

HUMAN-LEVEL ARTIFICIAL INTELLIGENCE & COGNITIVE SCIENCE

Nils J. Nilsson
Stanford AI Lab

<http://ai.stanford.edu/~nilsson>

Symbolic Systems 100, April 15, 2008

OUTLINE

Computation and Intelligence

Approaches Toward HLA

The Current Situation

THE COMPUTER: A UNIVERSAL MACHINE

“It can be shown that a single special machine of that type [a Turing machine] can be made to do the work of all. It could in fact be made to work as a model of any other machine. The special machine may be called the universal machine.”

—Alan Turing

“The importance of the universal machine is clear. We do not need to have an infinity of different machines doing different jobs. **A single one will suffice**. The engineering problem of producing various machines for various jobs is replaced by the office work of ‘programming’ the universal machine to do these jobs.”

—Alan Turing

“[Turing] decided the scope of the computable encompassed far more than could be captured by explicit instruction notes, and quite enough to include **all that human brains did, however creative or original**.”

—Andrew Hodges, a Turing Biographer

THE PHYSICAL SYMBOL SYSTEM HYPOTHESIS (PSSH)

“A physical symbol system [i.e., a computer] has the **necessary** and **sufficient** means for intelligent action.”

—Allen Newell and Herbert Simon

Relevance to Cognitive Science:

Computational processes can explain (be a theory of) human intelligence.

Relevance to AI:

Computational processes can implement human intelligence.

But, it's just a hypothesis!

THE NAYSAYERS CLAIM (Among Other Things) “THE BRAIN IS NOT A COMPUTER!”

Computation:

mainly serial

10^9 ops/sec

10^9 transistors

digital/discrete (even binary!)

disembodied

silicon

subject to crashes

...

The Brain:

highly parallel

10^3 ops/sec

10^{14} neurons; 10^{17} synapses

analog/continuous

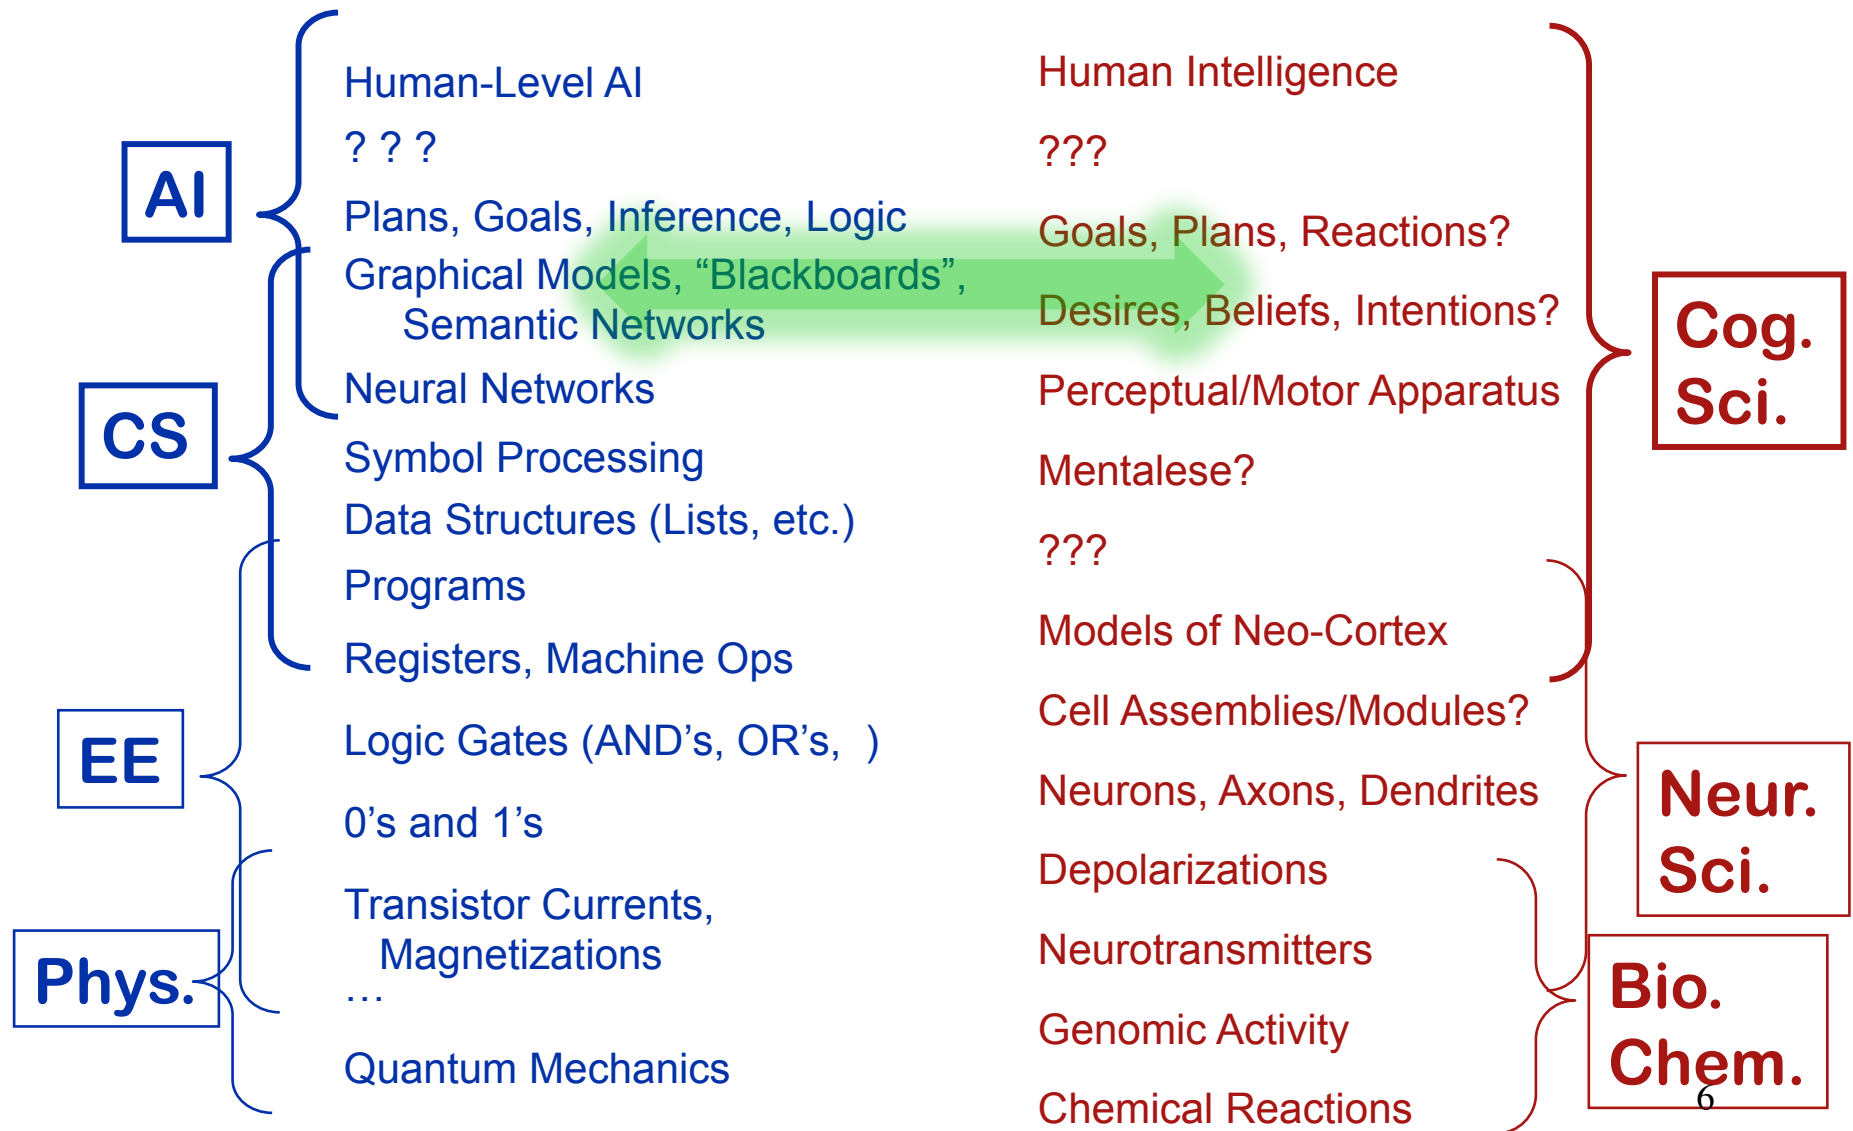
embodied

protein

fault-tolerant

...

WHAT IS A COMPUTER? WHAT IS A BRAIN?



CAN COMPUTATIONAL SYSTEMS
BE INTELLIGENT?

HOW WOULD WE KNOW?

The Turing Test

The “Employment” Test

SOME JOBS THAT HUMANS PERFORM*




Meeting and Convention Planner
Maid and Housekeeping Cleaner
Receptionist
Financial Examiner
Computer Programmer
Roofer's Helper
Library Assistant
Procurement and Sales Engineer
Farm, Greenhouse, Nursery Worker
Dishwasher
Home Health Aide

Small Engine Repairer
Paralegal
Lodging Manager
Proofreader
Tour Guide and Escort
Geographer
Engine and Other Machine Assembler
Security Guard
Retail Salesperson
Marriage and Family Counselor
Hand Packer and Packager

CAN THEY BE AUTOMATED?

*From "America's Job Bank," a list of more than 1,500 jobs.

Available at www.jobsearch.org/help/employer/SSONetJobCodeListbyCategory2.html



HOW TO PROCEED?

**LET'S LOOK AT SOME
THINGS AI HAS TRIED**

THINGS AI HAS TRIED

- Try to program some activities thought to require intelligence
- Try to program some fundamental processes thought to be involved in intelligence
- Try to imitate the brain
- Try to simulate the performance of ever more complex biological organisms
- Try to simulate biological evolution
- Try to “educate” simple (child-like) programs to make them more intelligent and capable

1. Programming Activities That Require Intelligence

Game playing

Theorem proving

Pattern recognition (images, speech, . . .)

Natural language understanding


Autonomous activity

(robots that can perceive, plan and act)

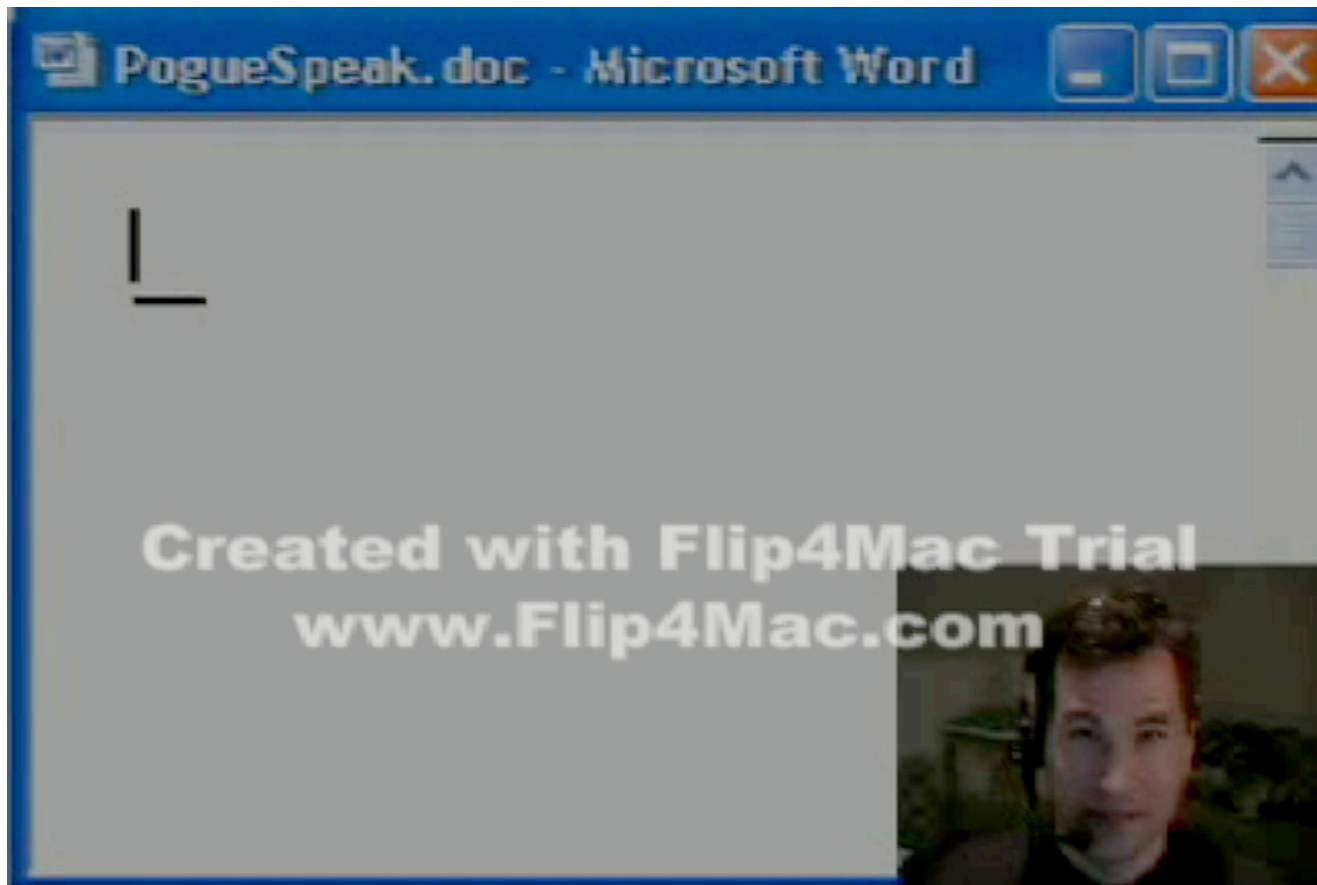
Expert judgment

. . .

Game-Playing



Speech Recognition




Autonomous Activity



Autonomous Activity



Expert Judgment



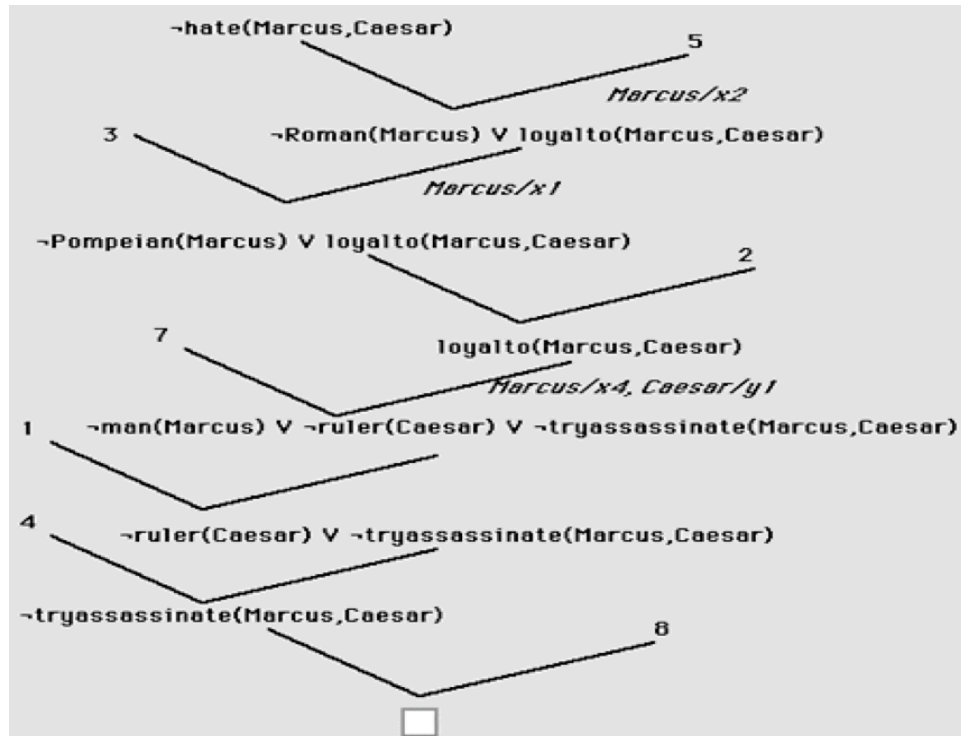

Part of a Large Bayesian Network
Used for the Diagnosis of Hepatobiliary Diseases

WOULD SOME COMBINATION
OF ALL OF THESE SKILLS
ADD UP TO HLAI?

2. Programming Processes Thought to be Involved in Intelligence

Logical reasoning
Probabilistic reasoning
Search
Image processing
Knowledge representation
Learning
Syntactic analysis
Planning
...

Logical Reasoning

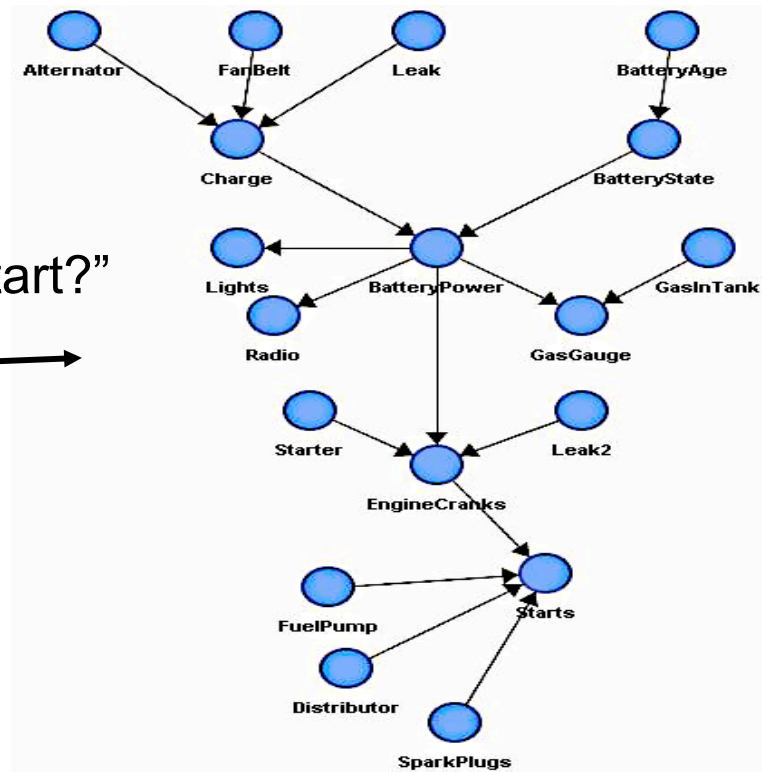


Resolution Theorem Proving
Propositional Satisfiability (SAT)

Probabilistic Reasoning


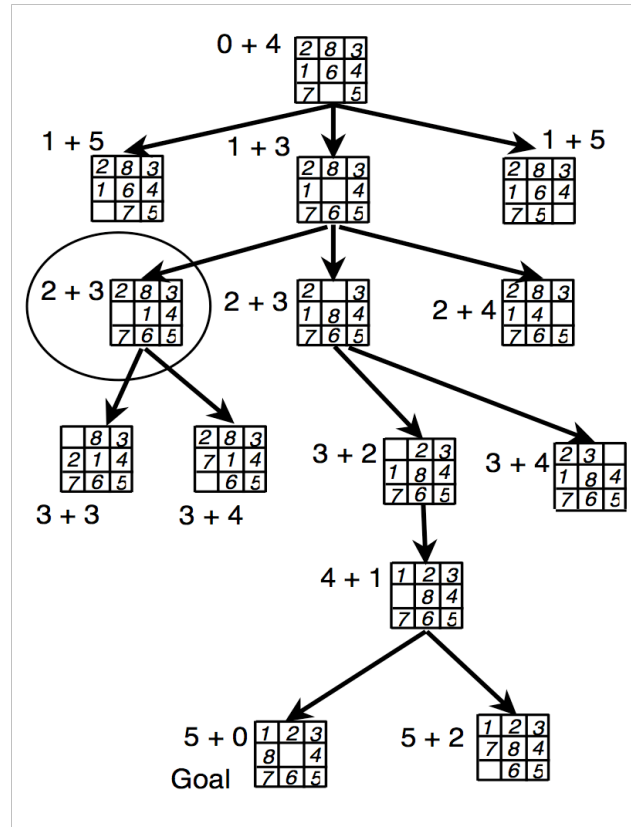


“Why won’t the car start?”




Bayesian Belief Networks

Search



**A*, Hill-Climbing
Recursive Back-Tracking**

Image Processing




E. g., Edge Extraction

Representing Knowledge



Semantic Networks, Cyc, WordNet


Learning



E. g., Reinforcement Learning

WOULD SOME COMBINATION
OF PROCESSES LIKE THESE
PRODUCE HLAI?

3. Trying to Imitate the Brain



Neural Networks


Models of the Neo-Cortex

Neural Networks

Input: Text Versions of English Words

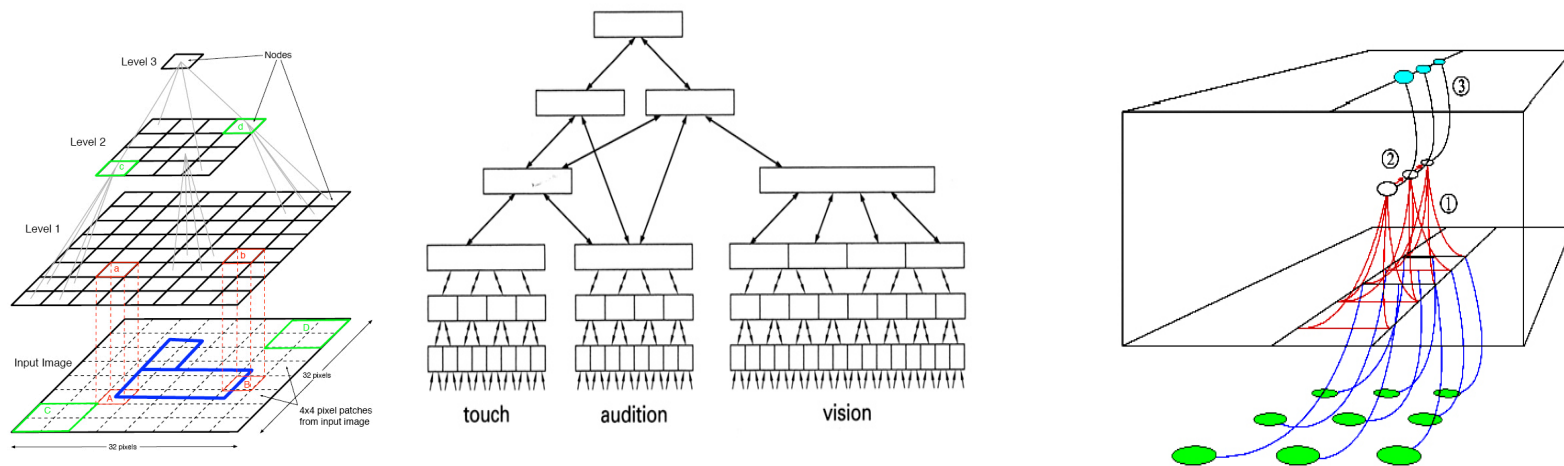
Output: Sound

Training: Change Weights to Make Sound More Correct



Sejnowski, T. J. and Rosenberg, C. R., Parallel networks that learn to pronounce English text, *Complex Systems* 1, 145-168 (1987).

Models of Neo-Cortex



Large Graphical Models


Jeff Hawkins, Tom Dean, David Mumford, Geoff Hinton, . . .

4. Simulating Simple Biological Organisms



Rod Brooks's Creepy-Crawly Things

5. Simulating Biological Evolution



Target-Seeking Demo

<http://www.cs.northwestern.edu/~fjs750/netlogo/final/gpdemo.html>

Truck-Backing Demo

<http://www.handshake.de/user/blickle/Truck/index.html>

Genetic Algorithms, Genetic Programming

6. “Educating” Educable Programs

Turing’s “Child Programme”

Cassimatis’s “Cognitive Substrate”

Lenat’s “CYC Bootstrapping”

WHICH OF THESE APPROACHES
WILL BE SUCCESSFUL?

WE'LL HAVE TO WAIT AND SEE!

IN THE MEANTIME:

**THESE EFFORTS
HAVE PRODUCED A
GROWING
ARMAMENTARIUM OF
TECHNICAL TOOLS**

Here is a Partial List:

Bayesian Belief Networks
Hidden Markov Models
Kalman Filtering
POMDP's
A* Global Search
Hill-Climbing Local Search
GA/GP
Resolution Theorem Prvg.
SAT Encodings/Solvers
Semantic Networks
Reinforcement Learning

Neural Networks
Backpropagation
Support Vector Machines
Blackboard Architectures
Monte Carlo Methods
Statistical Grammars
Expectation Maximization
Inductive Logic Programming
Teleo-Reactive Programs
Particle Filtering
Model-Based Vision

**Will They Help Us Achieve HLAI?
Are More Tools Needed?**

THE CURRENT SITUATION

The Tools Are Being Used to Solve Problems in Several Fields:

Biology
Genomics
Chemistry
Medicine
Aeronautics
Geology
Data Mining
Business
...

But Little Work is Being Done Toward HLAI. Why?

PETER HART'S QUESTION (AI's Progress)

