AI at 50: From Programs to Solvers
Models and Techniques for General Intelligence

Héctor Geffner
ICREA & Universitat Pompeu Fabra
Barcelona, Spain

Goals

- Address some changes that have taken place in **AI research** in last 20 years
  - ▷ Changes that have to do with move from **programs** to **solvers**

$$\textit{Problem} \Longrightarrow \boxed{\textit{Solver}} \Longrightarrow \textit{Solution}$$

  - ▷ Solvers are **general programs** whose scope defined in terms of a **model**
  - ▷ Challenge is computational: how to make the solvers **scale up**
- Explain its relevance to the old AI goals concerning **general intelligence** and **human cognition**

Motivation

- It is often assumed that no much has happened in AI Research since the 80's in relation to the grand old goals of AI

Motivation

- It is often assumed that no much has happened in AI Research since the 80's in relation to the grand old goals of AI
  - ▷ Marvin Minsky, a founding father, declares *AI to be brain-dead since the 70's*, Wired Magazine, Nov 8 2003

Motivation

- It is often assumed that no much has happened in AI Research since the 80's in relation to the grand old goals of AI
  - ▷ Marvin Minsky, a founding father, declares *AI to be brain-dead since the 70's*, Wired Magazine, Nov 8 2003
  - ▷ In *Complex Cognition*, by R. Sternberg and T. Ben-Zeev, Oxford Univ. 2001, AI chapter features 5 programs: Logic Theorist (1956), GPS (1958), Eliza (1963), Mycin (1975), Dendral (1976)

Motivation

- It is often assumed that no much has happened in AI Research since the 80's in relation to the grand old goals of AI
  - ▷ Marvin Minsky, a founding father, declares *AI to be brain-dead since the 70's*, Wired Magazine, Nov 8 2003
  - ▷ In *Complex Cognition*, by R. Sternberg and T. Ben-Zeev, Oxford Univ. 2001, AI chapter features 5 programs: Logic Theorist (1956), GPS (1958), Eliza (1963), Mycin (1975), Dendral (1976)
  - ▷ Many of the debates surrounding AI (GOFAI, Situated AI, Symbolic vs. Non-Symbolic) date back to the 80's and not revised since

Motivation

- It is often assumed that no much has happened in AI Research since the 80's in relation to the grand old goals of AI
  - ▷ Marvin Minsky, a founding father, declares *AI to be brain-dead since the 70's*, Wired Magazine, Nov 8 2003
  - ▷ In *Complex Cognition*, by R. Sternberg and T. Ben-Zeev, Oxford Univ. 2001, AI chapter features 5 programs: Logic Theorist (1956), GPS (1958), Eliza (1963), Mycin (1975), Dendral (1976)
  - ▷ Many of the debates surrounding AI (GOFAI, Situated AI, Symbolic vs. Non-Symbolic) date back to the 80's and not revised since
- I aim to show that this impression is wrong

Outline

- Some AI history
- The problem of generality in AI
- Models and Solvers:
  - ▷ Graphical Models: logical and probabilistic (SAT, CSPs, Bayesian Nets)
  - ▷ Planning Models: logical and probabilistic (Strips, POMDPs, Contingent)
- Lessons learned and why they matter
- Summary

Darmouth 1956

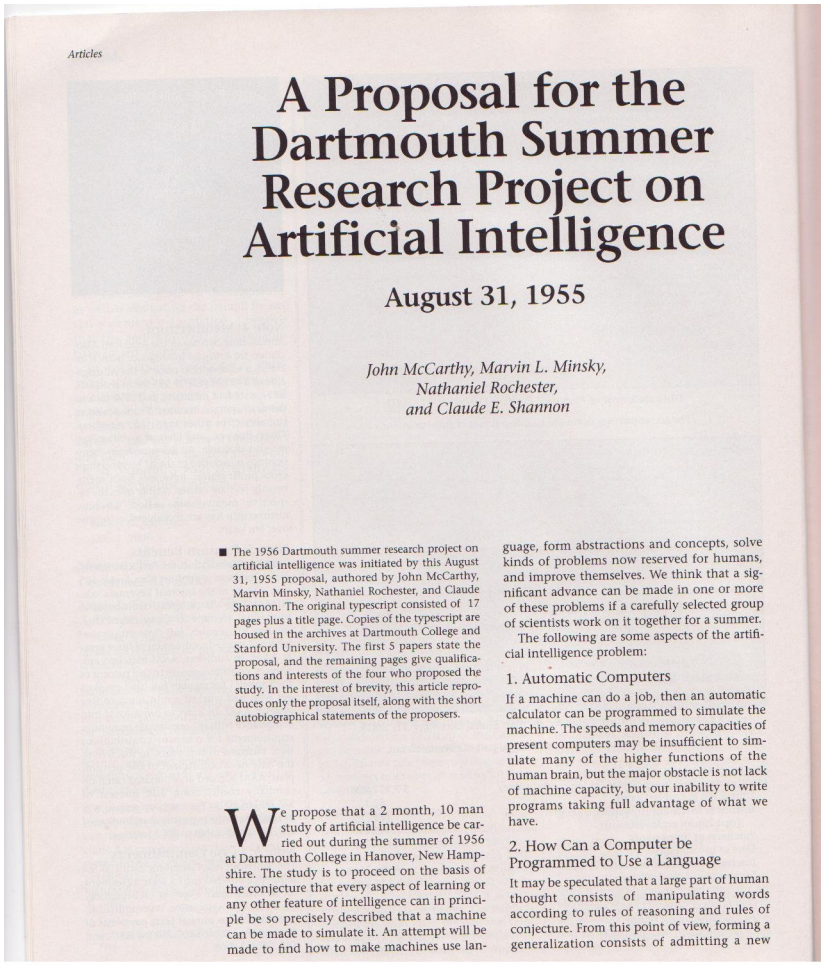

Articles

A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence

August 31, 1955

John McCarthy, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon

■ The 1956 Dartmouth summer research project on artificial intelligence was initiated by this August 31, 1955 proposal, authored by John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon. The original typescript consisted of 17 pages plus a title page. Copies of the typescript are housed in the archives at Dartmouth College and Stanford University. The first 5 papers state the proposal, and the remaining pages give qualifications and interests of the four who proposed the study. In the interest of brevity, this article reproduces only the proposal itself, along with the short autobiographical statements of the proposers.

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

The following are some aspects of the artificial intelligence problem:

1. Automatic Computers

If a machine can do a job, then an automatic calculator can be programmed to simulate the machine. The speeds and memory capacities of present computers may be insufficient to simulate many of the higher functions of the human brain, but the major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.

2. How Can a Computer be Programmed to Use a Language

It may be speculated that a large part of human thought consists of manipulating words according to rules of reasoning and rules of conjecture. From this point of view, forming a generalization consists of admitting a new

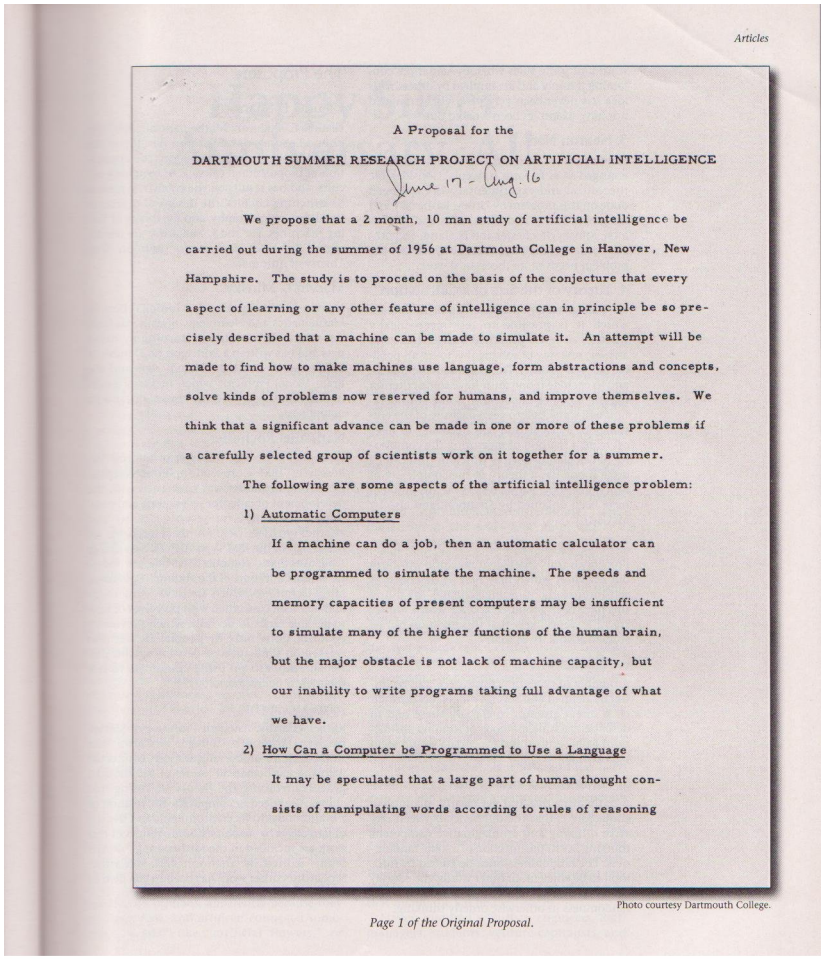

Articles

A Proposal for the

DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

June 17 - Aug. 16

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

The following are some aspects of the artificial intelligence problem:

1) Automatic Computers

If a machine can do a job, then an automatic calculator can be programmed to simulate the machine. The speeds and memory capacities of present computers may be insufficient to simulate many of the higher functions of the human brain, but the major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.

2) How Can a Computer be Programmed to Use a Language

It may be speculated that a large part of human thought consists of manipulating words according to rules of reasoning

Photo courtesy Dartmouth College.

Page 1 of the Original Proposal.

"The proposal (for the meeting) is to proceed on the basis of the conjecture that every aspect of . . . intelligence can in principle be so precisely described that a machine can be made to simulate it"

Computers and Thought 1963

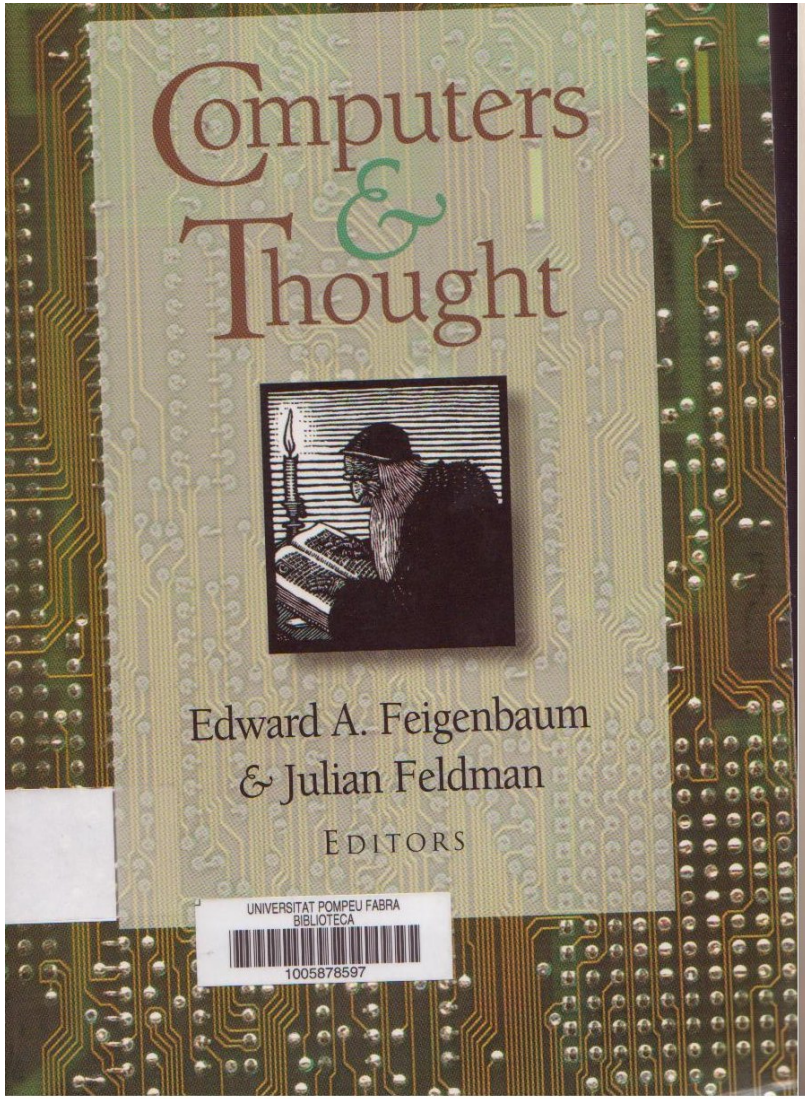

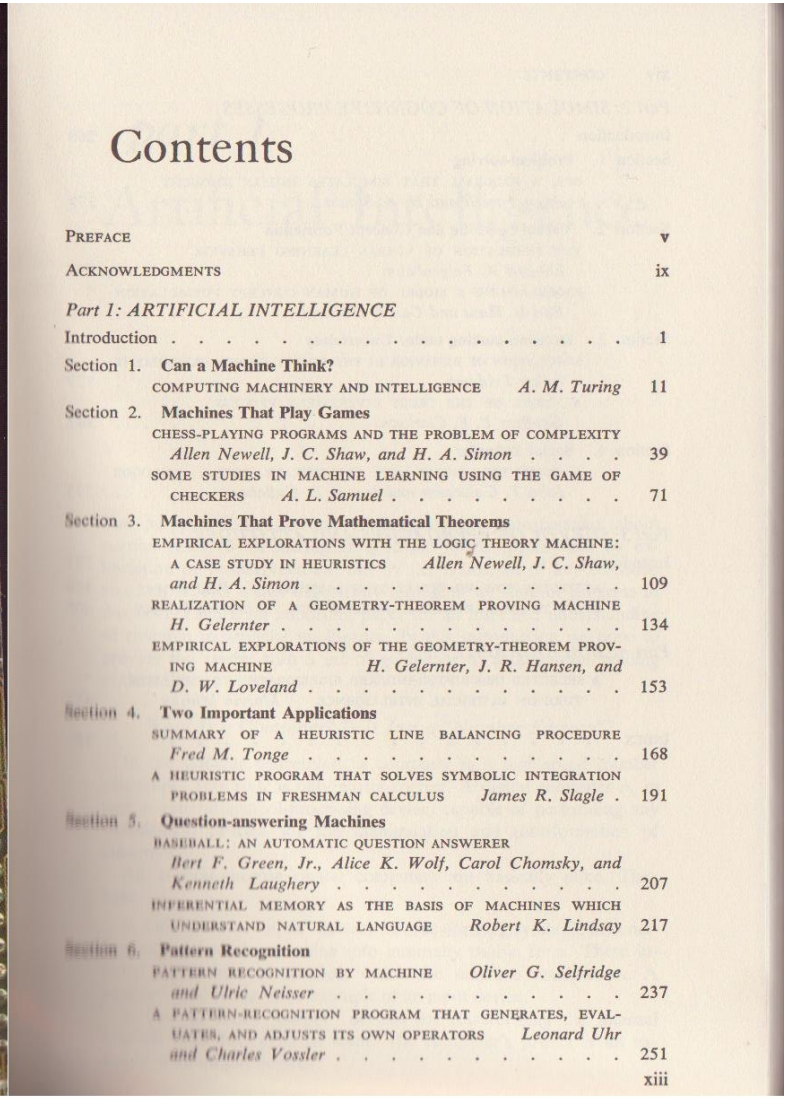

Contents

xiii

An early collection of AI papers and programs for playing chess and checkers, proving theorems in logic and geometry, planning, etc.

Importance of Programs in Early AI Work

In preface of 1963 edition of *Computers and Thought*

We have tried to focus on papers that report results. *In this collection, the papers . . . describe actual working computer programs . . . Because of the limited space, we chose to avoid the more speculative . . . pieces.*

In preface of 1995 AAAI edition

A critical selection criterion was that the paper had to describe . . . a running computer program . . . All else was talk, philosophy not science . . . (L)ittle has come out of the "talk".

AI, Programming, and AI Programming

Many of the key AI contributions in 60's, 70's, and early 80's had to to with **programming** and the **representation of knowledge** in **programs**:

- Lisp (Functional Programming)
- Prolog (Logic Programming)
- Rule-based Programming
- Interactive Programming Environments and Lisp Machines
- Frame, Scripts, Semantic Networks
- 'Expert Systems' Shells and Architectures

AI methodology: Theories as Programs

- For writing an AI dissertation in the 60's, 70's and 80's, it was common to:
  - ▷ pick up a task and domain X
  - ▷ analyze/introspect/find out how task is solved
  - ▷ capture this reasoning in a program
- The dissertation was then
  - ▷ a **theory** about X (scientific discovery, circuit analysis, computational humor, story understanding, etc), and
  - ▷ a **program** implementing the theory, **tested** over a few examples.

Many great ideas came out of this work . . . but there was a problem . . .

Methodological problem:

Theories expressed as programs cannot be proved wrong: when a program fails, it can always be blamed on 'missing knowledge'

Methodological problem:

Theories expressed as programs cannot be proved wrong: when a program fails, it can always be blamed on 'missing knowledge'

Three approaches to this problem

Methodological problem:

Theories expressed as programs cannot be proved wrong: when a program fails, it can always be blamed on 'missing knowledge'

Three approaches to this problem

- narrow the domain (expert systems)
  - ▷ problem: lack of generality

Methodological problem:

Theories expressed as programs cannot be proved wrong: when a program fails, it can always be blamed on 'missing knowledge'

Three approaches to this problem

- narrow the domain (expert systems)
  - ▷ problem: lack of generality
- accept the program is just an illustration, a demo
  - ▷ problem: limited scientific value

Methodological problem:

Theories expressed as programs cannot be proved wrong: when a program fails, it can always be blamed on 'missing knowledge'

Three approaches to this problem

- narrow the domain (expert systems)
  - ▷ problem: lack of generality
- accept the program is just an illustration, a demo
  - ▷ problem: limited scientific value
- fill up the missing knowledge (intuition, commonsense)
  - ▷ problem: not successful so far

Methodological problem:

Theories expressed as programs cannot be proved wrong: when a program fails, it can always be blamed on 'missing knowledge'

Three approaches to this problem

- narrow the domain (expert systems)
  - ▷ problem: lack of generality
- accept the program is just an illustration, a demo
  - ▷ problem: limited scientific value
- fill up the missing knowledge (intuition, commonsense)
  - ▷ problem: not successful so far

AI in the 80's

The knowledge-based approach reached an **impasse** in the 80's, a time also of debates and controversies:

- **Good Old Fashioned AI** is "rule application" but intelligence is not (Haugeland)
- **Situated AI:** representation not needed and gets in the way (Brooks)
- **Neural Networks:** inference needed is not logical but probabilistic (PDP Group)

AI in the 80's

The knowledge-based approach reached an **impasse** in the 80's, a time also of debates and controversies:

- **Good Old Fashioned AI** is "rule application" but intelligence is not (Haugeland)
- **Situated AI:** representation not needed and gets in the way (Brooks)
- **Neural Networks:** inference needed is not logical but probabilistic (PDP Group)

Many of these criticisms of mainstream AI at least partially valid then.

AI in the 80's

The knowledge-based approach reached an **impasse** in the 80's, a time also of debates and controversies:

- **Good Old Fashioned AI** is "rule application" but intelligence is not (Haugeland)
- **Situated AI:** representation not needed and gets in the way (Brooks)
- **Neural Networks:** inference needed is not logical but probabilistic (PDP Group)

Many of these criticisms of mainstream AI at least partially valid then.

How valid are they now?

AI Research in 2007

Recent issues of AIJ, JAIR, AAAI or IJCAI shows papers on:

1. **SAT and Constraints**
2. **Search and Planning**
3. **Probabilistic Reasoning**
4. **Probabilistic Planning**
5. Inference in First-Order Logic
6. Machine Learning
7. Natural Language
8. Vision and Robotics
9. Multi-Agent Systems

I'll focus on 1–4: these areas often deemed about **techniques**, but more accurate to regard them as **models** and **solvers**.

Example: Solver for Linear Equations

$$\textit{Problem} \Longrightarrow \boxed{\textit{Solver}} \Longrightarrow \textit{Solution}$$

Example: Solver for Linear Equations

$$\textit{Problem} \Longrightarrow \boxed{\textit{Solver}} \Longrightarrow \textit{Solution}$$

- **Problem:** The age of John is 3 times the age of Peter. In 10 years, it will be only 2 times. How old are John and Peter?

Example: Solver for Linear Equations

$$\textit{Problem} \Longrightarrow \boxed{\textit{Solver}} \Longrightarrow \textit{Solution}$$

- **Problem:** The age of John is 3 times the age of Peter. In 10 years, it will be only 2 times. How old are John and Peter?
- **Expressed as:** $J = 3P \;\; ; \;\; J + 10 = 2(P + 10)$
- **Solver:** Gauss-Jordan (Variable Elimination)
- **Solution:** $P = 10 \;\; ; \;\; J = 30$

Example: Solver for Linear Equations

$$\textit{Problem} \Longrightarrow \boxed{\textit{Solver}} \Longrightarrow \textit{Solution}$$

- **Problem:** The age of John is 3 times the age of Peter. In 10 years, it will be only 2 times. How old are John and Peter?
- **Expressed as:** $J = 3P \;\; ; \;\; J + 10 = 2(P + 10)$
- **Solver:** Gauss-Jordan (Variable Elimination)
- **Solution:** $P = 10 \;\; ; \;\; J = 30$

Solver is **general** as deals with any problem expressed as an instance of **model**

Linear Equations Model, however, is **tractable**, AI models are not . . .

AI Solvers

$$\textit{Problem} \Longrightarrow \boxed{\textit{Solver}} \Longrightarrow \textit{Solution}$$

- The basic models and task we will consider are
  - ▷ **Constraint Satisfaction/SAT**: find state that satisfies constraints
  - ▷ **Bayesian Networks:** find probability over variable given observations
  - ▷ **Planning Problems:** find actions that map given state into a final state
  - ▷ **Planning with Feedback:** find strategy for mapping state into final state

- All of these models are **intractable**, and some extremely powerful (POMDPs)
- The challenge is computational: **how to scale up**
- For this, solvers must **recognize and exploit structure** of the problems
- Methodology is **empirical**: benchmarks and competitions
- Significant progress in recent years

SAT and CSPs

- **SAT:** determine if there is a **truth assignment** that satisfies a set of clauses

$$x \vee \neg y \vee z \vee \neg w \vee \cdots$$

- Problem is NP-Complete, which in practice means worst-case behavior of SAT algorithms is **exponential** in number of variables
- Yet current SAT solvers manage to solve problems with **thousands of variables and clauses**, and used widely (circuit design, verification, planning, etc)
- Key is **efficient (poly-time) inference** in every node of search tree: **unit resolution**, **conflict-based learning**, . . .
- Many other ideas **logically possible**, but **do not work** (don't scale up).
- Same for **Constraint Satisfaction Problems (CSPs)**

Related Tasks: From SAT to Bayesian Networks

- **Weighted MAX-SAT**: find assignment σ that minimizes total cost $w(C)$ of violated clauses

$$\sum_{C:\sigma \not\models C} w(C)$$

- **Weighted Model Counting**: Adds up 'weights' of satisfying assignments:

$$\sum_{\sigma:\sigma \models T} \prod_{L \in \sigma} w(L)$$

SAT methods extended to these other tasks, closely connected to **probabilistic** reasoning tasks over **Bayesian Networks**:

- **Most Probable Explanation (MPE)** easily cast as Weighted MAX-SAT
- **Probability Assessment** $P(X|Obs)$ easily cast as Weighted Model Counting

Current best BN solvers built over this formulation (ACE, Weighted Cachet)

SAT, CSPs, and BNs: Complexity and Treewidth

- The underlying structure of SAT, CSPs, and BNets can be expressed by graph G

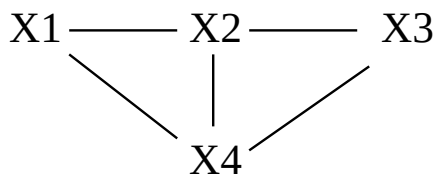

- A parameter called the **(induced) treewidth** $w(G)$ measures then how 'close' is G to a Tree, $w(G) = 2$ for G above, and $w(G) = 1$ if G is a tree.

- All SAT, CSP, and BN tasks are **exponential** in $w(G)$, and thus solvable in **polynomial time** for **bounded** $w(G)$ (e.g., trees)

- These models often referred to as **graphical models** (Dechter 03)

Action Selection in Planning Models

- **(Classical) Planning** concerned with finding a **sequence of actions** that transforms an **initial state** into a **goal state**. This is called a **plan**
- **States** are truth assignments as before, represented by the atoms that are true
- Actions **add** certain atoms and **delete** others, provided their **preconditions** hold
- A **planner** is a solver that takes a **planning problem** (initial and goal states, and actions) and outputs a **plan**
- The **cost** of a plan given by the **number of actions**

$$\textit{Init, Actions, Goals} \implies \boxed{\textit{Planner}} \implies \textit{Plan}$$

Example

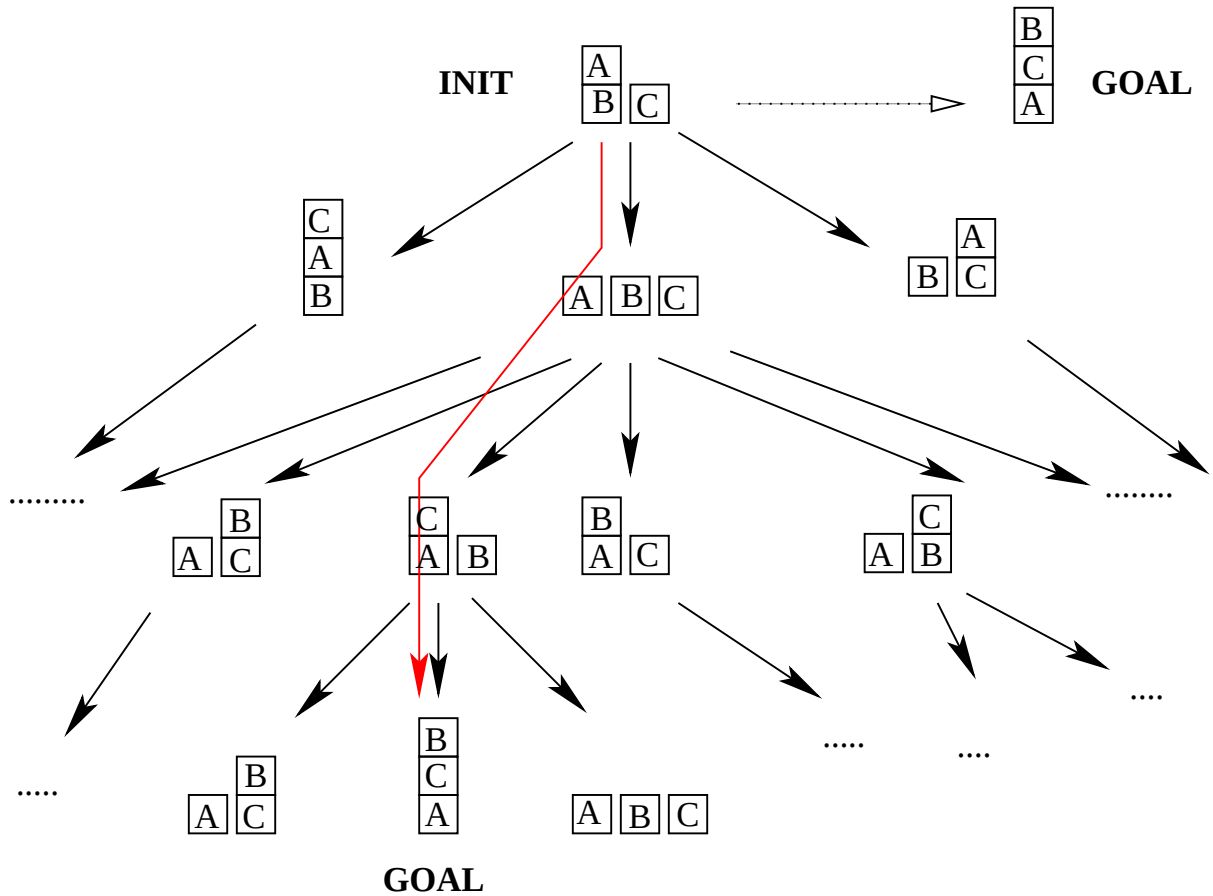

- Given the **actions** that move a 'clear' block to the table or onto a another 'clear' block, **find a plan** to achieve the goal
- Problem becomes **finding a path** in a **directed graph**

How planning problems are solved?

- How do we find a route in a map from Barcelona to Madrid?
- Need **sense of direction:** whether an action takes us towards the goal or not
- In AI, this is captured by **heuristic functions**: functions $h(s)$ that provide an **estimate of the cost** (number of actions) from any state s to the goal

How planning problems are solved?

- How do we find a route in a map from Barcelona to Madrid?
- Need **sense of direction:** whether an action takes us towards the goal or not
- In AI, this is captured by **heuristic functions**: functions $h(s)$ that provide an **estimate of the cost** (number of actions) from any state s to the goal
- Key new idea in planning is that useful heuristics $h(s)$ can be obtained **automatically** from the problem encoding
- **How?** Solving a **relaxed problem** where **deletes** are dropped
- Heuristic $h(s)$ is cost of solution found for **relaxed problem** in **poly-time**

How is our problem solved?

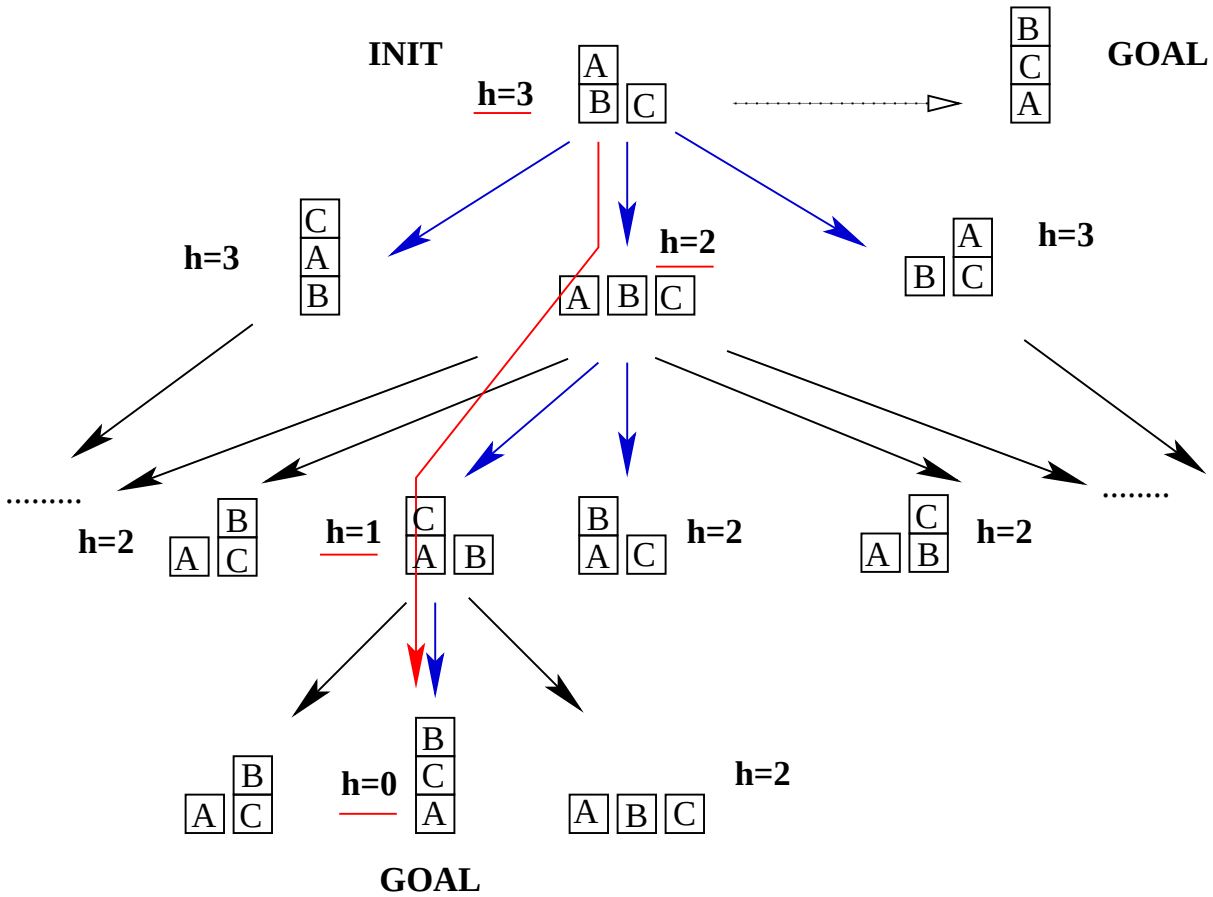

- Provided with the **heuristic** h, plan found without search by **hill-climbing**
- Actually, only the states reached by the actions in blue evaluated in FF

The appraisals $h(s)$ from a cognitive point of view

- they are **opaque** and thus cannot be **conscious**

  meaning of symbols in the relaxation is not the normal meaning; e.g., objects can be at many places at the same time as old locations not deleted

- they are **fast and frugal** (linear-time), but unlike the 'fast and frugal heuristics' of Gigerenzer et al. are **general**

  they apply to all problems fitting the model (planning problems)

- they play the role of **'gut feelings'** or **'emotions'** according to De Sousa 87, Damasio 94, Evans 2002, Gigerenzer 2007

  providing a guide to action while avoiding infinite regresses in the decision process

Status of Classical Planning

- The good news: **Classical planning works!**
  - ▷ *Large problems solved very fast (non-optimally)*
- **Model simple but useful**
  - ▷ *Operators not primitive; can be policies themselves*
  - ▷ *Fast closed-loop replanning able to cope with uncertainty sometimes*
- Not so good; **limitations:**
  - ▷ *Does not model* **Uncertainty** *(no probabilities)*
  - ▷ *Does not deal with* **Incomplete Information** *(no sensing)*
  - ▷ . . .

Beyond Classical Planning: Two Strategies

1. **Top-down:** Develop solver for **more general class of models;** e.g., Markov Decision Processes (MDPs), Partial Observable MDPs (POMDPs), . . .

   +: generality
   −: complexity

2. **Bottom-up:** Extend the scope of **current 'classical' solvers**

   +: efficiency
   −: generality

The Trouble with Incomplete State Information

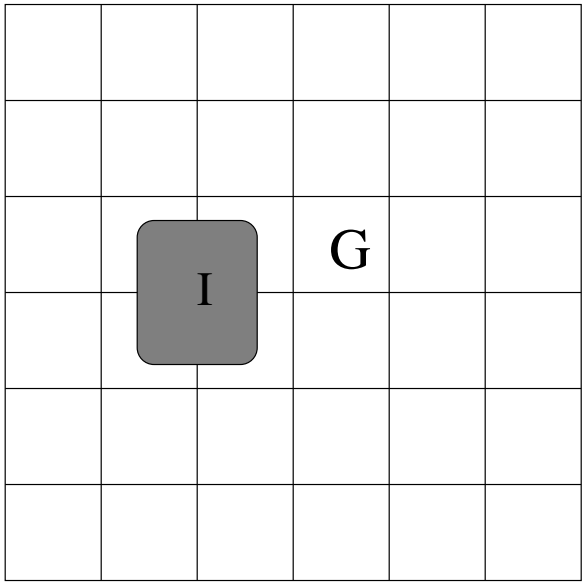

Problem: A robot must move from an **uncertain** I into G with **certainty**, one cell at a time, in a grid nxn. No sensing is available.

- Conformant and classical planning look similar except for uncertain I
- Plans, however, may be quite different: best **conformant plan** above **must move robot to a corner first!** (to localize)

Wumpus World PEAS description

Performance measure

gold +1000, death -1000

-1 per step, -10 for using the arrow

Environment

Squares adjacent to wumpus are smelly

Squares adjacent to pit are breezy

Glitter iff gold is in the same square

Shooting kills wumpus if you are facing it

Shooting uses up the only arrow

Grabbing picks up gold if in same square

Releasing drops the gold in same square

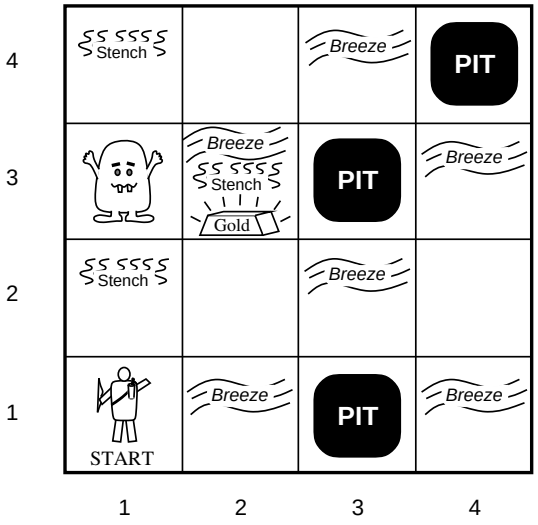

Actuators Left turn, Right turn,
Forward, Grab, Release, Shoot

Sensors Breeze, Glitter, Smell

How to solve planning problems with incomplete information?

- Action selection problem in **classical planning** (full state info, deterministic actions) solved by automatic extraction of informative **heuristic functions** $h(s)$ that estimate cost to goal

- Action selection problem P in **planning with sensing** pursued in two ways:

  - ▷ *extract informative* **heuristics** $h(b)$ *automatically over* **beliefs** b
  - ▷ *convert automatically into problem* $K(P)$ *with no incomplete information*

- Both approaches **optimal** for optimal heuristic $h(b)$ and complete translation $K(P)$; although fast approximations used instead

- A number of planning **competitions** exist to evaluate ideas and push the state of the art . . .

Summary

- A **research agenda** that has emerged in last 20 years: **solvers** for a range of **intractable models**

- **Solvers** unlike other programs are **general** as they do not target individual problems but families of problems (**models**)

- The challenge is **computational** and the methodology **empirical**

- Consistent **progress:**
  - ▷ efficient but effective inference methods (derivation of h, conflict-learning)
  - ▷ islands of tractability (treewidth methods and relaxations)
  - ▷ transformations (compiling away incomplete info, extended goals, . . .)

- While the agenda is technical, resulting ideas likely to be **relevant** for understanding general intelligence and human cognition

General Solvers and People

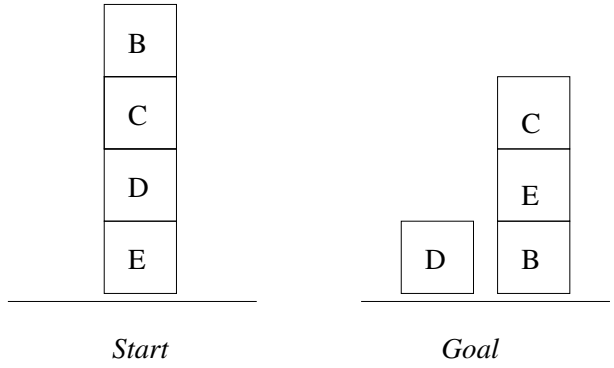

How people solve this problem?

- Blocks looks too easy; psychologists preferred puzzles like Tower of Hanoi. Yet
  - ▷ Blocks is **not trivial for a general solver**, and indeed
  - ▷ Language and Perception easy for people but not **computationally**
- Are these problems solved using 'domain-knowledge'?
  - ▷ **not clear:** easy to introduce variations where knowledge does not apply
  - ▷ **not necessary:** can be solved provided 'right inferences' are captured

Bayesian Networks

A Bayesian Network (Pearl 1988) is a compact representation of a joint probability distribution over a set of variables $X_1, \ldots, X_n$ made up of:

- a DAG where the nodes are the variables $X_1, \ldots, X_n$
- Conditional probability tables $prob(X_i|pa(X_i))$, $i = 1, \ldots, n$, where $pa(X_i)$ refers to the parents of X_i in the DAG

The DAG implicitly defines a set of **independences** that result in joint distrib.

$$P(X_1, \ldots, X_n) = \prod_{i=1,n} P(X_i|pa(X_i))$$

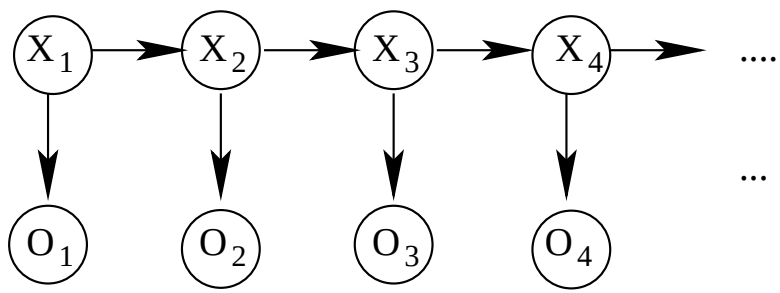

A **Hidden Markov Model** is a **Bayesian Tree** solvable in **linear-time.**

Executing Policies vs Finding the Policies

- **Solutions** to planning problems (Strips, MDPs, POMDPs,. . .) can be **expressed** as **stimulus-action rules** (and many other ways)

- Yet, **expressing solutions** (to Strips, MDPs, . . . , ..) is different than **finding solutions**

- In **Situated-AI** (Brooks 90), it is assumed that the **programmer solves the problem** in his/her head, and the robot **just executes this solution**

- In **Model-based AI**, solutions computed by solver.

Executing Policies vs Finding the Policies

- **Solutions** to planning problems (Strips, MDPs, POMDPs,. . .) can be **expressed** as **stimulus-action rules** (and many other ways)

- Yet, **expressing solutions** (to Strips, MDPs, . . . , ..) is different than **finding solutions**

- In **Situated-AI** (Brooks 90), it is assumed that the **programmer solves the problem** in his/her head, and the robot **just executes this solution**

- In **Model-based AI**, solutions computed by solver.

  Is this necessary?

Executing Policies vs Finding the Policies

- **Solutions** to planning problems (Strips, MDPs, POMDPs,. . .) can be **expressed** as **stimulus-action rules** (and many other ways)

- Yet, **expressing solutions** (to Strips, MDPs, . . . , ..) is different than **finding solutions**

- In **Situated-AI** (Brooks 90), it is assumed that the **programmer solves the problem** in his/her head, and the robot **just executes this solution**

- In **Model-based AI**, solutions computed by solver.

  Is this necessary?

  - ▷ in lower animals, solutions **hardwired** by evolution
  - ▷ in humans, evolution and culture have produced **solution methods** as well

- **Learning** is the third approach to get solutions, different than **Model-based Solver** or **Programming/Evolution**