrow$ Blocked_W_0) $\wedge$ (Bloc d_W_0 $\Rightarrow$Wall_0,1))
- $\neg$At_1,1_0 v (($\neg$Wall_0,1 v Blocked_W_0) $\wedge$ ($\neg$Blocked_W_0 v Wall_0,1))
- ($\neg$At_1,1_0 v $\neg$Wall_0,1 v Blocked_W_0) $\wedge$ ($\neg$At_1,1_0 v $\neg$Blocked_W_0 v Wall_0,1)

Replace biconditional by two implications

Replace $\alpha \Rightarrow \beta$ by $\neg\alpha$ v β

Distribute v over $\wedge$

Logical Agent Vocab

Model

- Complete assignment of symbols to True/False

Sentence

- Logical statement
- Composition of logic symbols and operators

KB

- Collection of sentences representing facts and rules we know about the world

Query

- Sentence we want to know if it is *provably* True, *provably* False, or *unsure*.

Provably True, Provably False, or Unsure

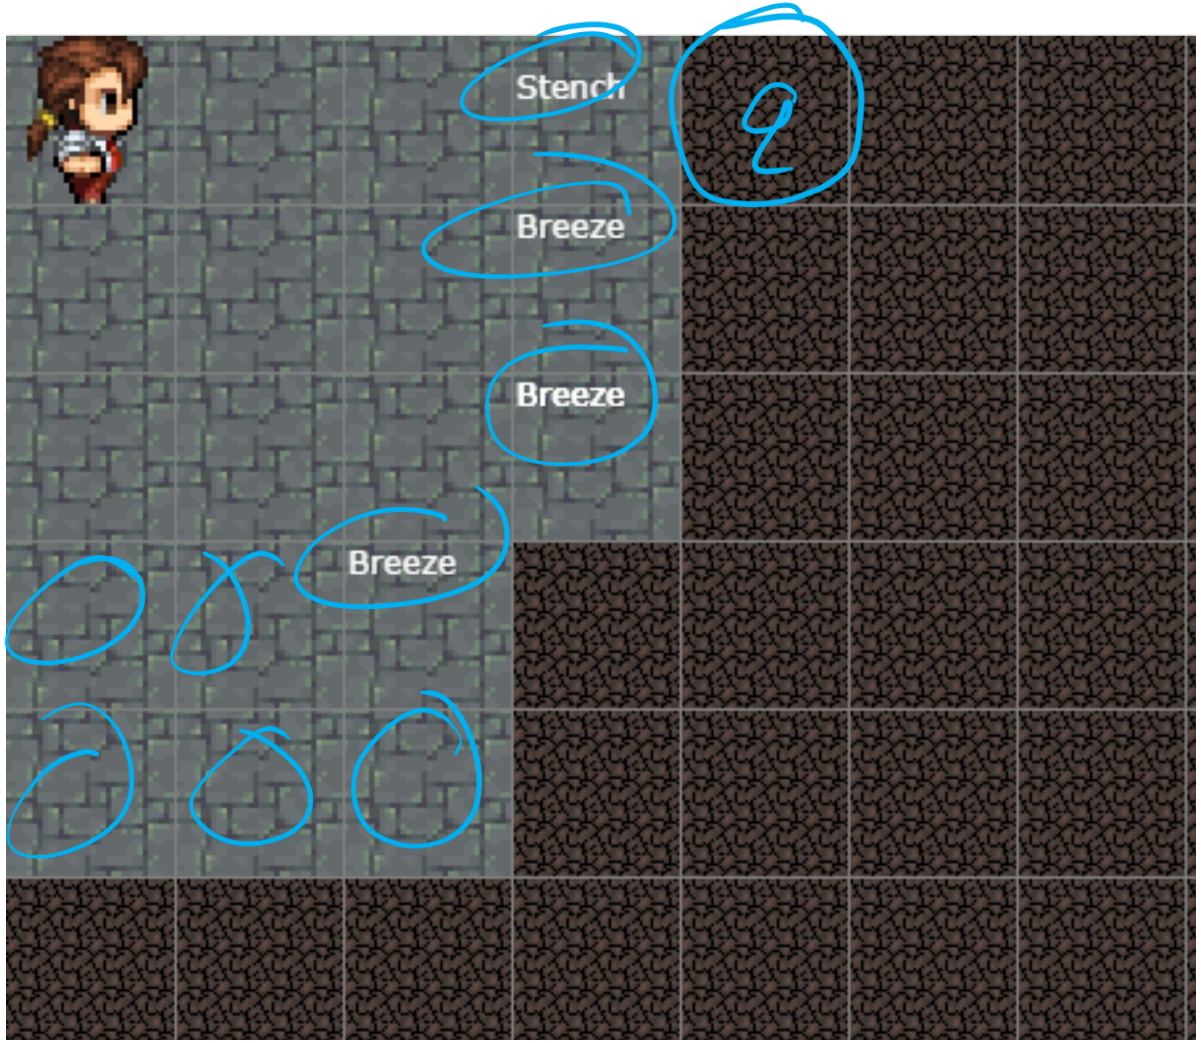

$KB \models q$

entails

if yes

q is

definitely

True

http://thiagodnf.github.io/wumpus-world-simulator/

Logical Agent Vocab

Entailment

- Input: sentence1, sentence2
- Each model that satisfies sentence1 must also satisfy sentence2
- "If I know 1 holds, then I know 2 holds"
- (ASK), TT-ENTAILS, FC-ENTAILS, RESOLUTION-ENTAILS

Satisfy (model, sentence)

satisfiable

- Input: model, sentence
- Is this sentence true in this model?
- Does this model satisfy this sentence
- "Does this particular state of the world work?'
- PL-TRUE

Logical Agent Vocab

Satisfiable

- Input: sentence
- Can find at least one model that satisfies this sentence
  - (We often want to know what that model is)
- "Is it possible to make this sentence true?"
- DPLL

Valid

- Input: sentence
- sentence is true in all possible models

Outline

Logical Agent Algorithms

- Vocab
- PL_TRUE
- Entailment
  - Model checking: Truth table entailment
  - Theorem proving:
  - (Forward chaining), resolution
- Satisfiability: DPLL
- Planning with logic

Propositional Logic

Check if sentence is true in given model

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

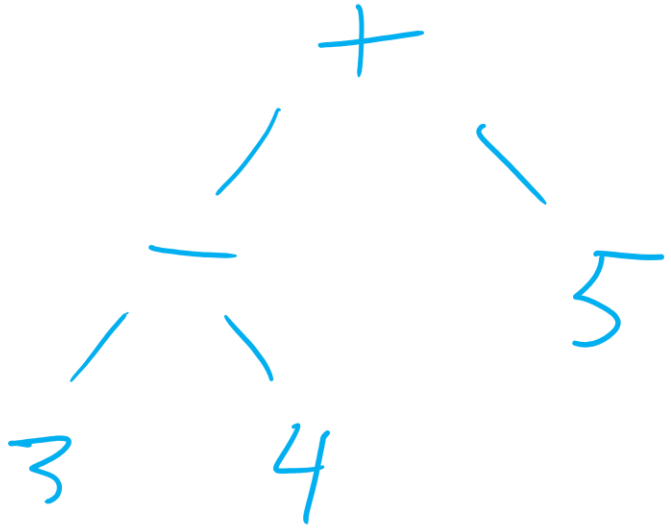

function PL-TRUE?(α,model) returns true or false

But are models and propositional logic sentences α represented?

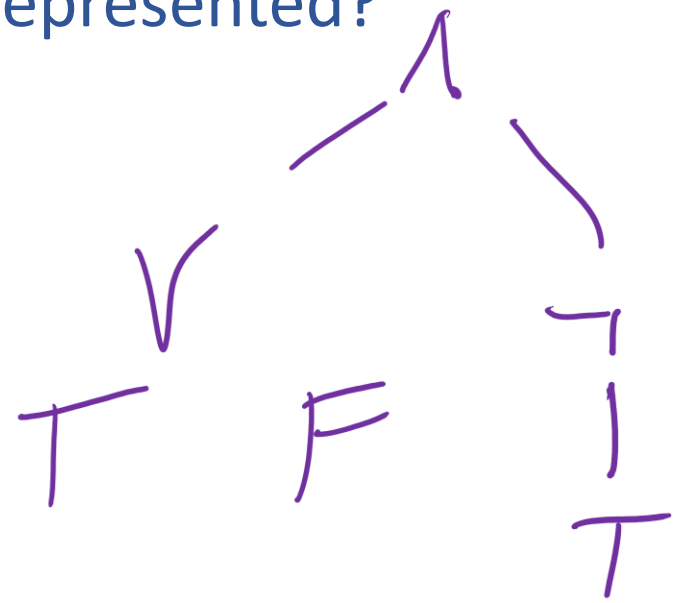

Propositional Logic

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”)

Outline

Logical Agent Algorithms

- Vocab
- PL_TRUE
- Entailment
  - Model checking: Truth table entailment
  - Theorem proving:
  - Forward chaining, resolution
- Satisfiability: DPLL
- Planning with logic

Inference: Proofs

A proof is a *demonstration* of entailment between α and β

Method 1: *model-checking*

- For every possible world, if α is true make sure that is β true too
- OK for propositional logic (finitely many worlds); not easy for first-order logic

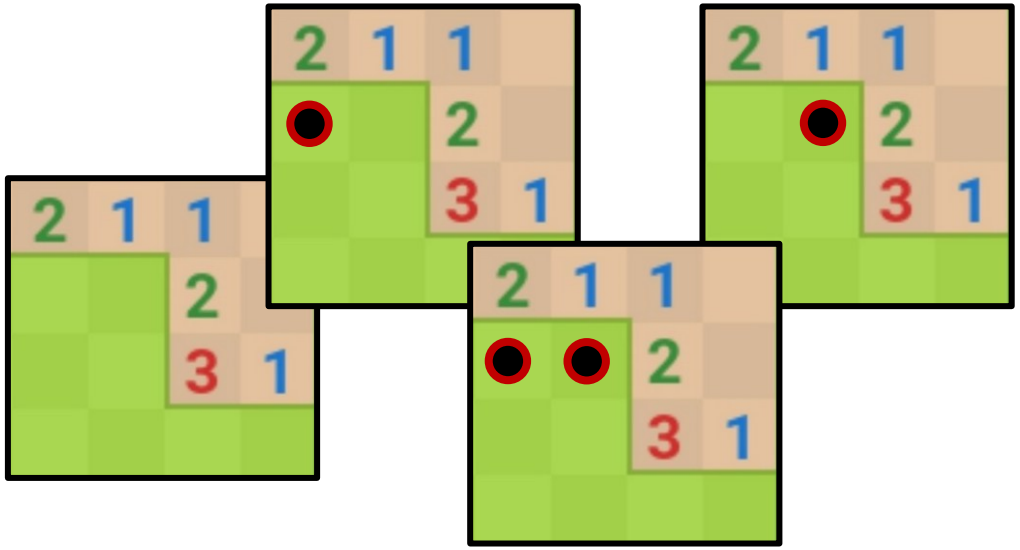

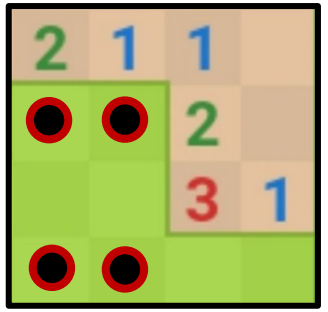

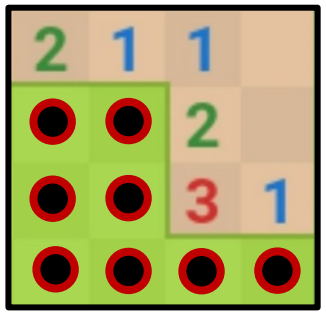

Simple Model Checking

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

Simple Model Checking

function TT-ENTAILS?(KB, α)

Same recursion as backtracking

$KB \models \alpha$

every model

for

where

PL_TRUE(model, KB)

then

PL_TRUE(model, α)

must be true

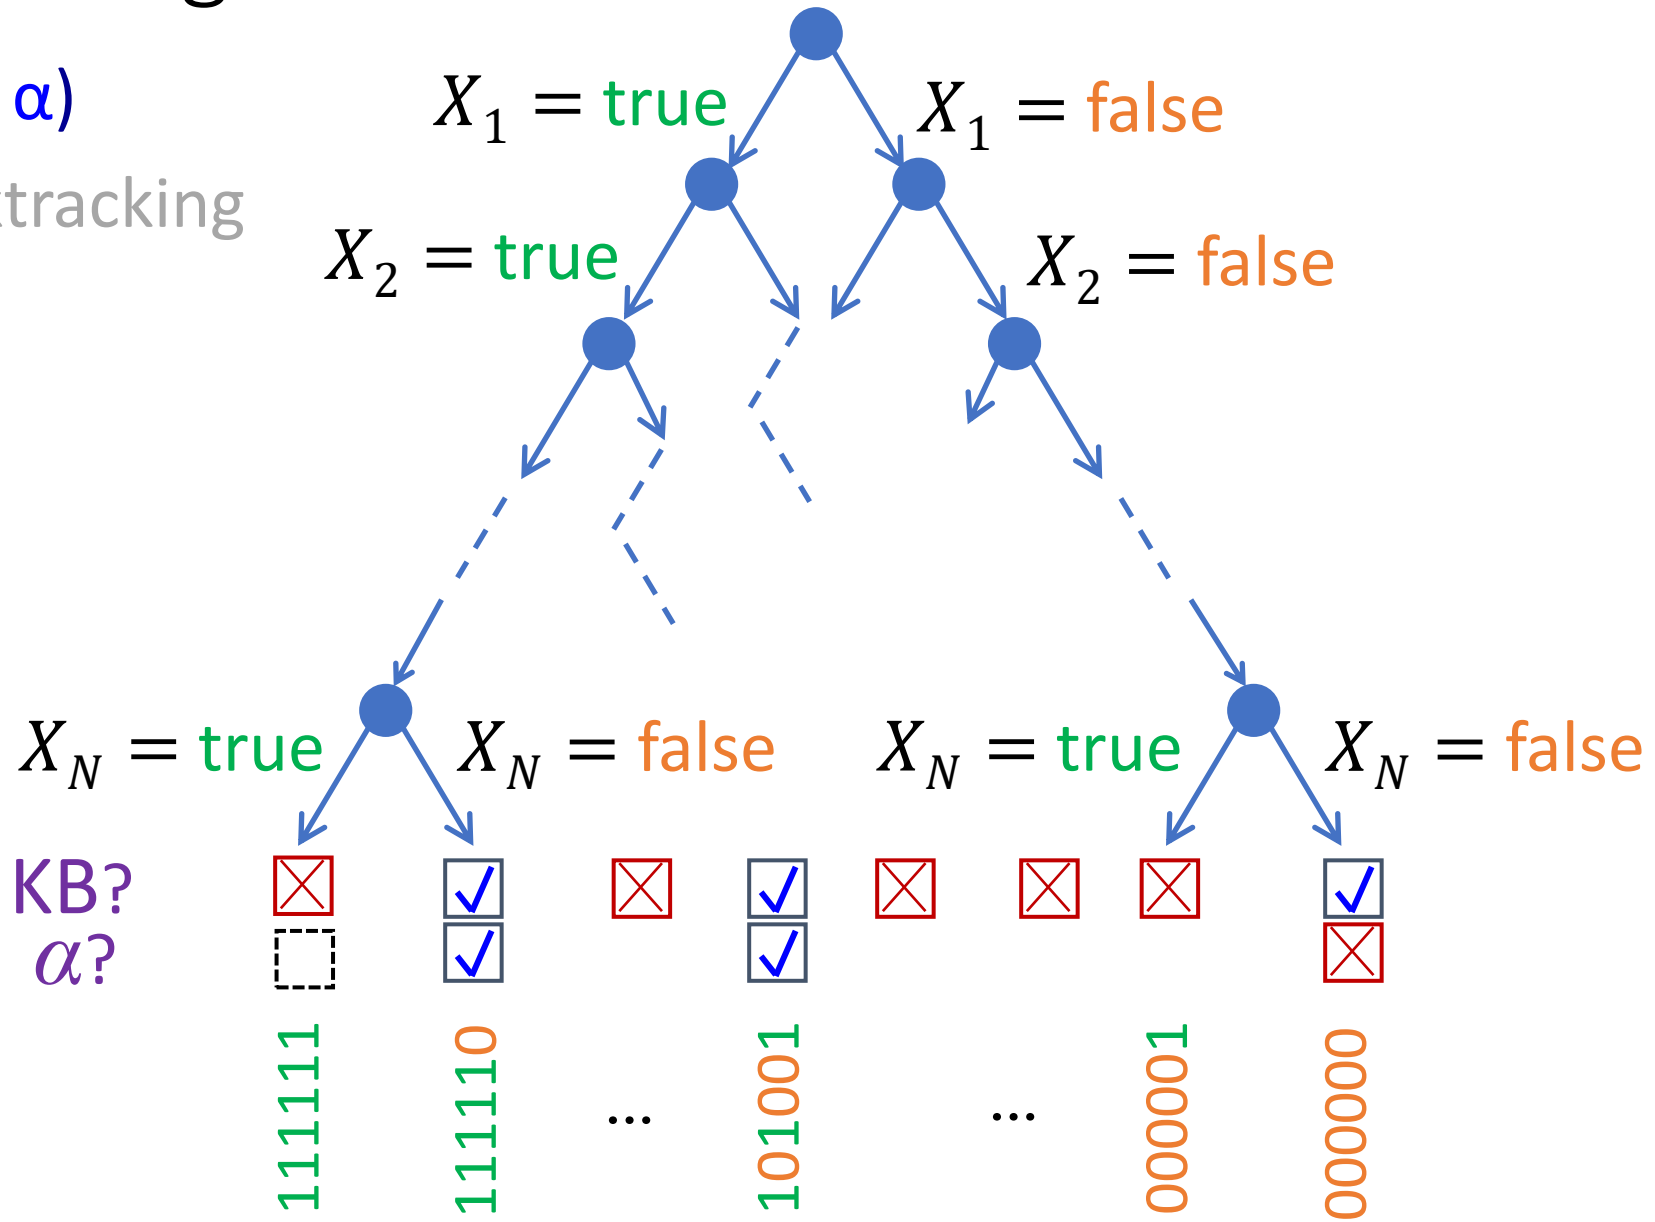

Simple Model Checking

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

    return TT-CHECK-ALL(KB, α, symbols(KB) ∪ symbols(α), {})

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

    Recursively check to make sure all models

    that satisfy the KB also satisfy α

Simple Model Checking

```
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
        X_i ← first(symbols)
        rest ← rest(symbols)
        return  and ( TT-CHECK-ALL(KB, α, rest, model ∪ {X_i = true})
                      TT-CHECK-ALL(KB, α, rest, model ∪ {X_i = false}) )
```

Simple Model Checking

Same recursion as backtracking

$O(2^N)$ time, linear space

O(N)

DFS

Can we do better?

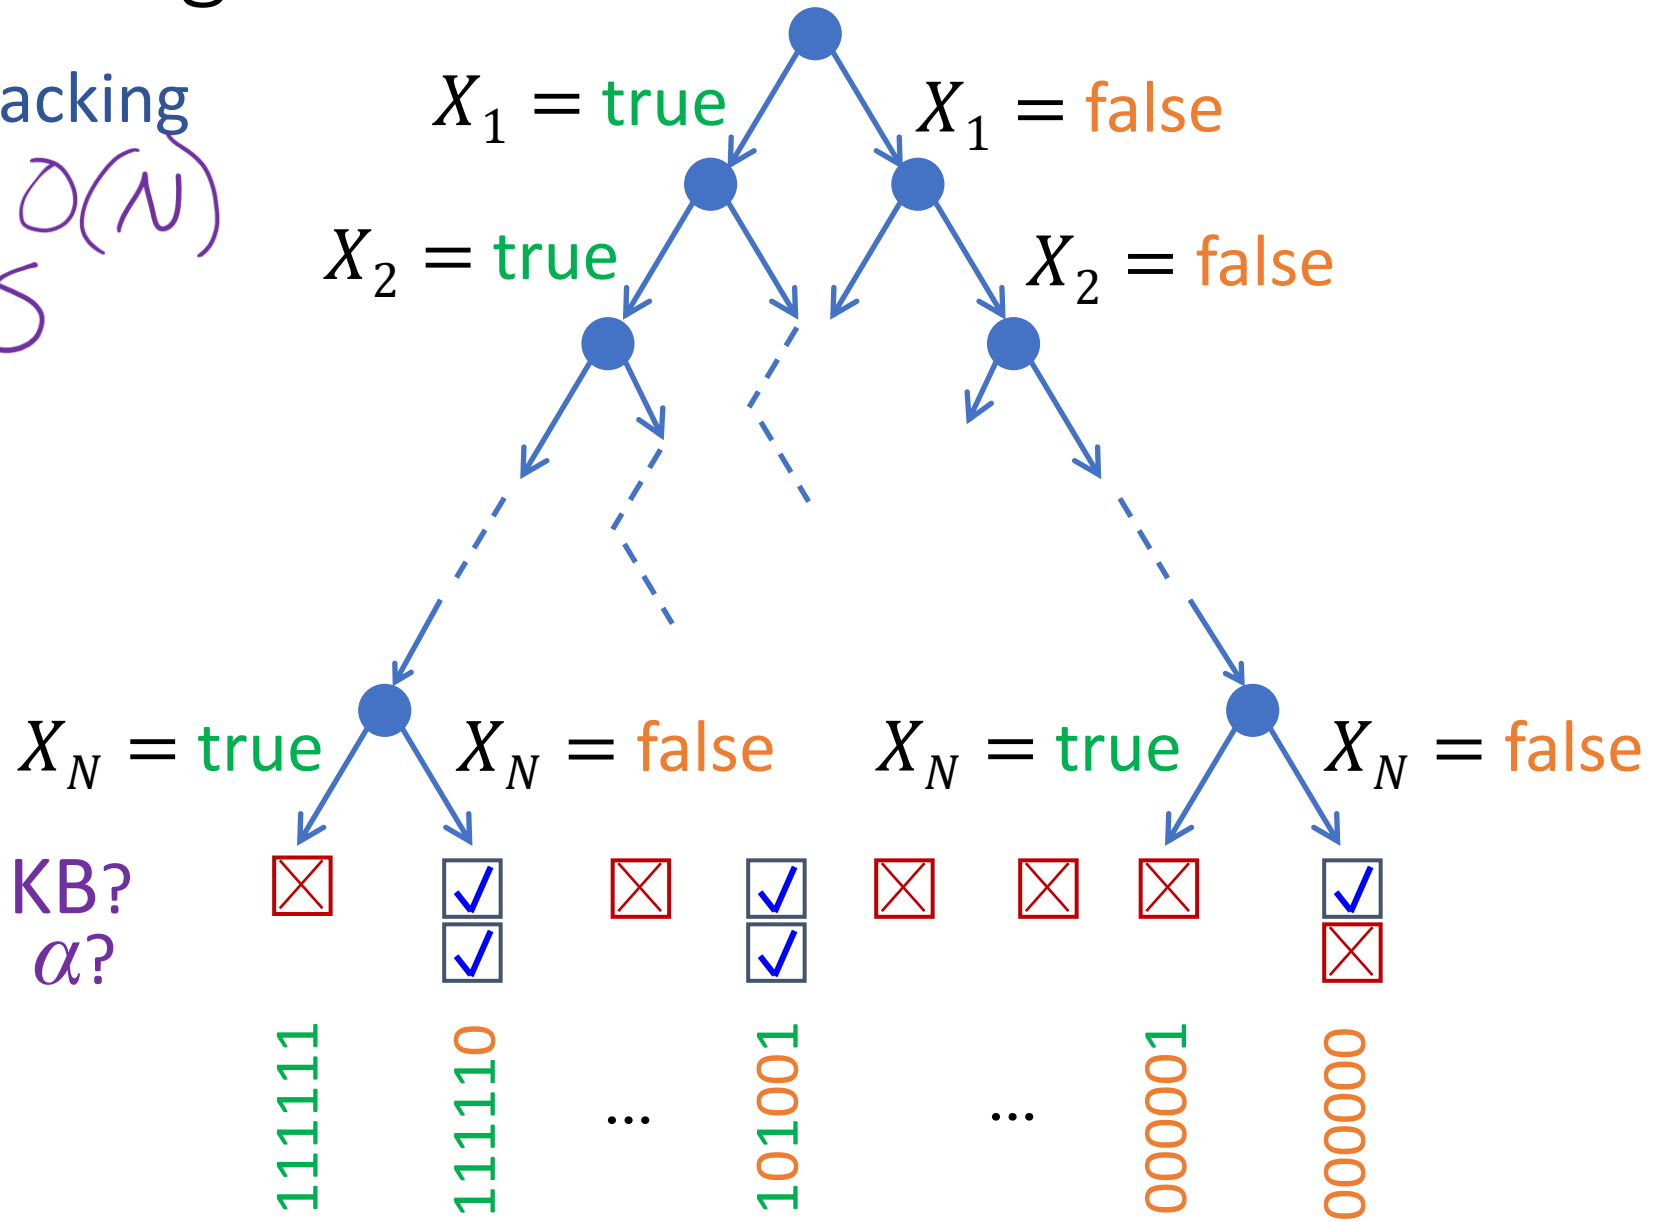

Inference: Proofs

A proof is a *demonstration* of entailment between α and β

Method 1: *model-checking*

- For every possible world, if α is true make sure that is β true too
- OK for propositional logic (finitely many worlds); not easy for first-order logic

Method 2: *theorem-proving*

- Search for a sequence of proof steps (applications of *inference rules*) leading from α to β
- E.g., from $P \wedge (P \Rightarrow Q)$, infer Q by *Modus Ponens*

Properties

- *Sound* algorithm: everything it claims to prove is in fact entailed
- *Complete* algorithm: every sentence that is entailed can be proved

Simple Theorem Proving: Forward Chaining

Forward chaining applies Modus Ponens to generate new facts:

- Given $X_1 \wedge X_2 \wedge \ldots X_n \Rightarrow Y$ and X_1, X_2, …, X_n
- Infer Y

Forward chaining keeps applying this rule, adding new facts, until nothing more can be added

Requires KB to contain only *definite clauses*:

- (Conjunction of symbols) $\Rightarrow$ symbol; or
- A single symbol (note that X is equivalent to True $\Rightarrow X$)

Forward Chaining Algorithm

function PL-FC-ENTAILS?(KB, q) Returns true or false

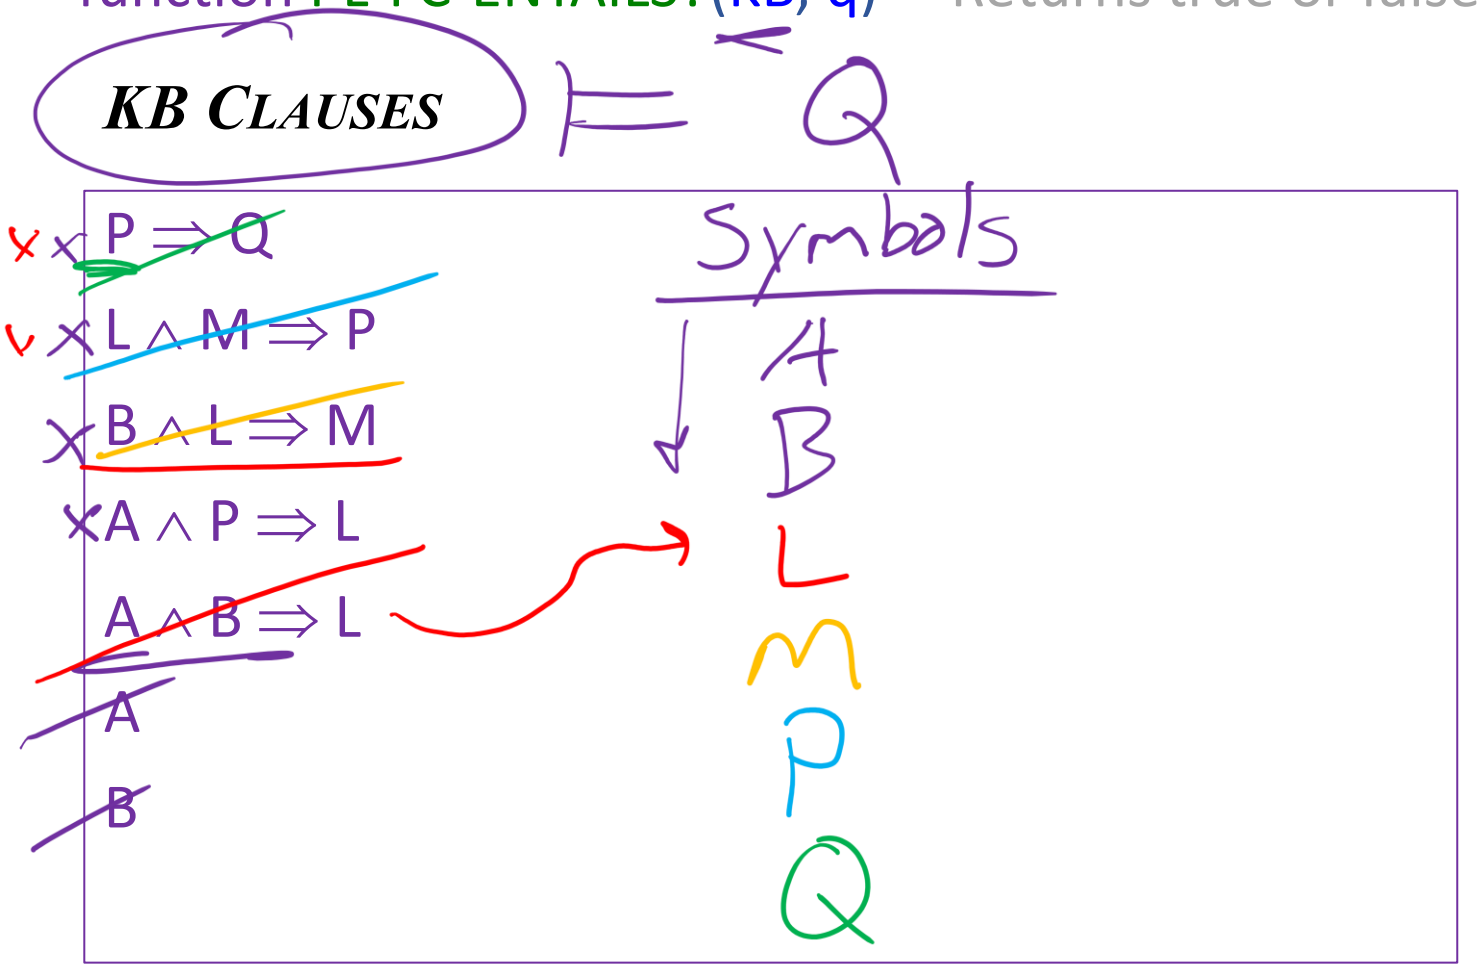

Forward Chaining Algorithm

function PL-FC-ENTAILS?(KB, q) returns true or false

count ← a table, where count[c] is the number of symbols in c's premise

inferred ← a table, where inferred[s] is initially false for all s

agenda ← a queue of symbols, initially symbols known to be true in KB

CLAUSES	COUNT	INFERRED	AGENDA
$P \Rightarrow Q$	1	A false	
$L \wedge M \Rightarrow P$	2	B false	
$B \wedge L \Rightarrow M$	2	L false	
$A \wedge P \Rightarrow L$	2	M false	
$A \wedge B \Rightarrow L$	2	P false	
A	0	Q false	
B	0		

Forward Chaining Example: Proving Q

Clauses	*Count*	*Inferred*
$P \Rightarrow Q$	1 / 0	A false true
$L \wedge M \Rightarrow P$	2 / 1 / 0	B false true
$B \wedge L \Rightarrow M$	2 / 1 / 0	L false true
$A \wedge P \Rightarrow L$	2 / 1 / 0	M false true
$A \wedge B \Rightarrow L$	2 / 1 / 0	P false true
A	0	Q false true
B	0	

Agenda

A B L M P Q

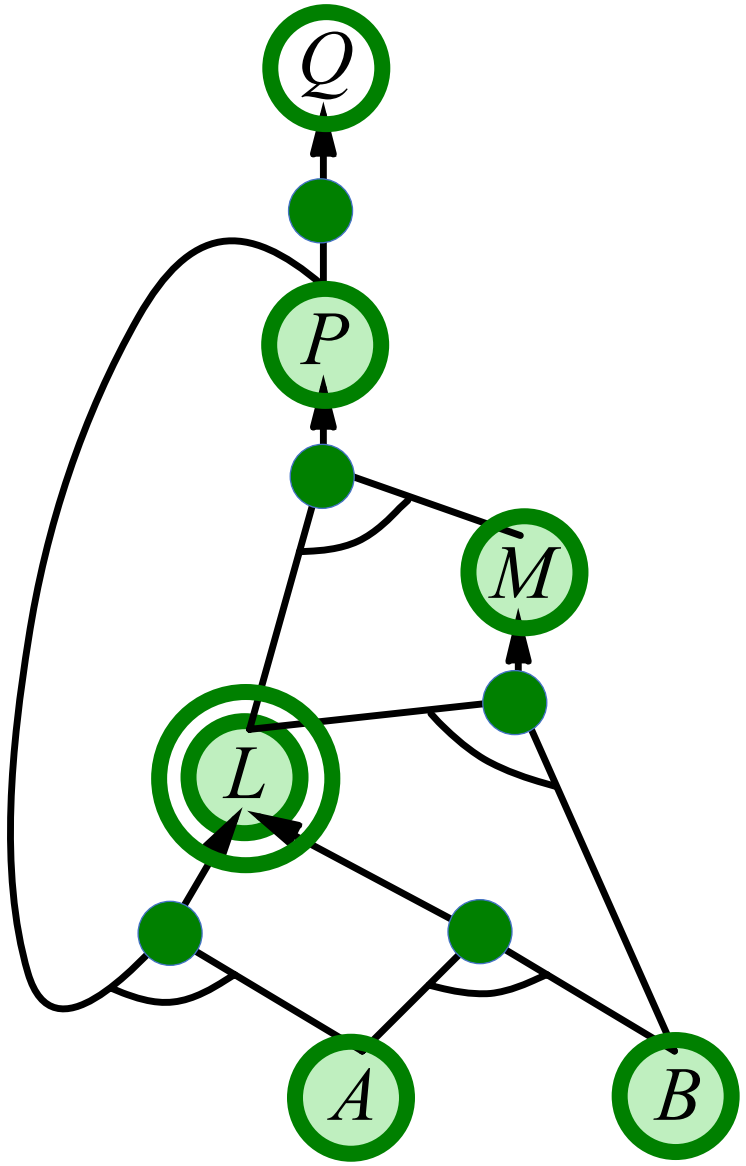

Forward Chaining Algorithm

```
function PL-FC-ENTAILS?(KB, q) returns true or false
    count ← a table, where count[c] is the number of symbols in c's premise
    inferred ← a table, where inferred[s] is initially false for all s
    agenda ← a queue of symbols, initially symbols known to be true in KB
    while agenda is not empty do
        p ← Pop(agenda)
        if p = q then return true
        if inferred[p] = false then
            inferred[p]←true
            for each clause c in KB where p is in c.premise do
                decrement count[c]
                if count[c] = 0 then add c.conclusion to agenda
    return false
```

Properties

Forward Chaining is:

- Sound and complete for definite-clause KBs
- Complexity: linear time ☺

Resolution is another theorem-proving algorithm that is:

- Sound and complete for any PL KBs!
- Complexity: exponential time ☹

Vocab Reminder

Literal

- Atomic sentence:
  T, F, Symbol, ¬Symbol

Clause

- Disjunction of literals:
  $A \vee B \vee \neg C$

Definite clause

- Disjunction of literals, *exactly one* is positive
  $\neg A \vee B \vee \neg C$

Inference Rules

Notation Alert!

Modus Ponens

$$\frac{\alpha \Rightarrow \beta, \quad \alpha}{\beta}$$

$$[(\alpha \Rightarrow \beta) \wedge \alpha] \Rightarrow \beta$$

Unit Resolution

$$\frac{a \vee b, \quad \neg b \vee c}{a \vee c}$$

resolve only one symbol at a time

General Resolution

$$\frac{z \vee a_1 \vee \cdots \vee a_m \vee b, \quad \neg b \vee c_1 \vee \cdots \vee c_n \vee \neg z}{a_1 \vee \cdots \vee a_m \vee c_1 \vee \cdots \vee c_n}$$

Resolution

Algorithm Overview

function PL-RESOLUTION?(KB, α) returns true or false

We want to prove that KB entails α

In other words, we want to prove that we cannot satisfy (KB and **not** α)

1. Start with a set of CNF clauses, including the KB as well as $\neg\alpha$
2. Keep resolving pairs of clauses until
   A. You resolve the empty clause
      Contradiction found!
      KB $\wedge \neg\alpha$ cannot be satisfied
      Return true, KB entails α
   B. No new clauses added
      Return false, KB does not entail α

$KB \models \alpha$
entails(KB, α)
result = SAT($KB \wedge \neg\alpha$)
if ! result
return TRUE

Resolution

Example trying to prove $\neg P_{1,2}$

General Resolution

$$\frac{a_1 \vee \cdots \vee a_m \vee b, \quad \neg b \vee c_1 \vee \cdots \vee c_n}{a_1 \vee \cdots \vee a_m \vee c_1 \vee \cdots \vee c_n}$$

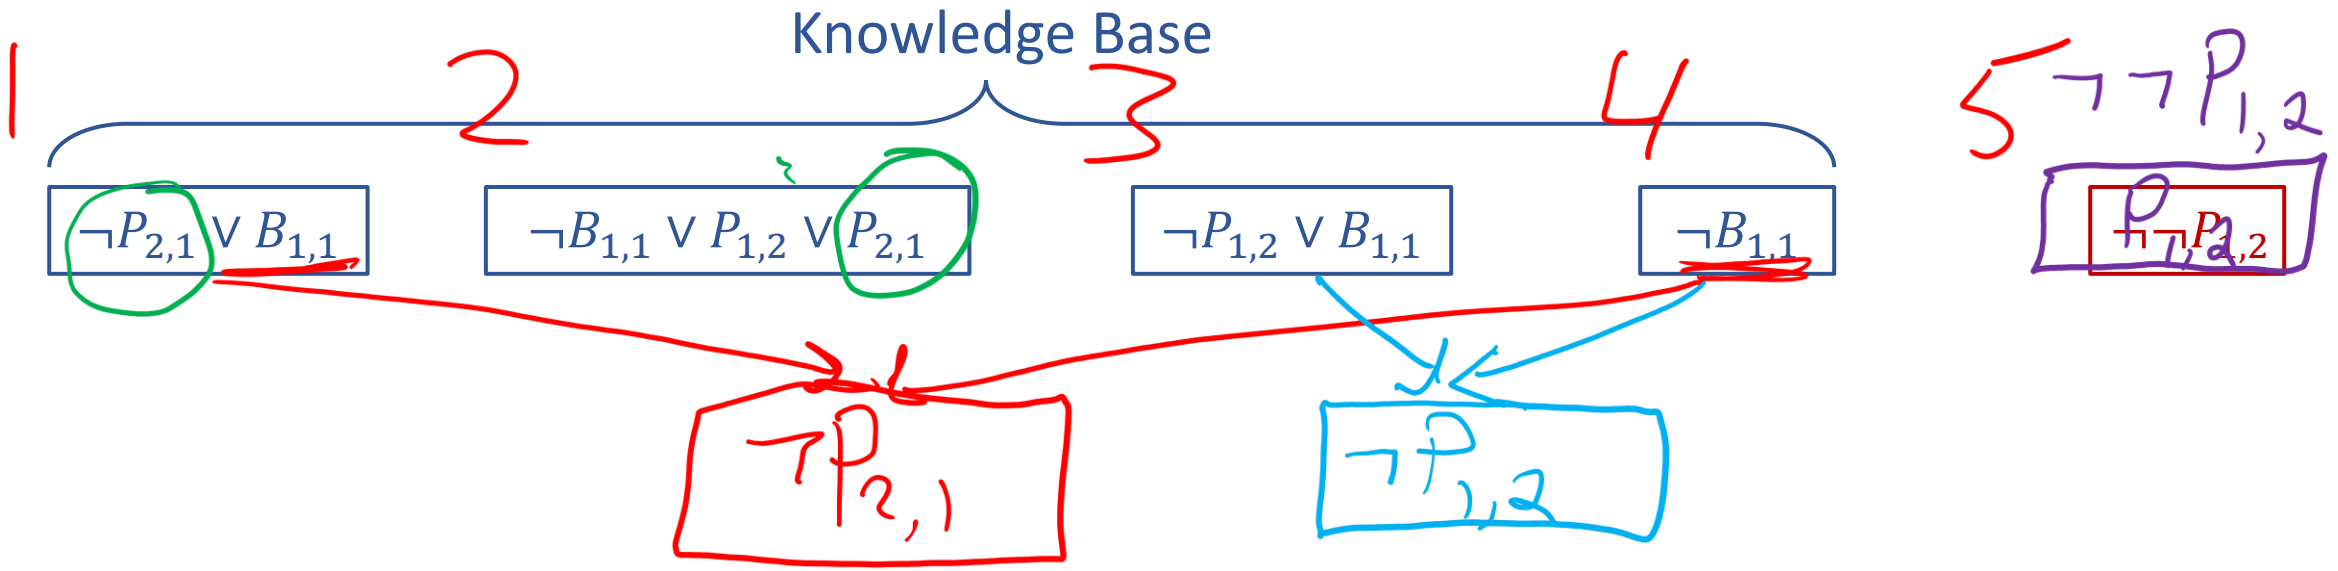

Resolution

Example trying to prove $\neg P_{1,2}$

General Resolution

$$\frac{a_1 \vee \cdots \vee a_m \vee b, \quad \neg b \vee c_1 \vee \cdots \vee c_n}{a_1 \vee \cdots \vee a_m \vee c_1 \vee \cdots \vee c_n}$$

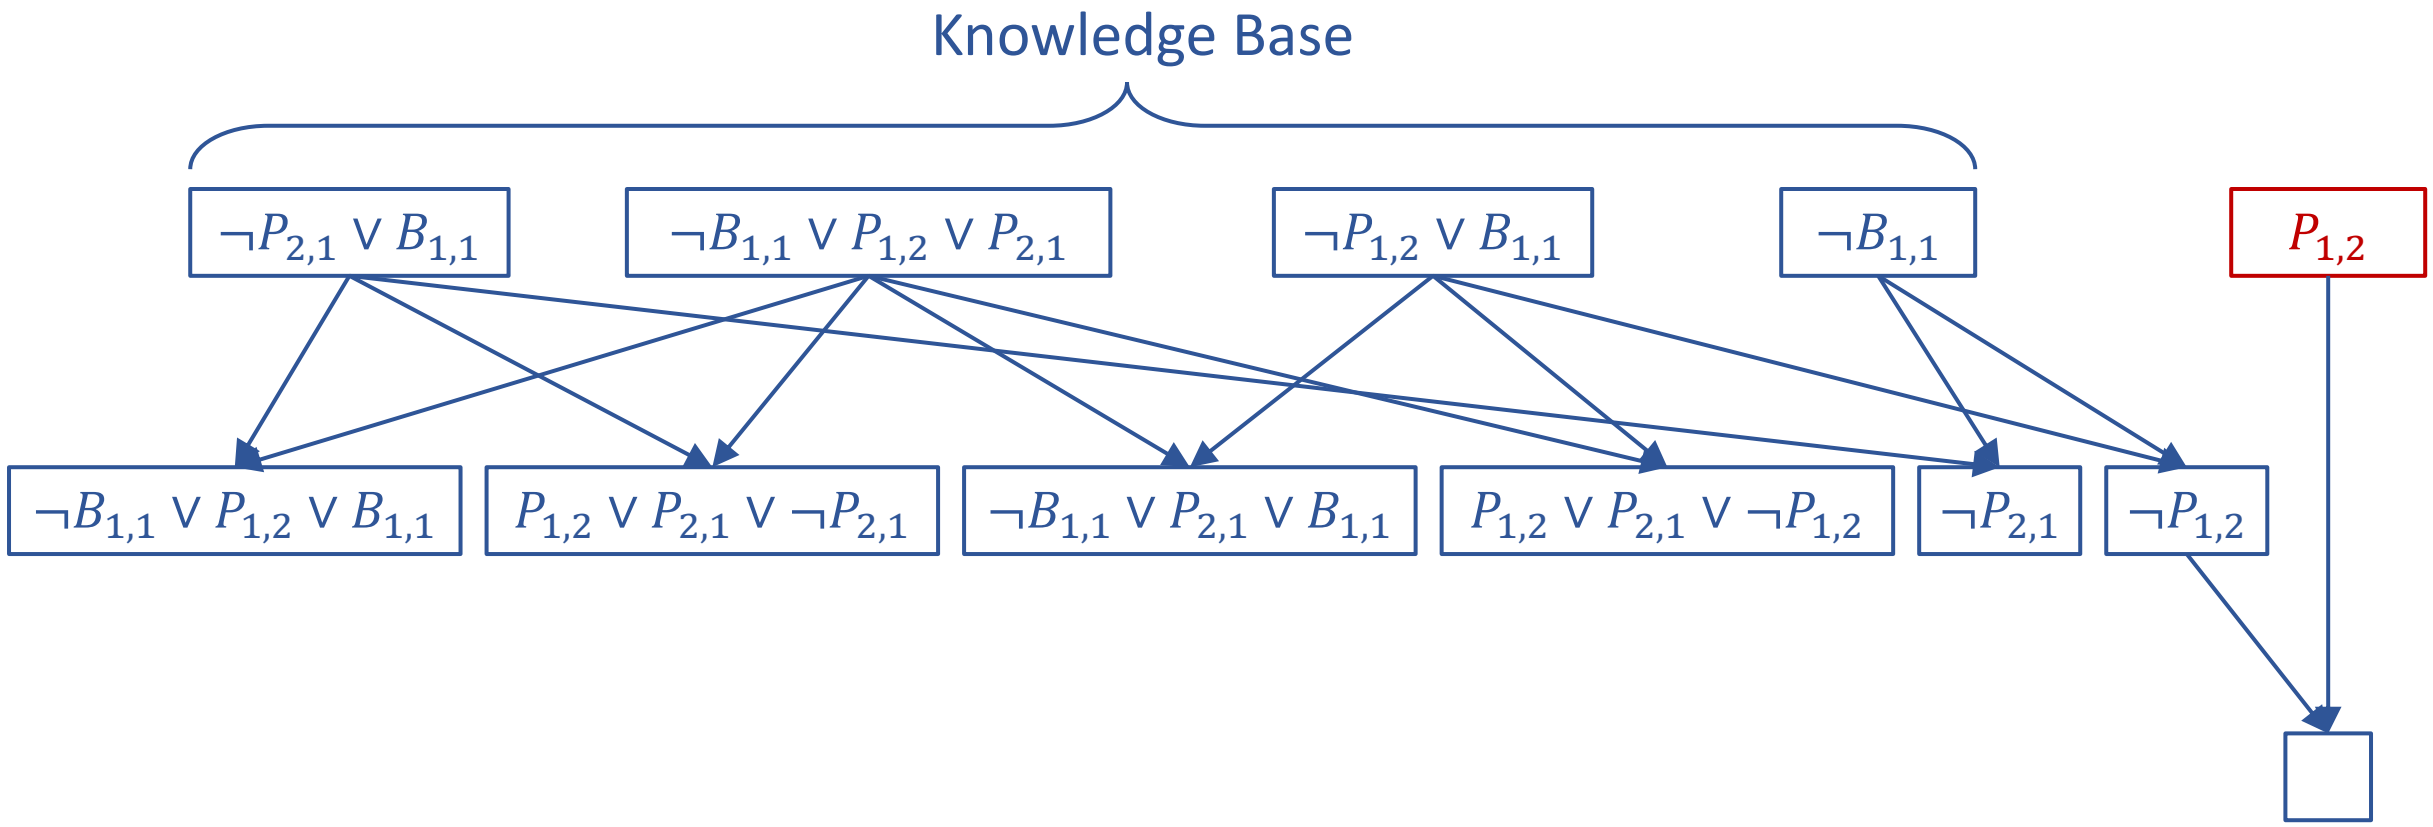

Resolution

function PL-RESOLUTION?(KB, α) returns true or false

clauses ← the set of clauses in the CNF representation of KB $\wedge \neg\alpha$

new ← { }

loop do

for each pair of clauses C_i, C_j in clauses do

resolvents ← PL-RESOLVE(C_i, C_j)

if resolvents contains the empty clause then

return true

new ← new ∪ resolvants

if new ⊆ clauses then

return false

clauses ← clauses ∪ new

Properties

Forward Chaining is:

- Sound and complete for definite-clause KBs
- Complexity: linear time ☺

Resolution is another theorem-proving algorithm that is:

- Sound and complete for any PL KBs!
- Complexity: exponential time ☹

Vocab Reminder

Literal

- Atomic sentence:
  T, F, Symbol, ¬Symbol

Clause

- Disjunction of literals:
  $A \vee B \vee \neg C$

Definite clause

- Disjunction of literals, *exactly one* is positive
  $\neg A \vee B \vee \neg C$

Outline

Logical Agent Algorithms

- Vocab
- PL_TRUE
- Entailment
  - Model checking: Truth table entailment
  - Theorem proving:
  - Forward chaining, resolution
- Satisfiability: DPLL
- Planning with logic

(sentence) → None or model

KB ∧ ¬q

list of clauses, symbols, model

Satisfiability and Entailment

A sentence is *satisfiable* if it is true in at least one world (e.g. CSPs!)

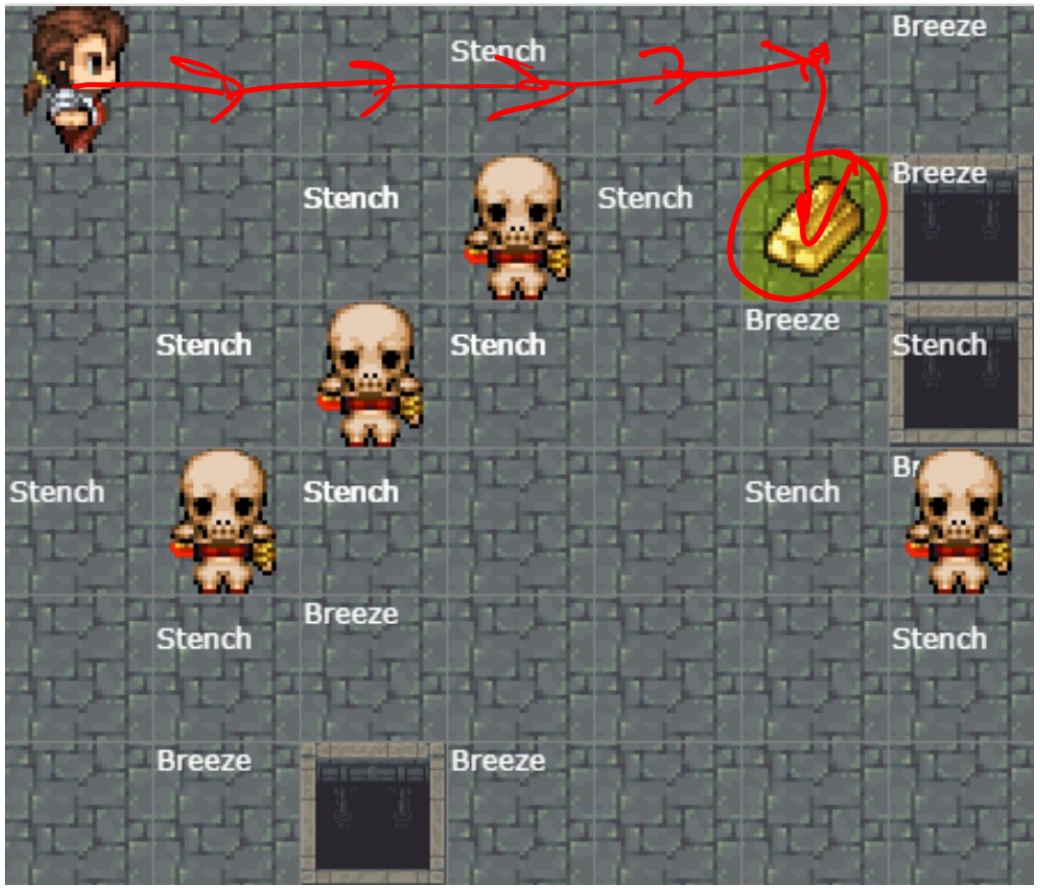

http://thiagodnf.github.io/wumpus-world-simulator/

Satisfiability and Entailment

A sentence is *satisfiable* if it is true in at least one world (cf CSPs!)

Suppose we have a hyper-efficient SAT solver; how can we use it to test entailment?

- Suppose $\alpha \models \beta$
- Then $\alpha \Rightarrow \beta$ is true in all worlds
- Hence $\neg(\alpha \Rightarrow \beta)$ is false in all worlds
- Hence $\alpha \wedge \neg\beta$ is false in all worlds, i.e., unsatisfiable

So, add the negated conclusion to what you know, test for (un)satisfiability; also known as reductio ad absurdum

Efficient SAT solvers operate on ***conjunctive normal form***

Efficient SAT solvers

DPLL (Davis-Putnam-Logemann-Loveland) is the core of modern solvers

Essentially a backtracking search over models with some extras:

- *Early termination*: stop if
  - all clauses are satisfied; e.g., $(A \vee B) \wedge (A \vee \neg C)$ is satisfied by {A=true}
  - any clause is falsified; e.g., $(A \vee B) \wedge (A \vee \neg C)$ is satisfied by {A=false, B=false}

- *Pure literals*: if all occurrences of a symbol in as-yet-unsatisfied clauses have the same sign, then give the symbol that value
  - E.g., A is pure and positive in $(A \vee B) \wedge (A \vee \neg C) \wedge (C \vee \neg B)$ so set it to true

- *Unit clauses*: if a clause is left with a single literal, set symbol to satisfy clause
  - E.g., if A=false, $(A \vee B) \wedge (A \vee \neg C)$ becomes $(\text{false} \vee B) \wedge (\text{false} \vee \neg C)$, i.e. $(B) \wedge (\neg C)$
  - Satisfying the unit clauses often leads to further propagation, new unit clauses, etc.

DPLL algorithm

```
function DPLL(clauses, symbols, model) returns true or false
    if every clause in clauses is true in model then return true
    if some clause in clauses is false in model then return false

    P, value ←FIND-PURE-SYMBOL(symbols, clauses, model)
    if P is non-null then return DPLL(clauses, symbols–P, modelU{P=value})

    P, value ←FIND-UNIT-CLAUSE(clauses, model)
    if P is non-null then return DPLL(clauses, symbols–P, modelU{P=value})

    P ← First(symbols)
    rest ← Rest(symbols)

    return or(DPLL(clauses, rest, modelU{P=true}),
              DPLL(clauses, rest, modelU{P=false}))
```

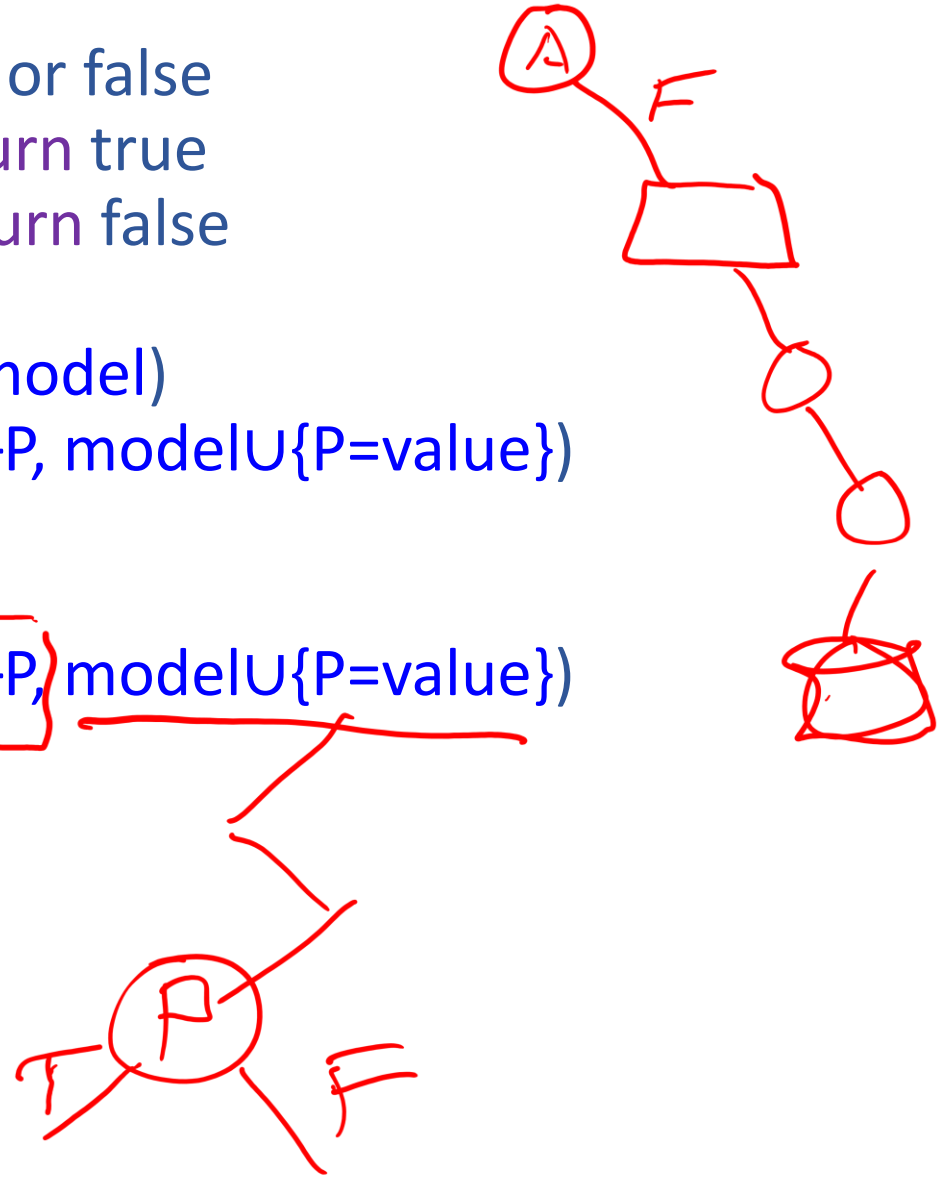

Outline

Logical Agent Algorithms

- Vocab
- PL_TRUE
- Entailment
  - Model checking: Truth table entailment
  - Theorem proving:
  - (Forward chaining), resolution
- Satisfiability: DPLL
- Planning with logic

Planning as Satisfiability

Given a hyper-efficient SAT solver, can we use it to make plans?

Yes, for fully observable, deterministic case: planning problem is solvable iff there is some satisfying assignment for actions etc.

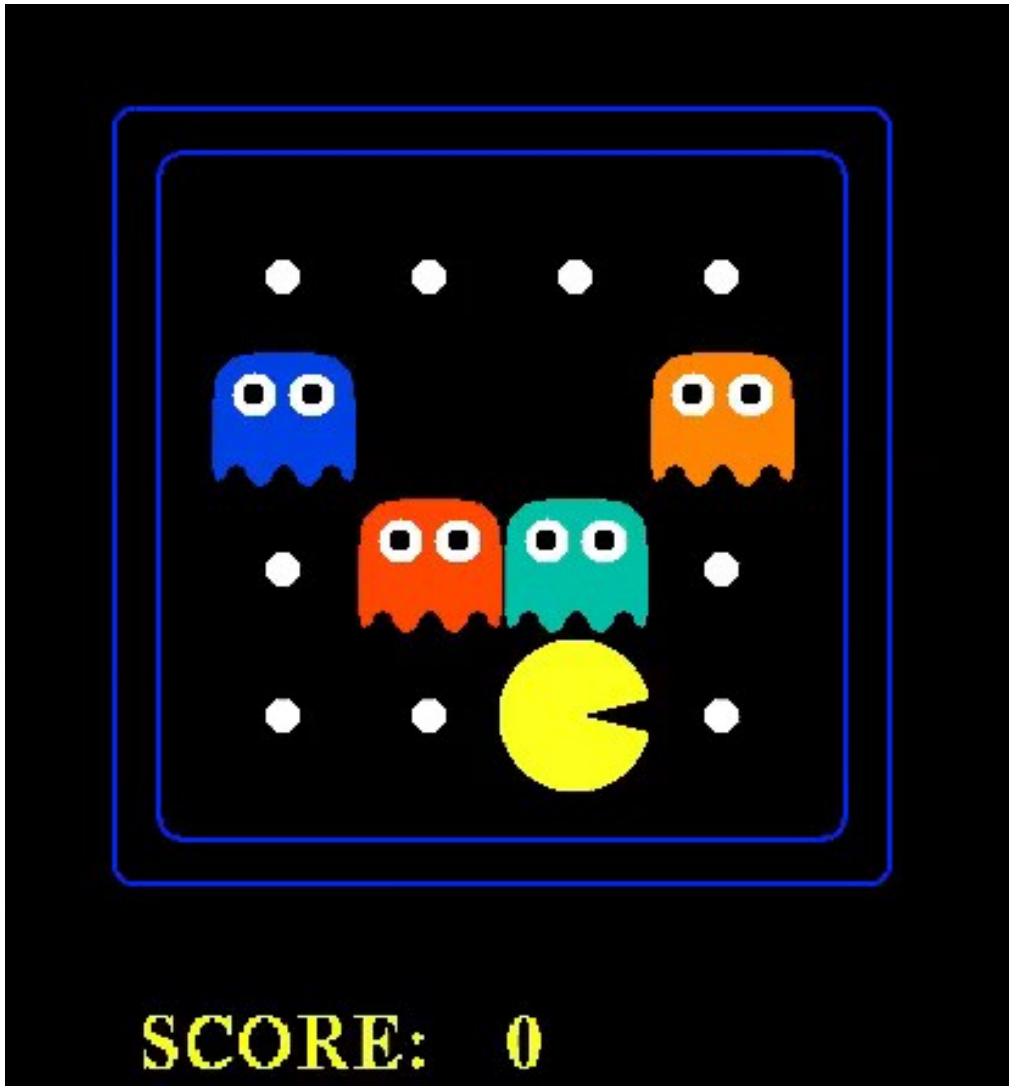
SCORE: 0

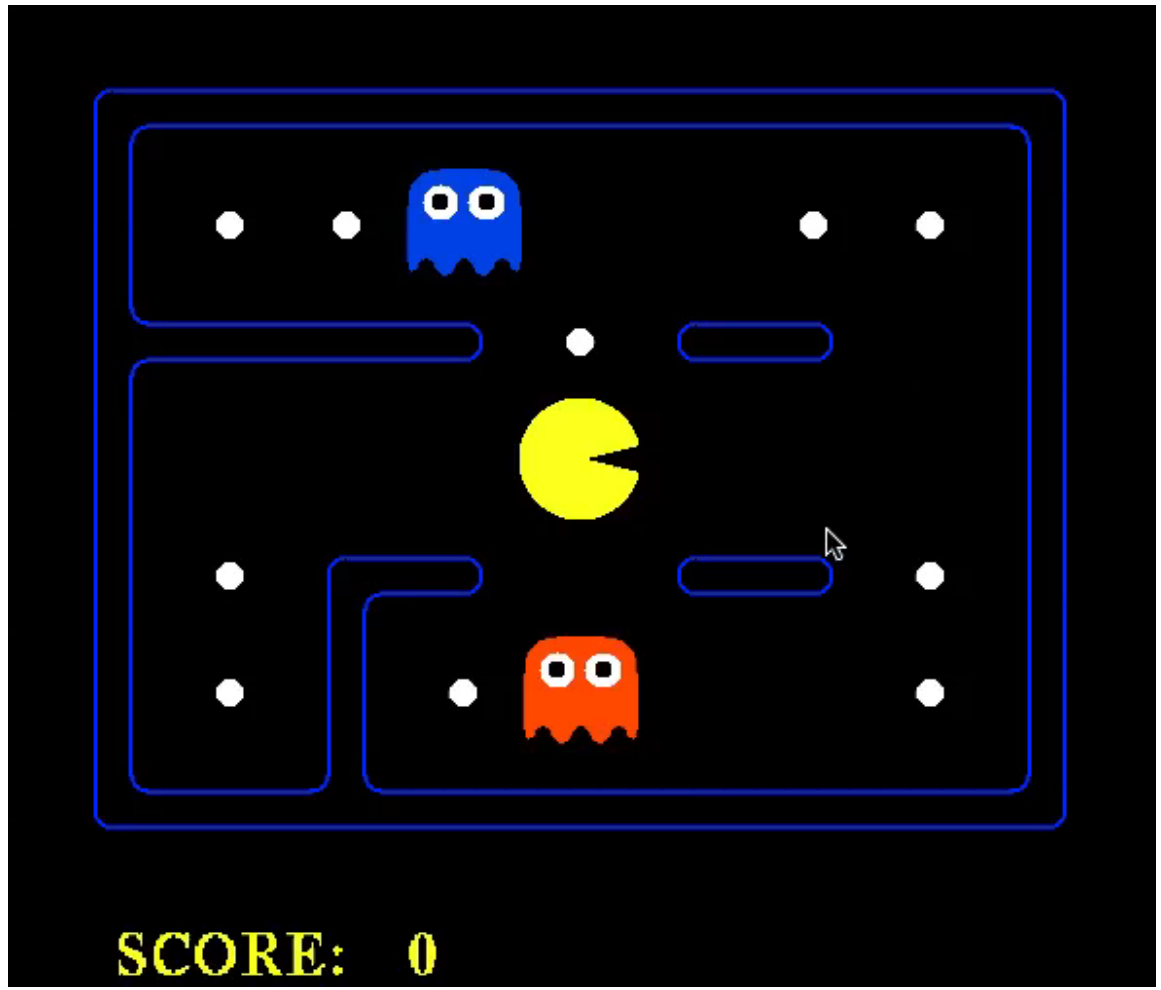
SCORE: 0

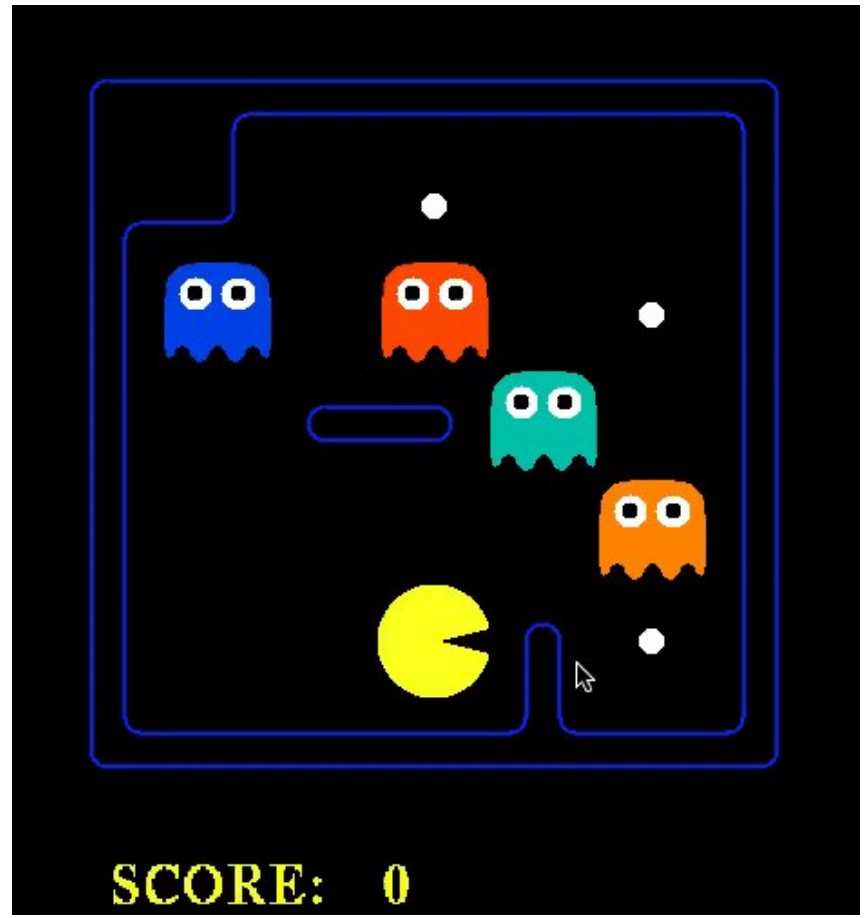

2^N

Planning as Satisfiability

Given a hyper-efficient SAT solver, can we use it to make plans?

Yes, for fully observable, deterministic case: planning problem is solvable iff there is some satisfying assignment for actions etc.

$P_{1,2,t}$

For T = 1 to infinity, set up the KB as follows and run SAT solver:

- Initial state, domain constraints
- Transition model sentences up to time T
- Goal is true at time T
- *Precondition axioms*: $At\_1,1\_0 \wedge N\_0 \Rightarrow \neg Wall\_1,2$ etc.
- *Action exclusion axioms*: $\neg(N\_0 \wedge W\_0) \wedge \neg(N\_0 \wedge S\_0) \wedge ..$ etc.

Initial State

The agent may know its initial location:

- At_1,1_0

Or, it may not:

- At_1,1_0 v At_1,2_0 v At_1,3_0 v … v At_3,3_0

We also need a *domain constraint* – cannot be in two places at once!

- $\neg$(At_1,1_0 $\wedge$ At_1,2_0) $\wedge$ $\neg$(At_1,1_0 $\wedge$ At_1,3_0) $\wedge$ …
- $\neg$(At_1,1_1 $\wedge$ At_1,2_1) $\wedge$ $\neg$(At_1,1_1 $\wedge$ At_1,3_1) $\wedge$ …
- …

Fluents and Effect Axioms

state ∧ act ⇒ state

A *fluent* is a state variable that changes over time

How does each *state variable* or *fluent* at each time gets its value?

Fluents for PL Pacman are Pacman_x,y_t , e.g., Pacman _3,3_17

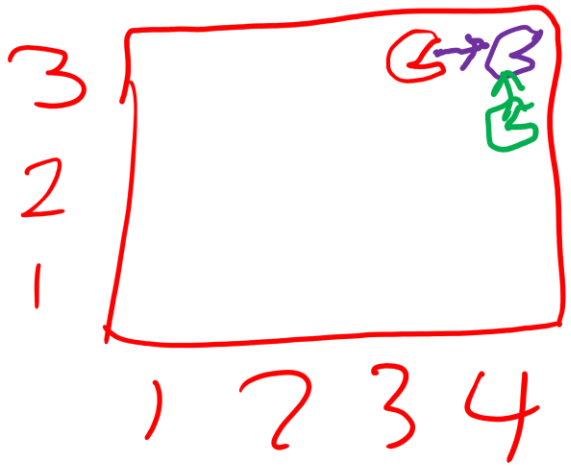

$P_{3,3,17} \wedge E_{17} \Leftrightarrow P_{4,3,18}$

Fluents and Effect Axioms

A *fluent* is a state variable that changes over time

How does each *state variable* or *fluent* at each time gets its value?

Fluents for PL Pacman are Pacman_x,y_t , e.g., Pacman _3,3_17

Fluents and Successor-state Axioms

A *fluent* is a state variable that changes over time

How does each *state variable* or *fluent* at each time gets its value?

Fluents for PL Pacman are Pacman_x,y_t , e.g., Pacman _3,3_17

A state variable gets its value according to a *successor-state axiom*

- $X_t \Leftrightarrow [X_{t-1} \wedge \neg$(some action$_{t-1}$ made it false)] v
  $[\neg X_{t-1} \wedge$ (some action$_{t-1}$ made it true)]

Fluents and Successor-state Axioms

A *fluent* is a state variable that changes over time

How does each *state variable* or *fluent* at each time gets its value?

Fluents for PL Pacman are Pacman_x,y_t , e.g., Pacman _3,3_17

A state variable gets its value according to a *successor-state axiom*

- $X_t \Leftrightarrow [X_{t-1} \wedge \neg(\text{some action}_{t-1} \text{ made it false})]$ v
  $[\neg X_{t-1} \wedge (\text{some action}_{t-1} \text{ made it true})]$

For Pacman location:

- Pacman _3,3_17 $\Leftrightarrow$ [Pacman _3,3_16 $\wedge \neg((\neg$Wall_3,4 $\wedge$ N_16) v ($\neg$Wall_4,3 $\wedge$ E_16) v …)]
  v [$\neg$ Pacman _3,3_16 $\wedge$ ((Pacman _3,2_16 $\wedge \neg$Wall_3,3 $\wedge$ N_16) v
  (Pacman _2,3_16 $\wedge \neg$Wall_3,3 $\wedge$ N_16) v …)]