



# Artificial Intelligence in the Open World

Eric Horvitz

Presidential Address  
July 2008



# Rich Intellectual History of AI



18<sup>th</sup>

- Julien Offray de La Mettrie  
*L'homme machine* (1747)

19<sup>th</sup>

- Charles Babbage  
*Difference Machine, Analytical Engine*


20<sup>th</sup>

- Turing, von Neumann, Weiner, *et al.*

- Newell, Simon, Rochester, McCarthy, *et al.*

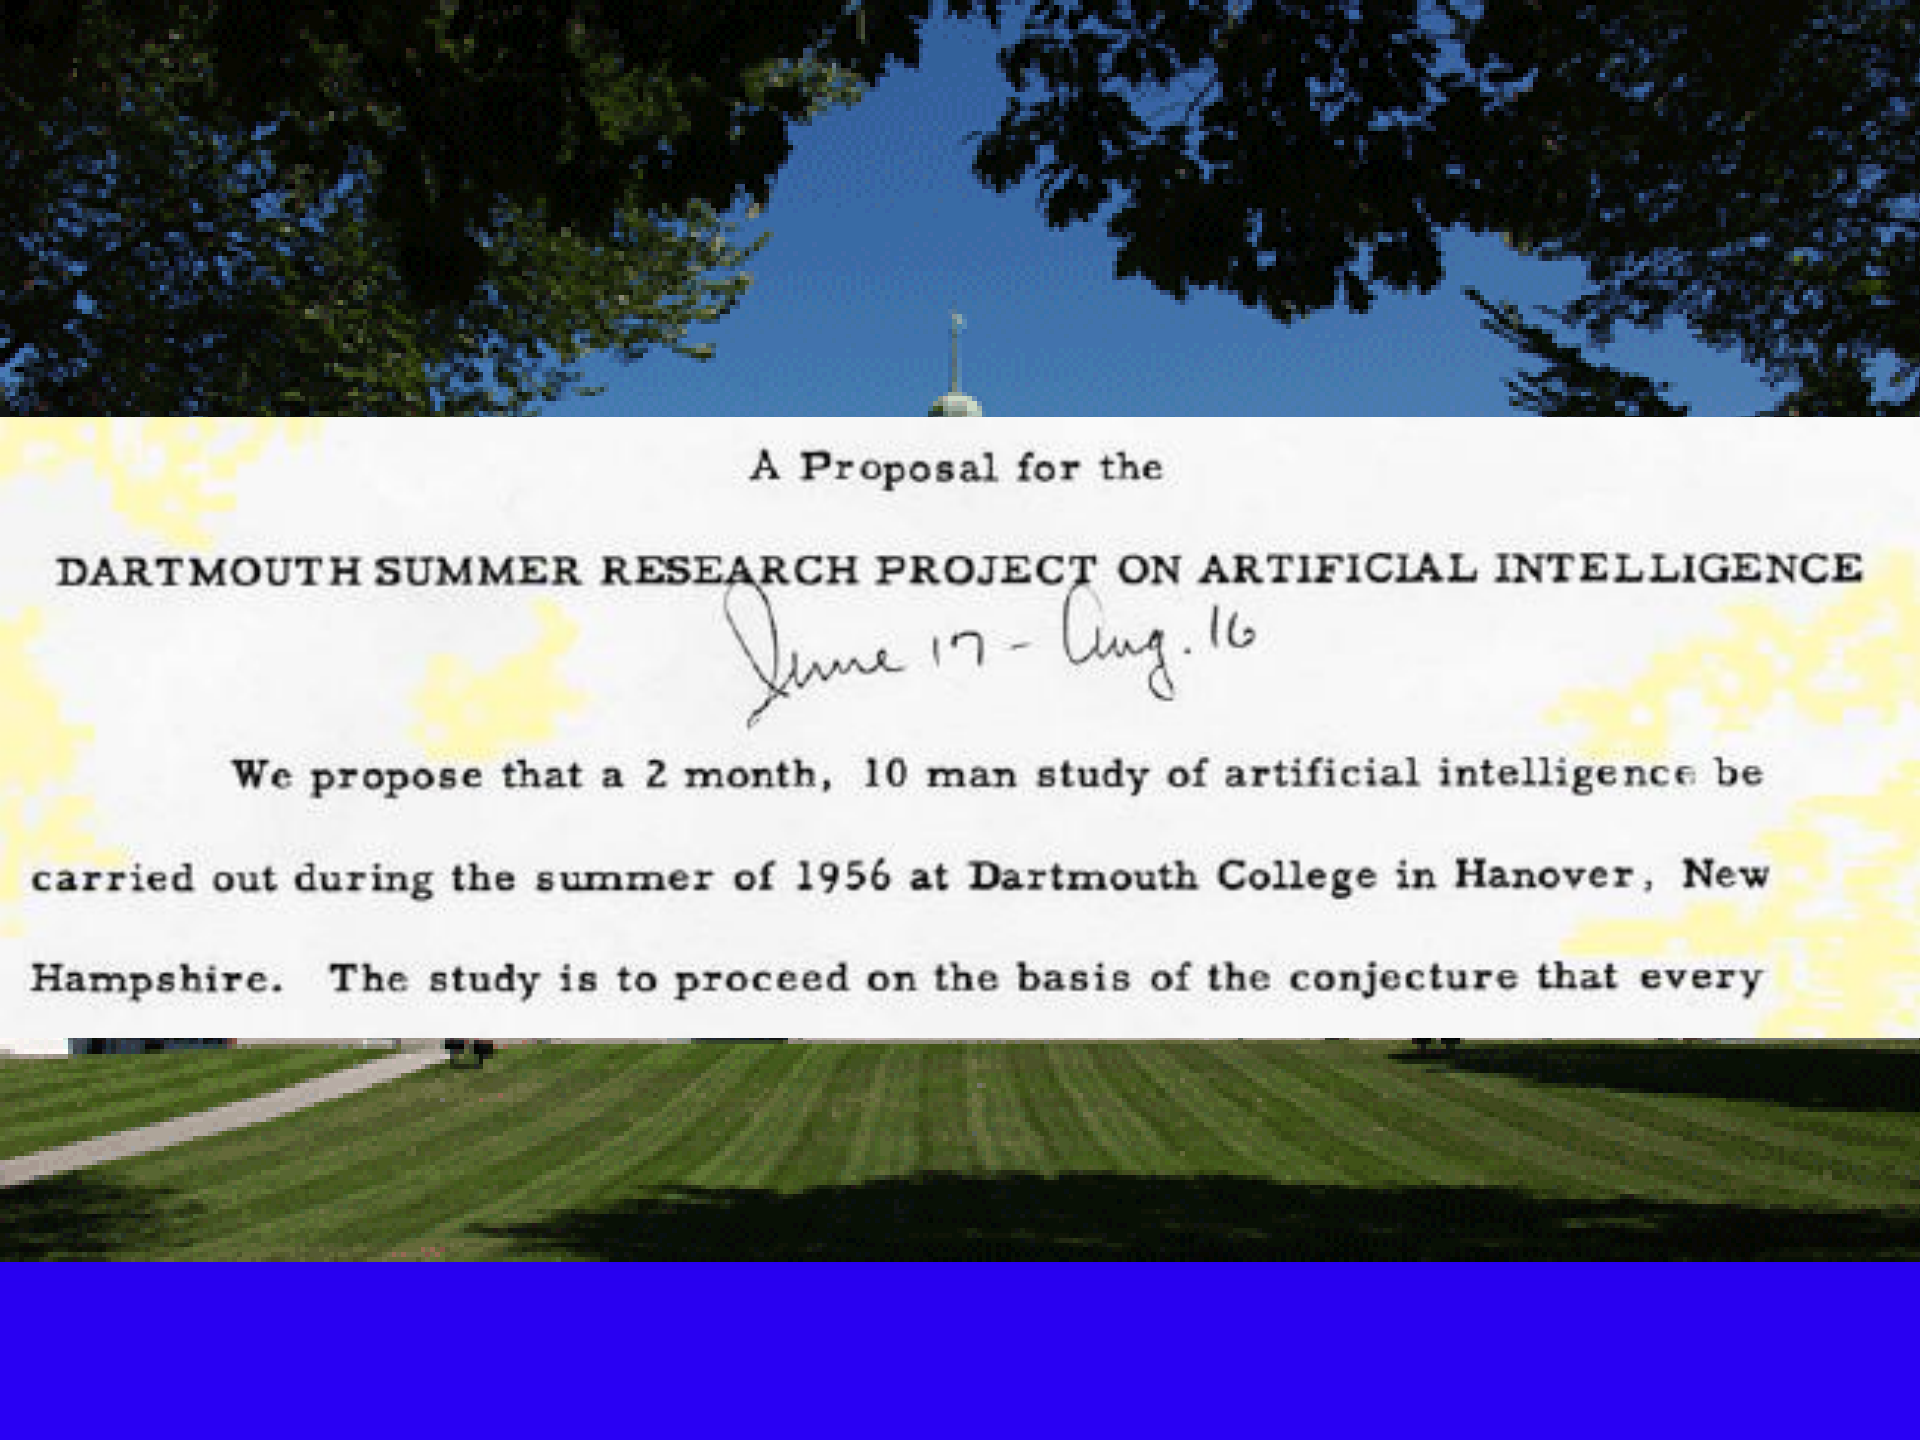
# Computation as Basis of Thought & Intelligent Behavior

- ◆ Theories of computability
- ◆ General purpose computer









A Proposal for the

DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

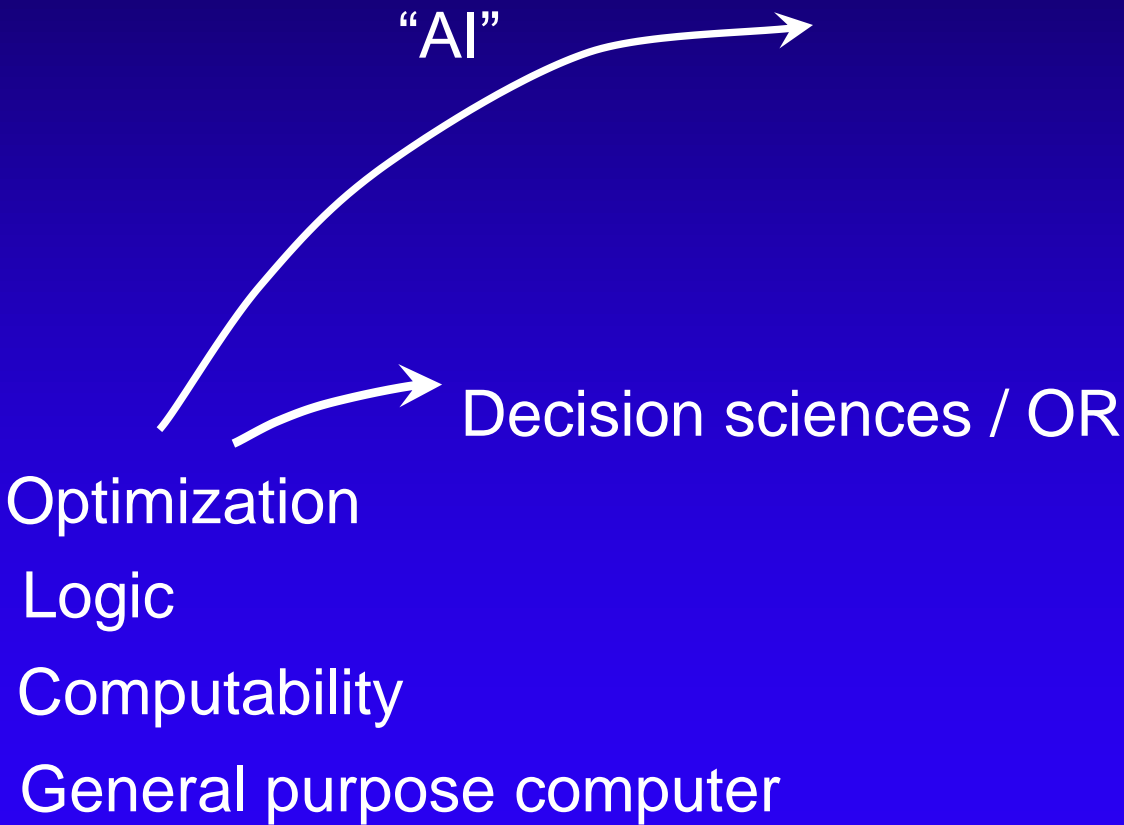
*June 17 - Aug. 16*

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

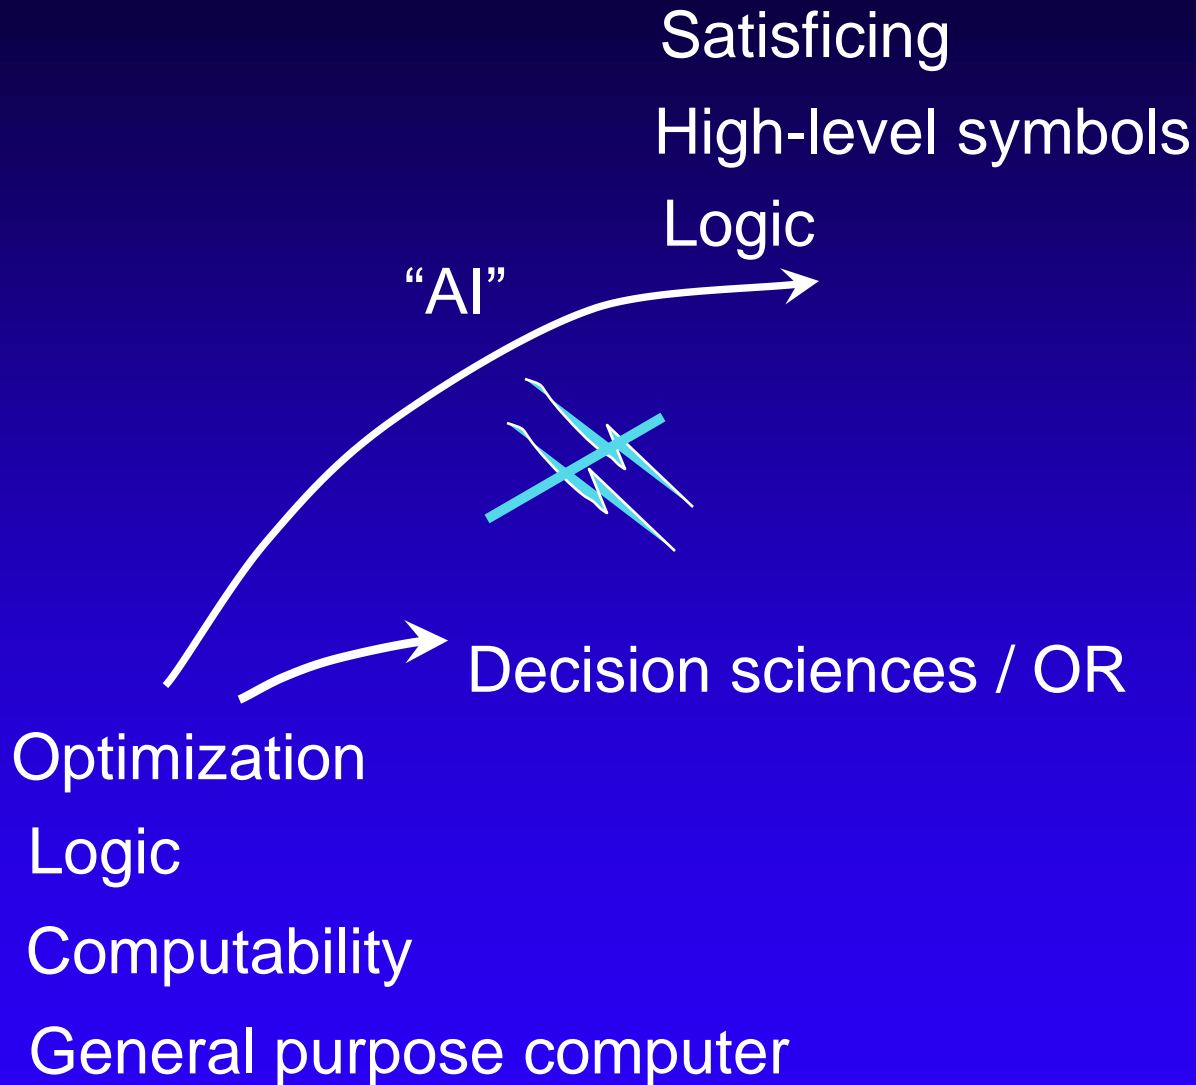
# Dartmouth Meeting

- ◆ Machine methods of forming abstractions from sensory and other data
- ◆ Carrying out activities which may best be described as self-improvement
- ◆ Manipulating words according to rules of reasoning and rules of conjecture
- ◆ Developing a theory of the complexity for various aspects of intelligence

# Paradigms of Computational Intelligence



# Paradigms of Computational Intelligence





# Grappling with Incompleteness in an Open World



# Grappling with Incompleteness in an Open World



# Grappling with Incompleteness in an Open World

*Bounded rationality (Simon, et al).*


Hopelessly incomplete knowledge of ...

- Preferences
- State of world
- Outcomes of action

# Enduring Perspective: Intelligence amidst Inescapable Incompleteness

Limited agents immersed in complex universes

- Limited representations
- Limited time and memory





# Stepping into the Open World

- ◆ **Key technical challenges**
- ◆ **AI moving into the world**
- ◆ **AI research community**

# Stepping into the Open World

- ◆ **Key technical challenges**
- ◆ **AI moving into the world**
- ◆ **AI research community**

# Stepping into the Open World

## THE ROBOT'S DILEMMA

The Frame Problem  
In Artificial Intelligence

# Stepping into the Open World





# Relevance & Attention in Open Worlds

## ■ Frame problem

*How to limit the scope of the reasoning required to derive the consequences of an action?*

## ■ Qualification problem

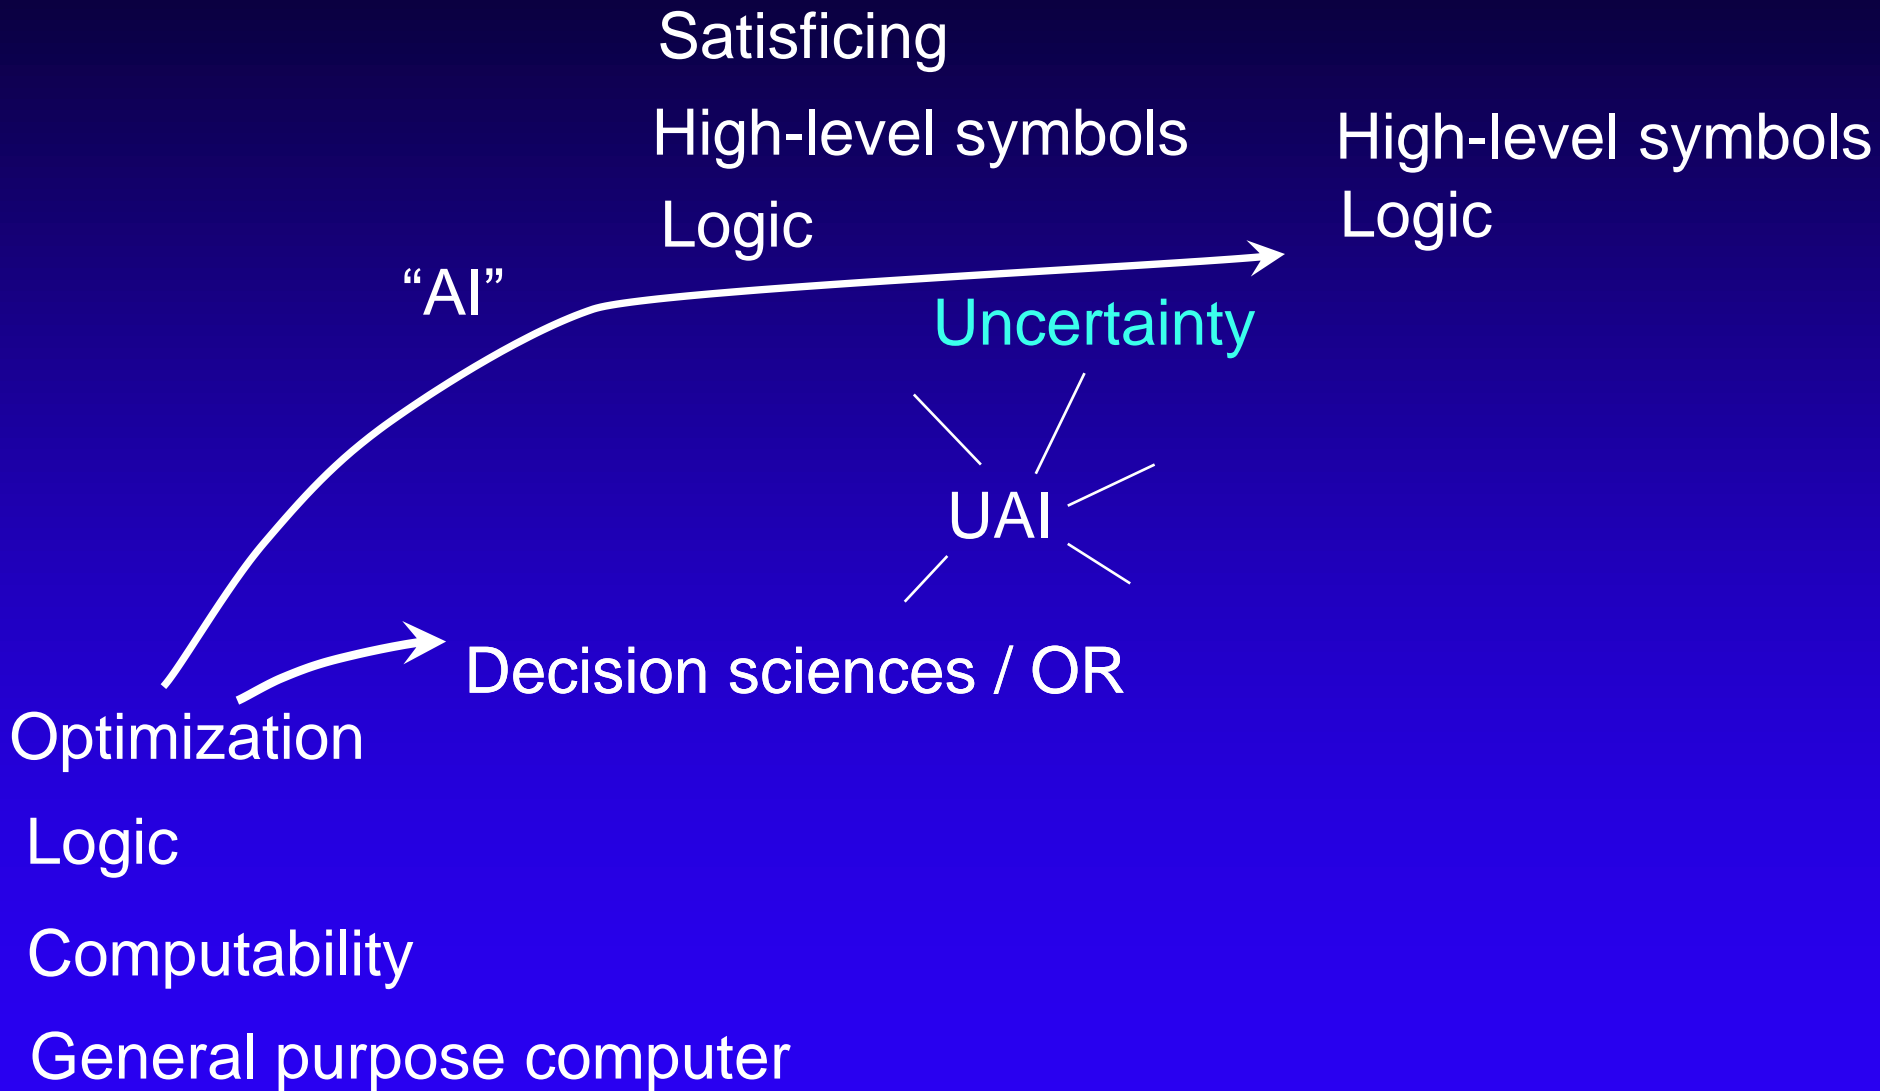
■ *All preconditions required for actions to have intended effects*

## ■ Ramification problem

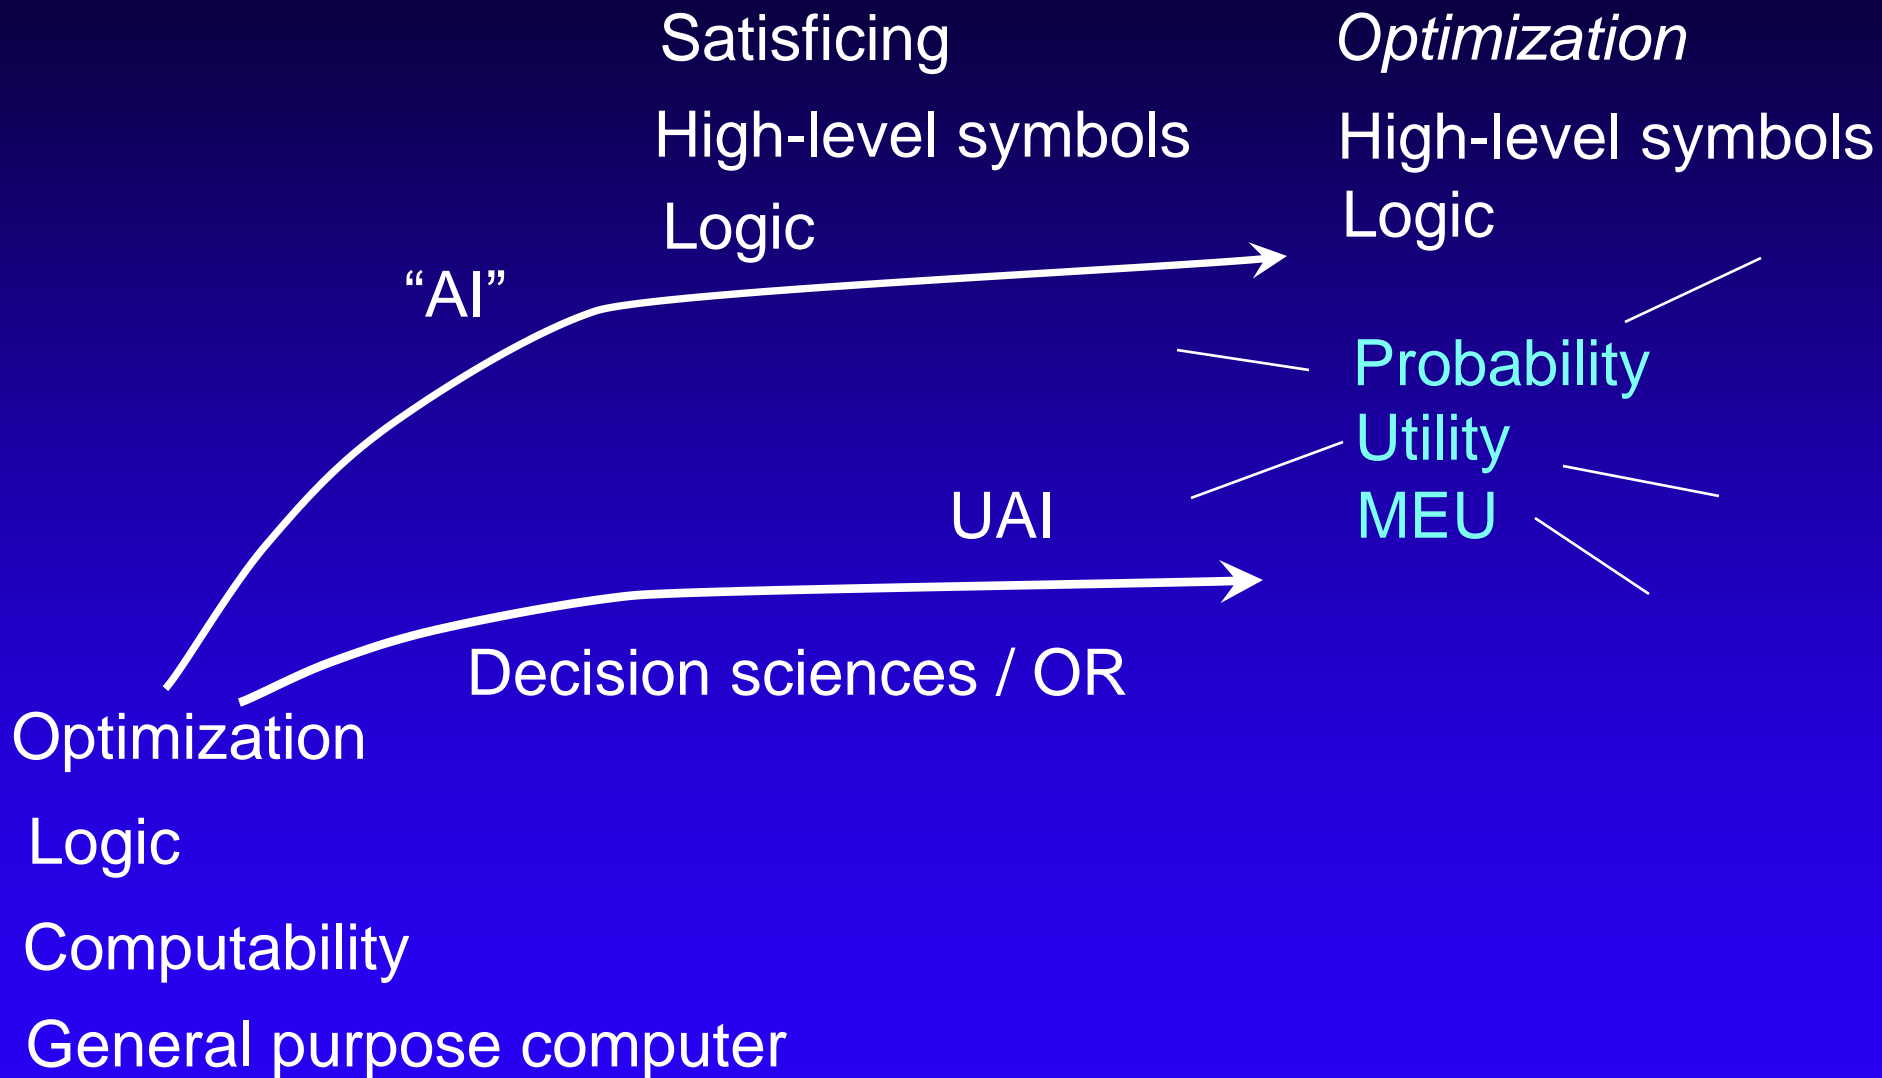
■ *All effects of action*

*Nonmonotonic logics, representation of fluents*

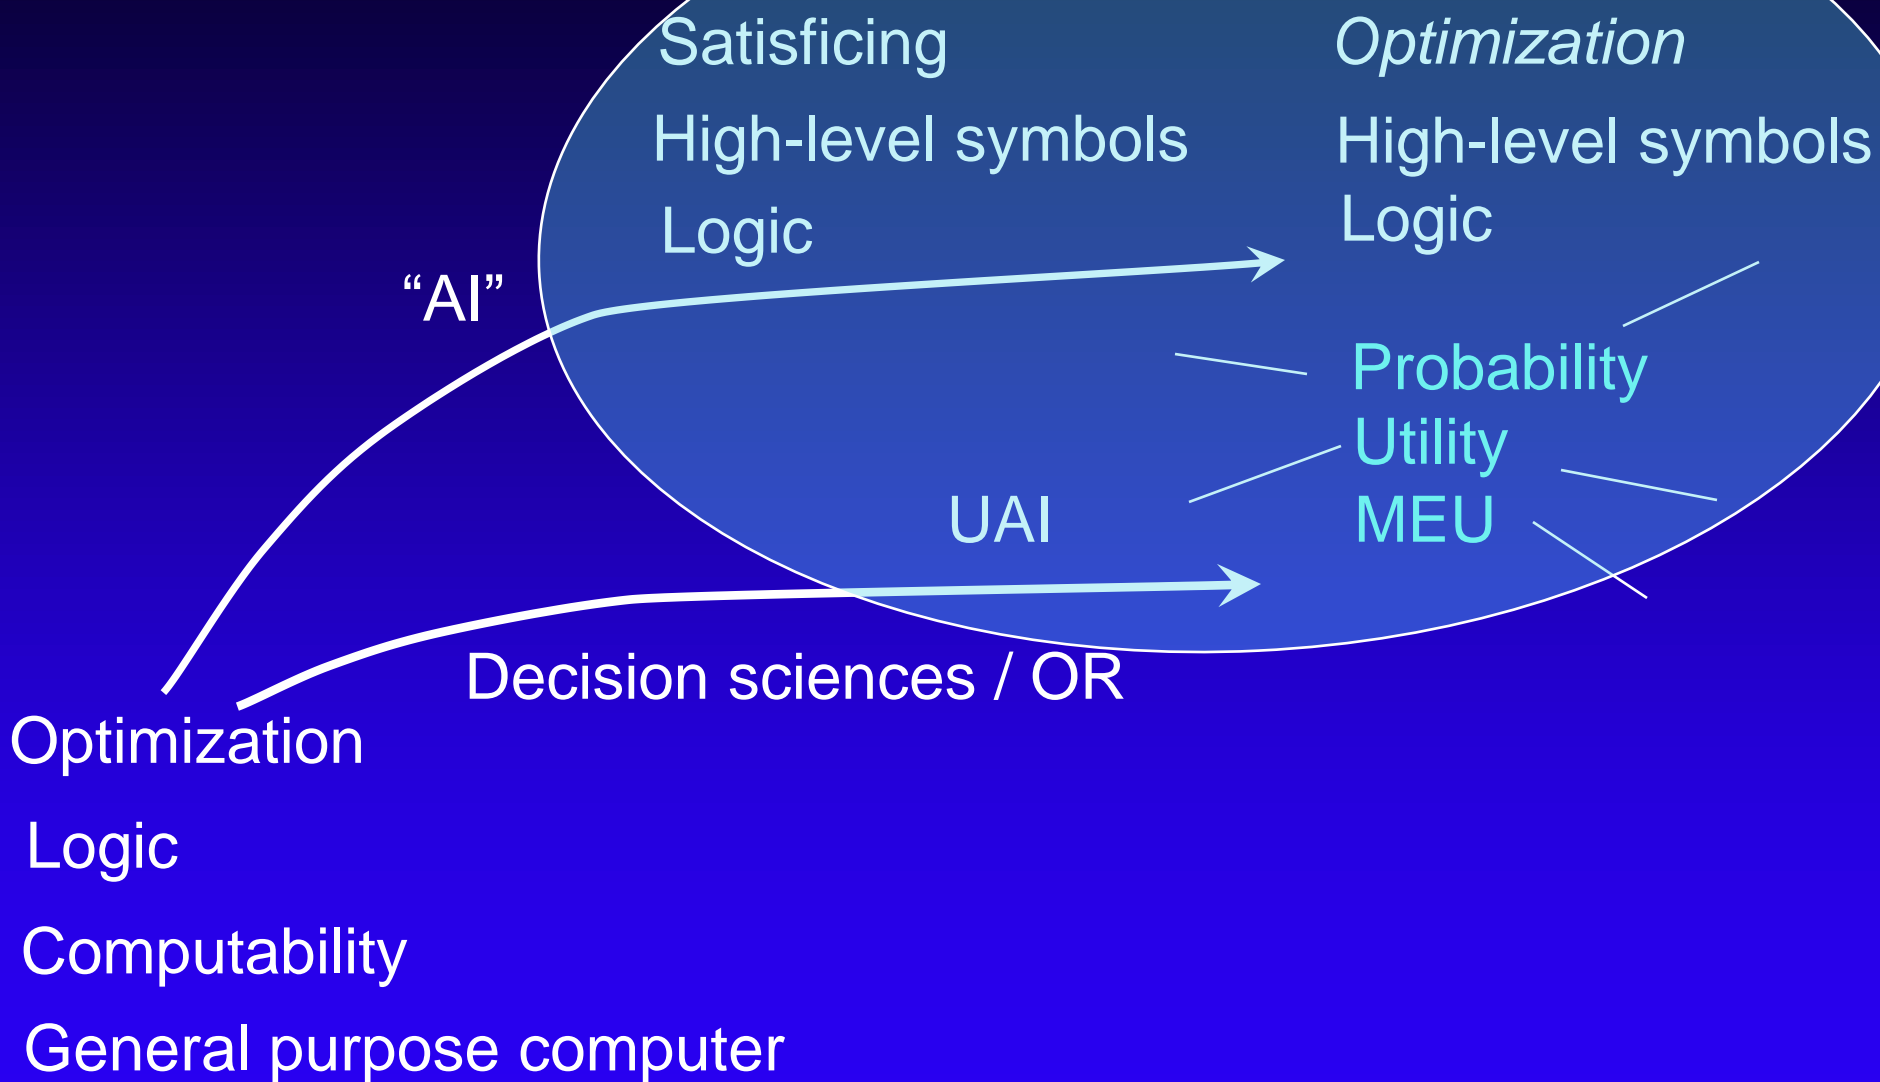
# Paradigms of Computational Intelligence



# Paradigms of Computational Intelligence

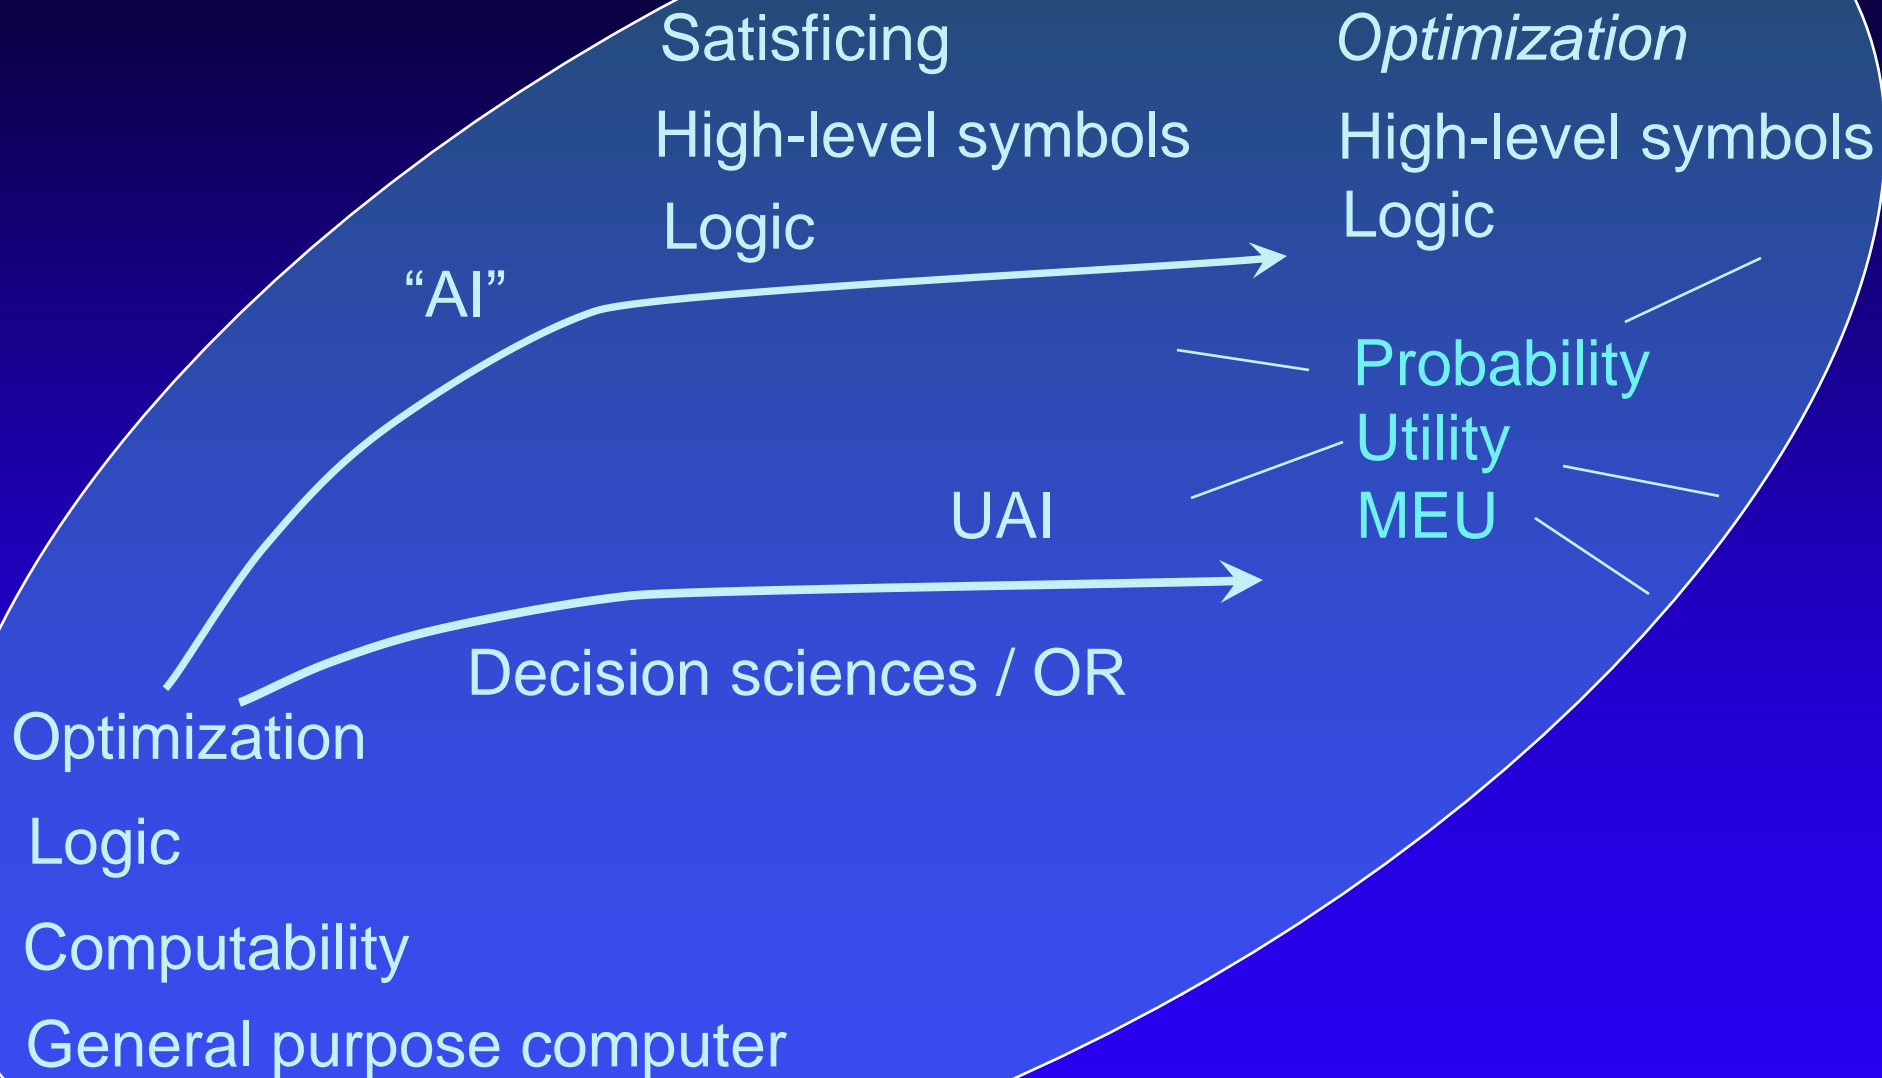


# Paradigms of Computational Intelligence





# Paradigms of Computational Intelligence



# Uncertainty as Organizing Principle

- Incompleteness is inescapable
- Uncertainty is ubiquitous
  - ❖ State of world
  - ❖ Outcome of action
  - ❖ Problem solving itself

*Push unknown & unrepresented details  
into probabilities and propagate*

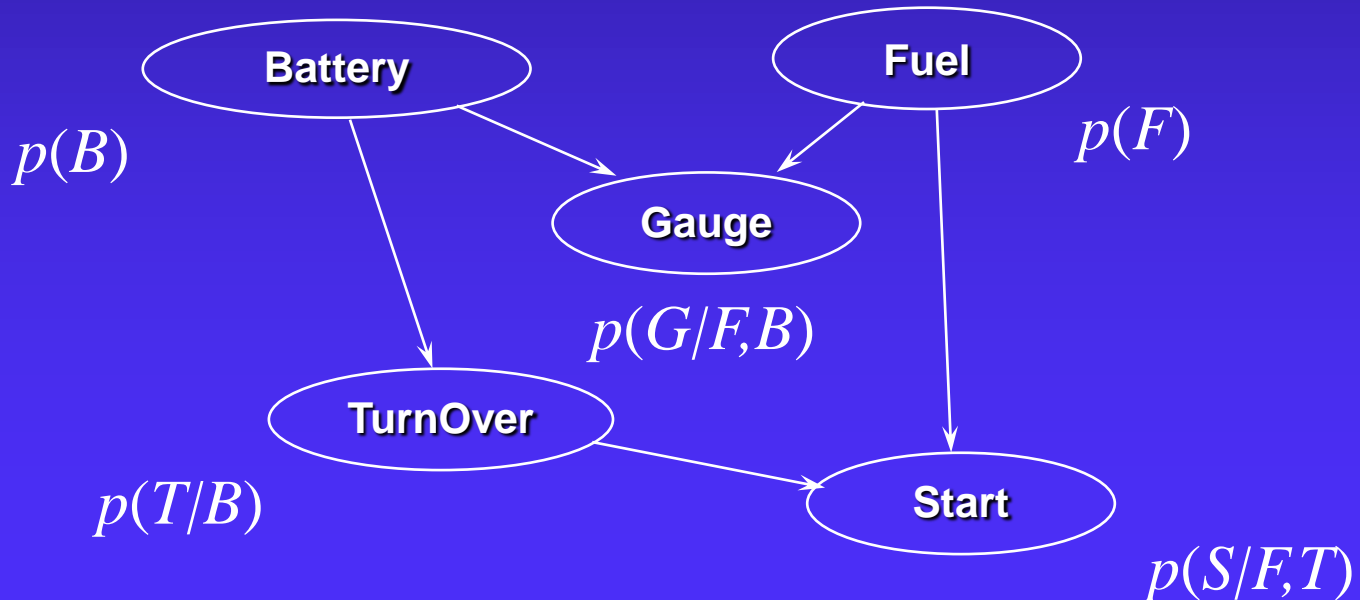
# Uncertainty as Organizing Principle

**→ Machinery for handling  
uncertainty & resource limitations  
foundational in intelligence**

# Expressive Representations of Uncertainty

## ◆ Graphical models for representing beliefs

- $d$ -separation
- Sound and complete algorithm for identifying all independencies entailed by the graph.



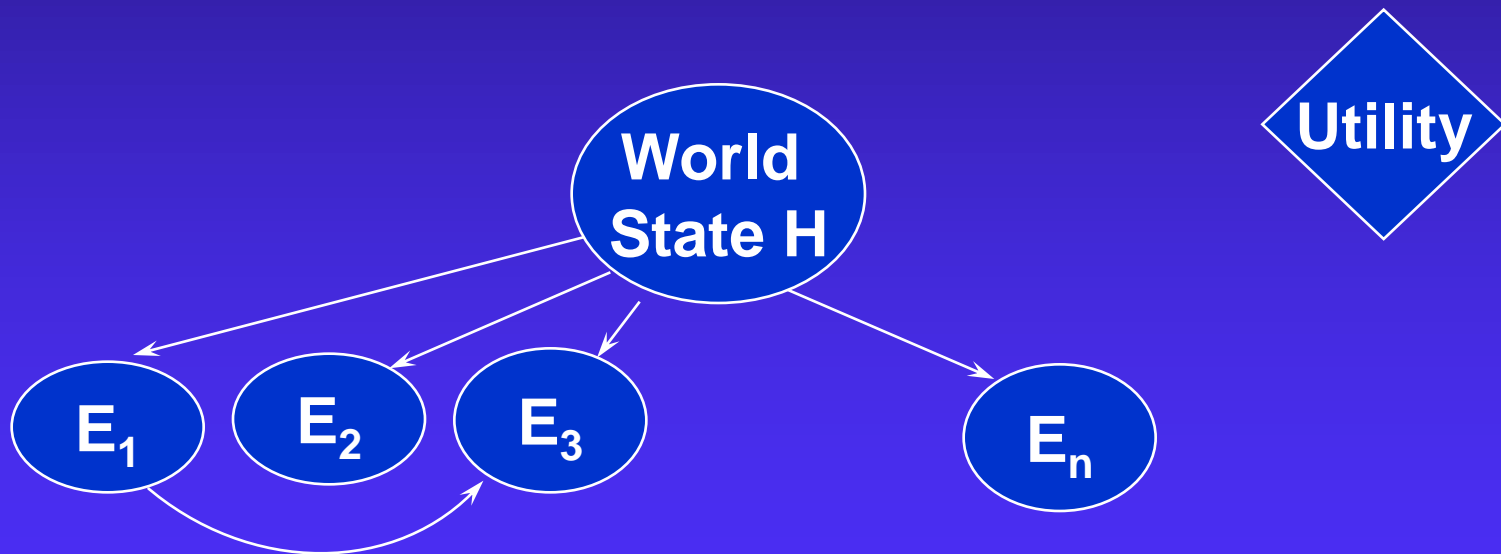
# Representing Problems of Action

- ◆ **Maximum expected utility (MEU)**
- ◆ **Influence diagrams (Howard & Matheson, 1975)**



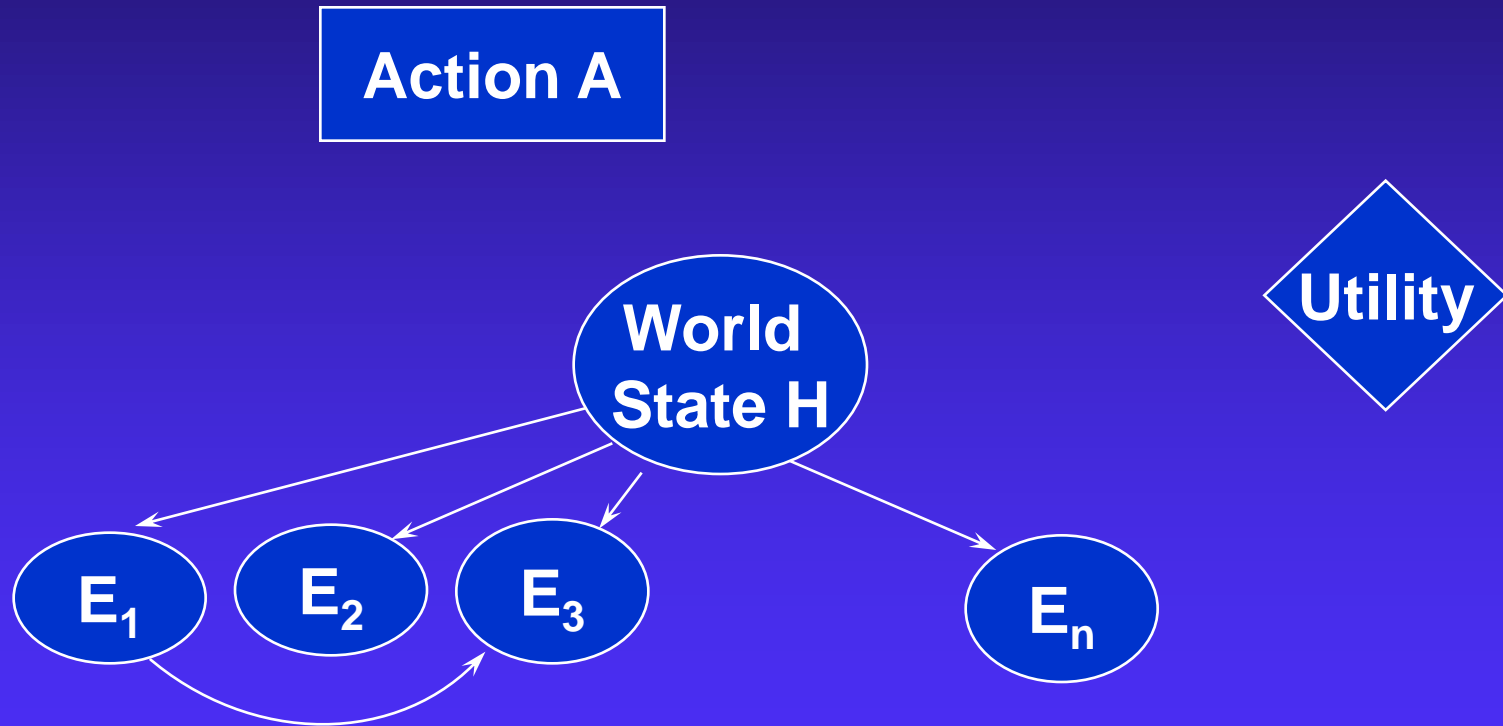
# Representing Problems of Action

- ◆ Maximum expected utility (MEU)
- ◆ Influence diagrams (Howard & Matheson, 1975)



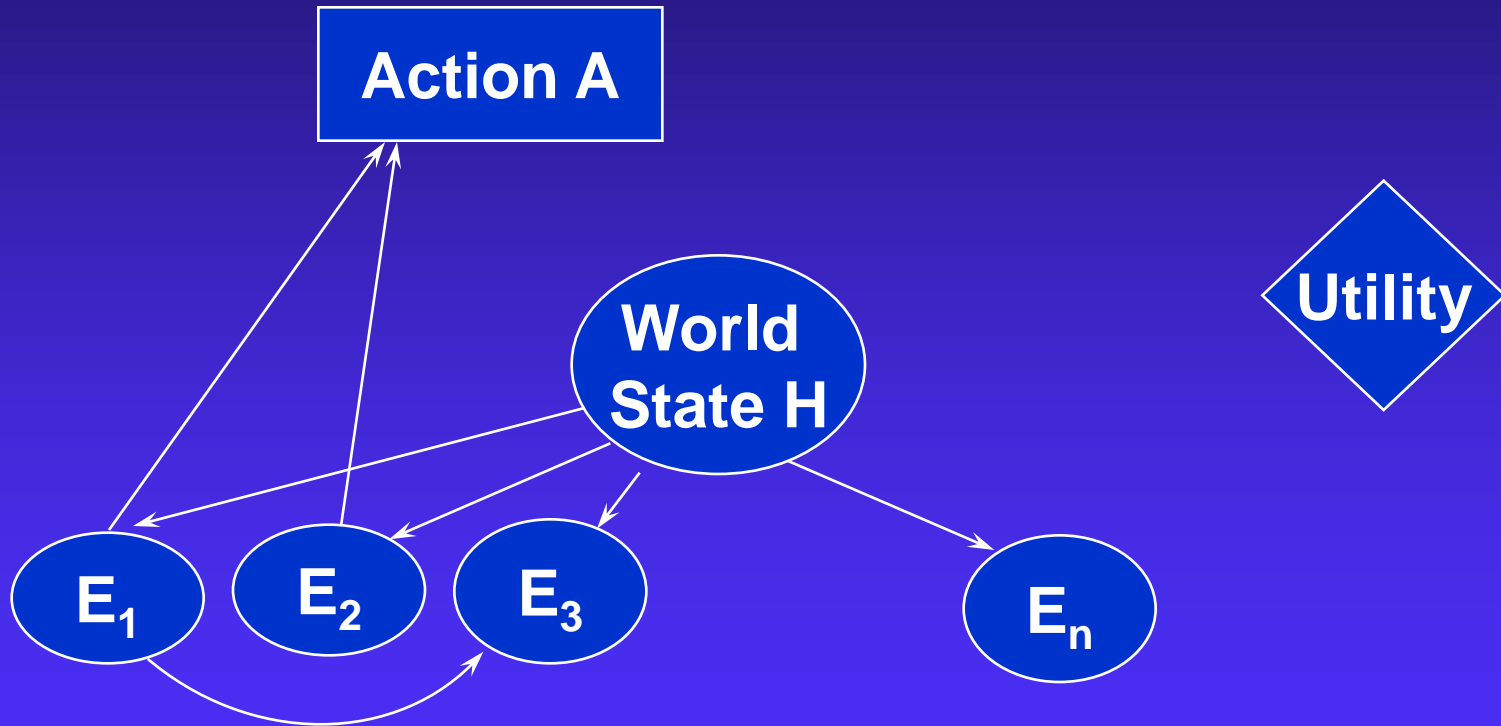
# Representing Problems of Action

- ◆ Maximum expected utility (MEU)
- ◆ Influence diagrams (Howard & Matheson, 1975)



# Representing Problems of Action

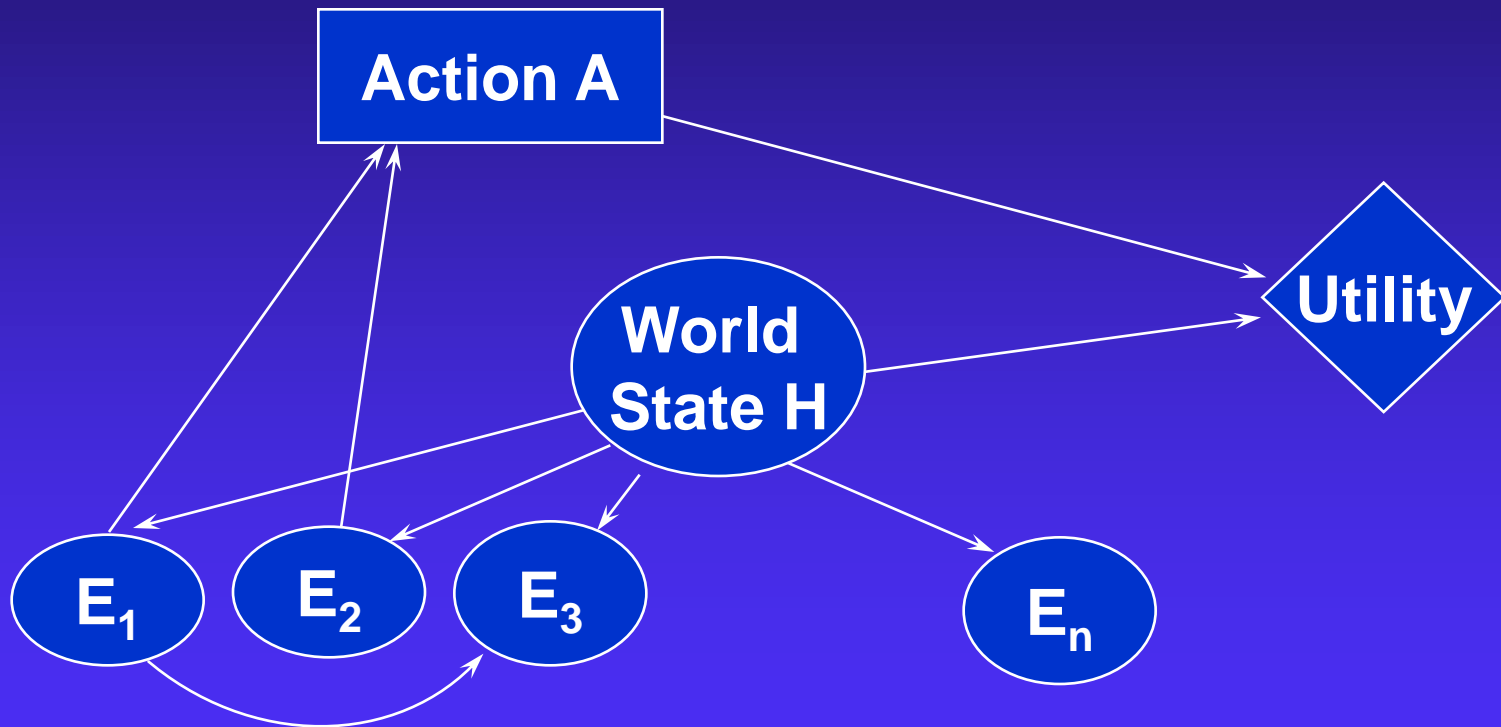
- ◆ Maximum expected utility (MEU)
- ◆ Influence diagrams (Howard & Matheson, 1975)



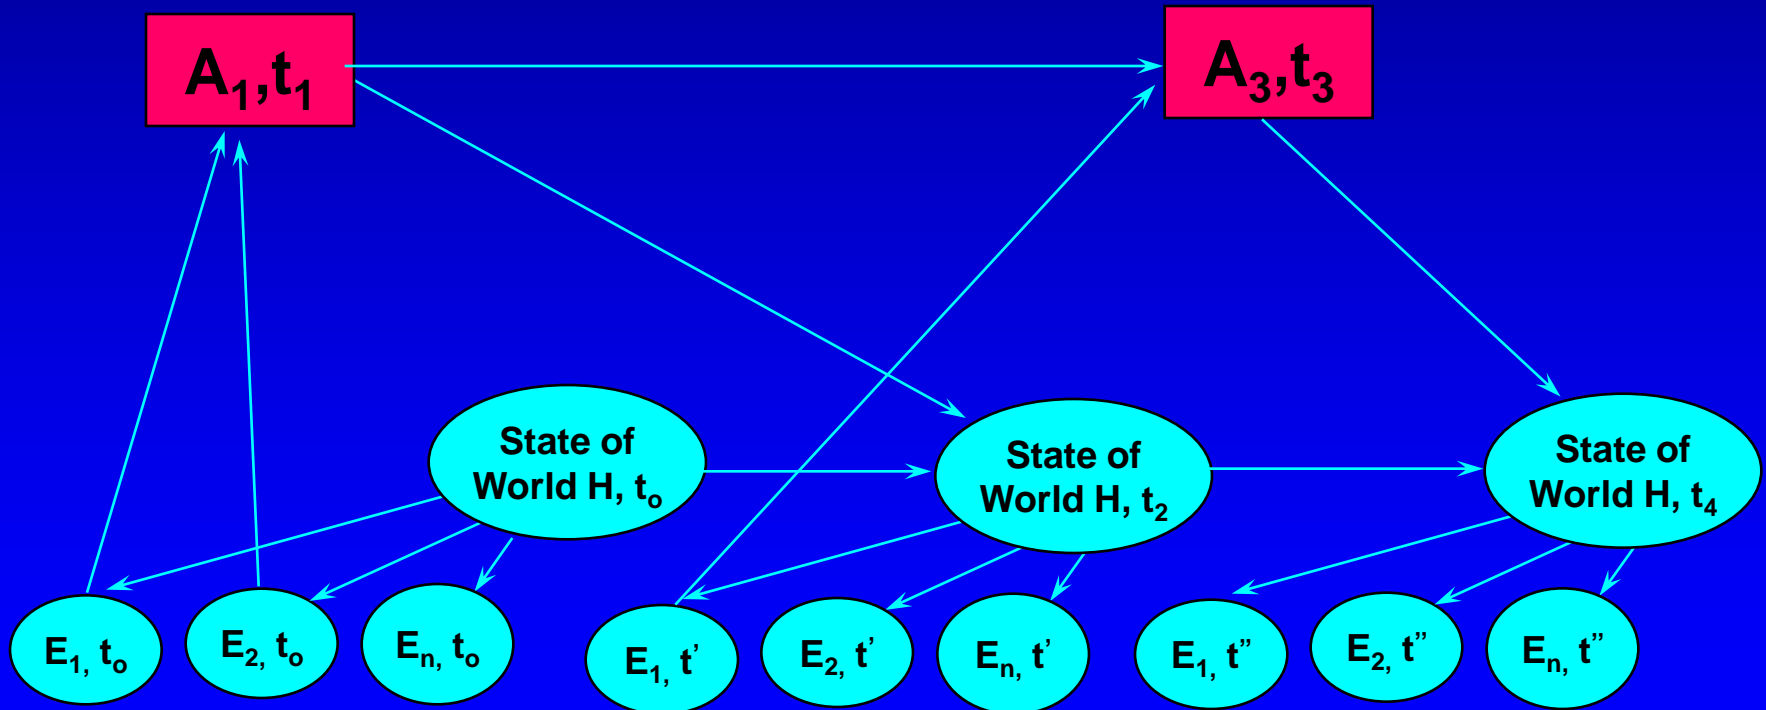


# Representing Problems of Action

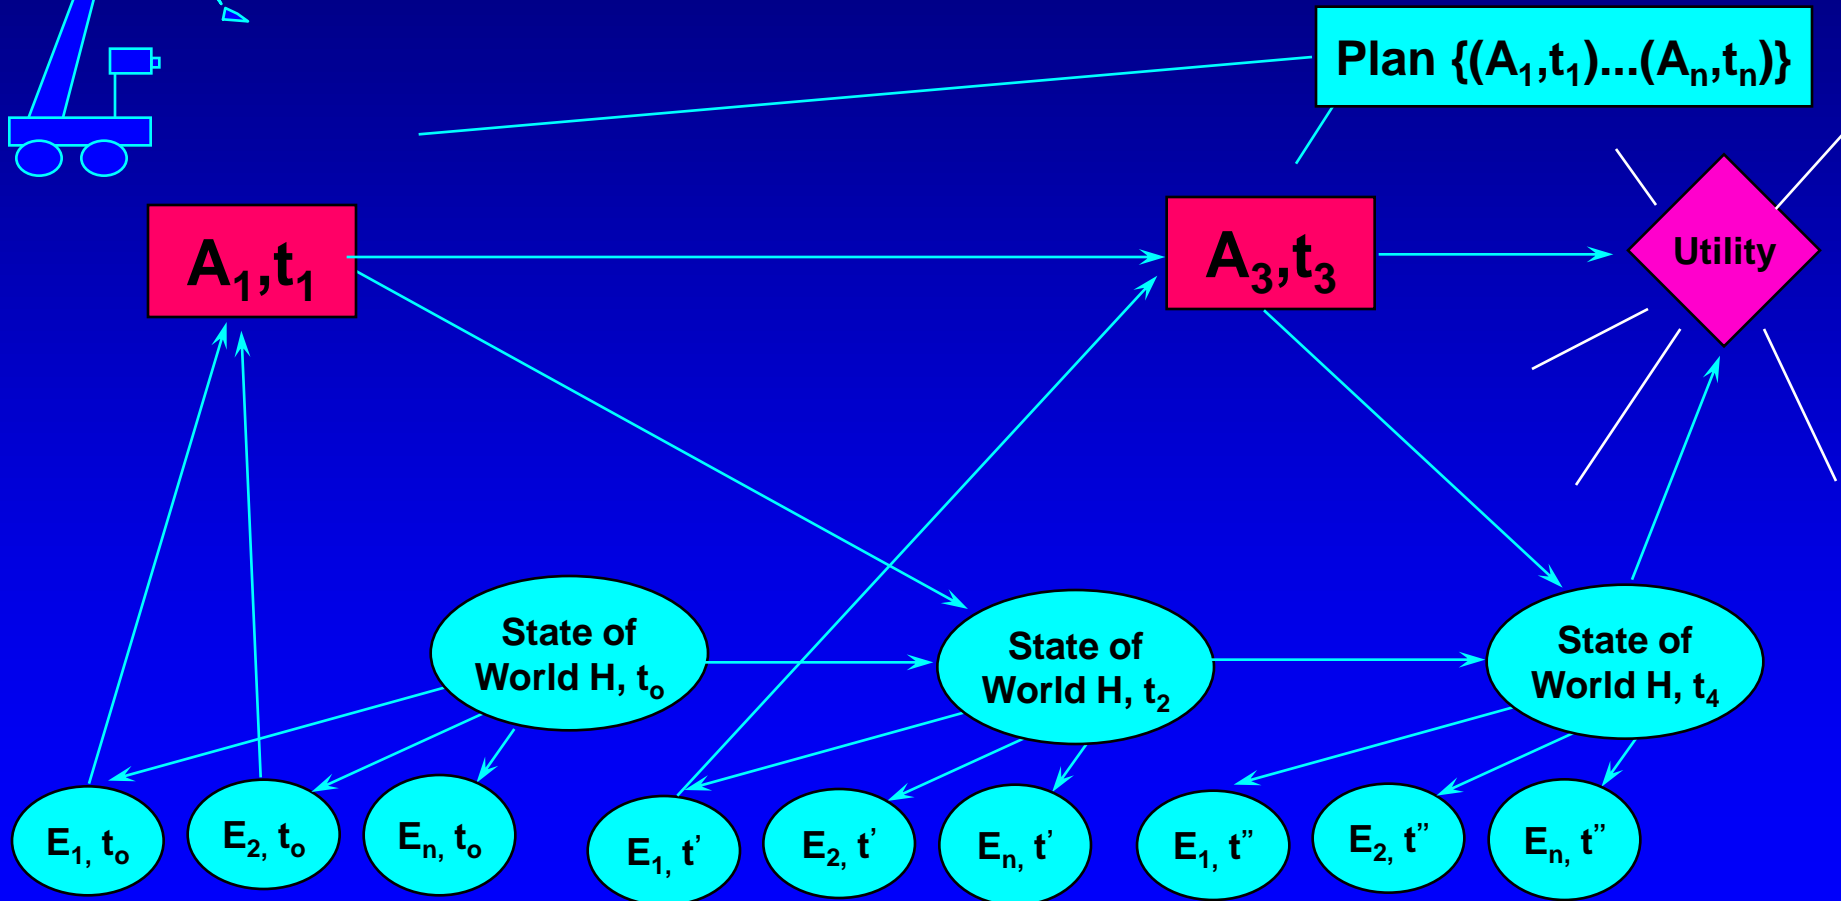

- ◆ Maximum expected utility (MEU)
- ◆ Influence diagrams (Howard & Matheson, 1975)



# Reasoning about Beliefs and Actions over Time




# Reasoning about Beliefs and Actions over Time






# Acceleration of Machine Learning: Discovering Structure and Concepts

- ◆ *e.g.*, Structure search
- ◆ Measure of model likelihood



# Identification of Hidden Variables

- ◆ “Unknown” variables discovered



# Beauty ...and the Bottleneck

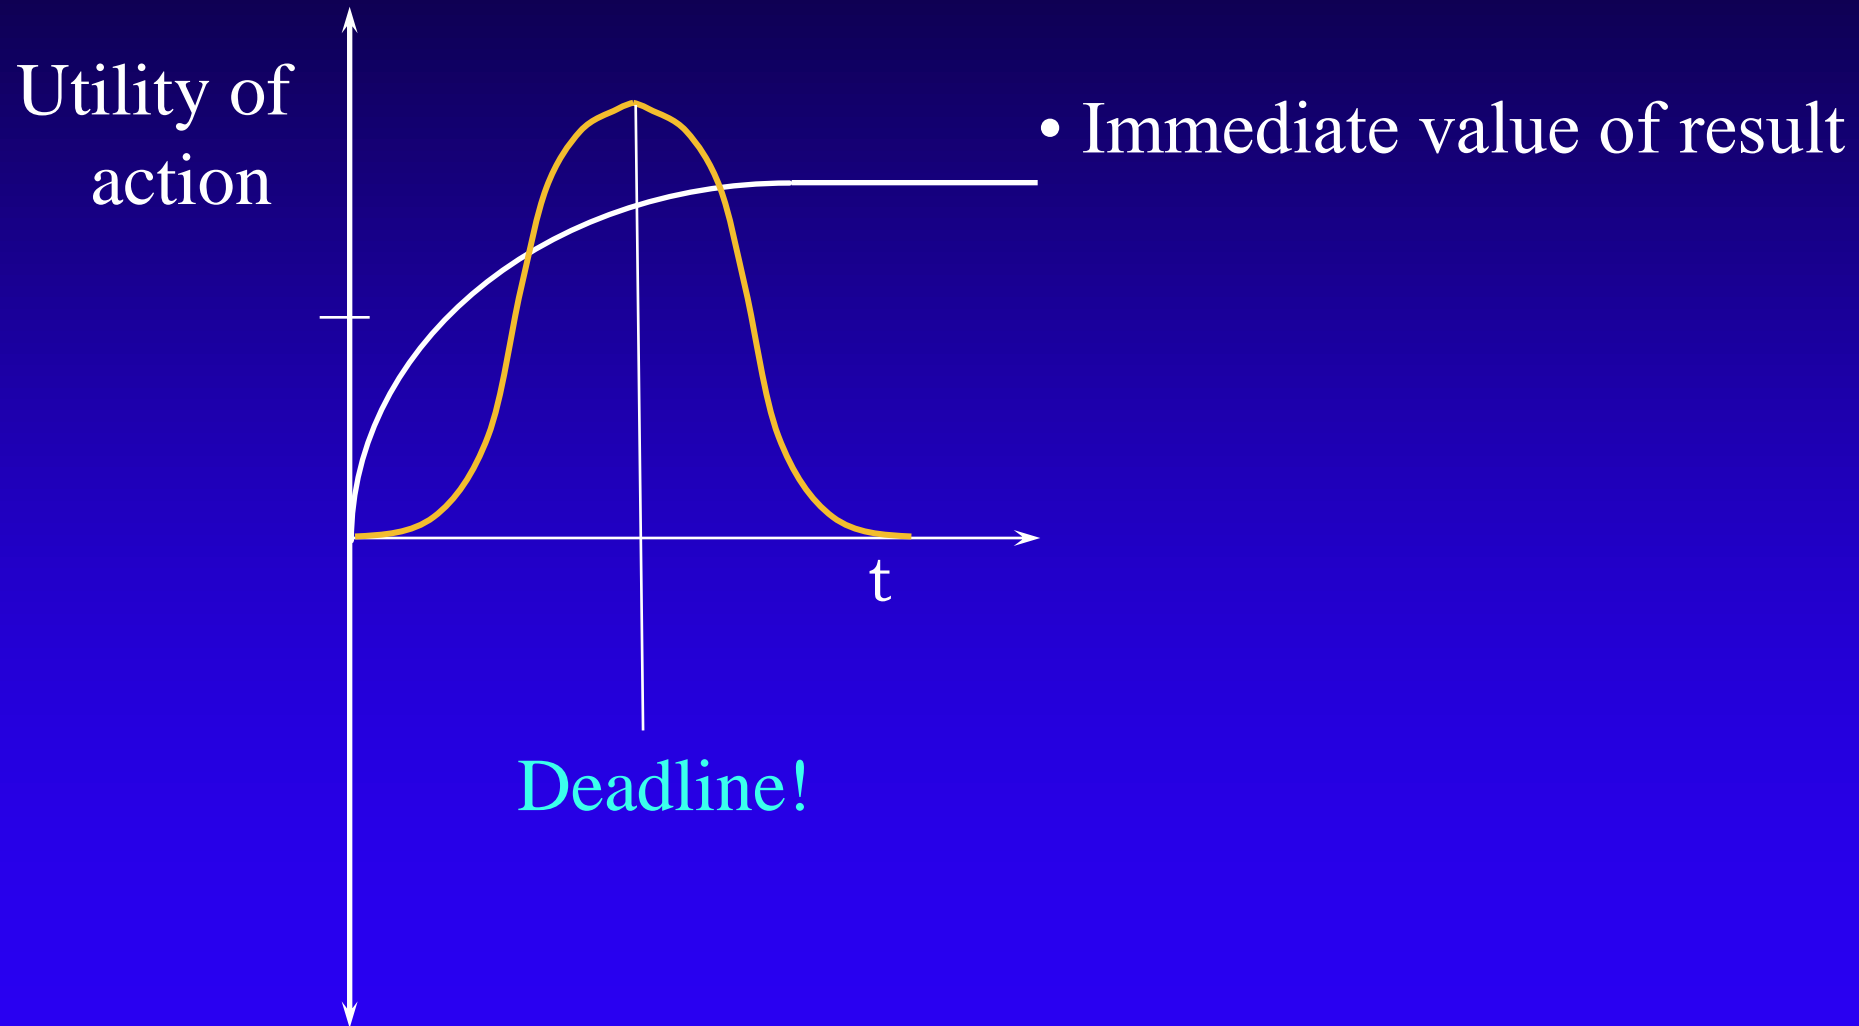
*Insights arising in tight situations*

- ◆ **Bounded rationality** → *Bounded optimality*
- ◆ **Economics of computation**
- ◆ **Flexible computational strategies**
- ◆ **Principles of reflection**

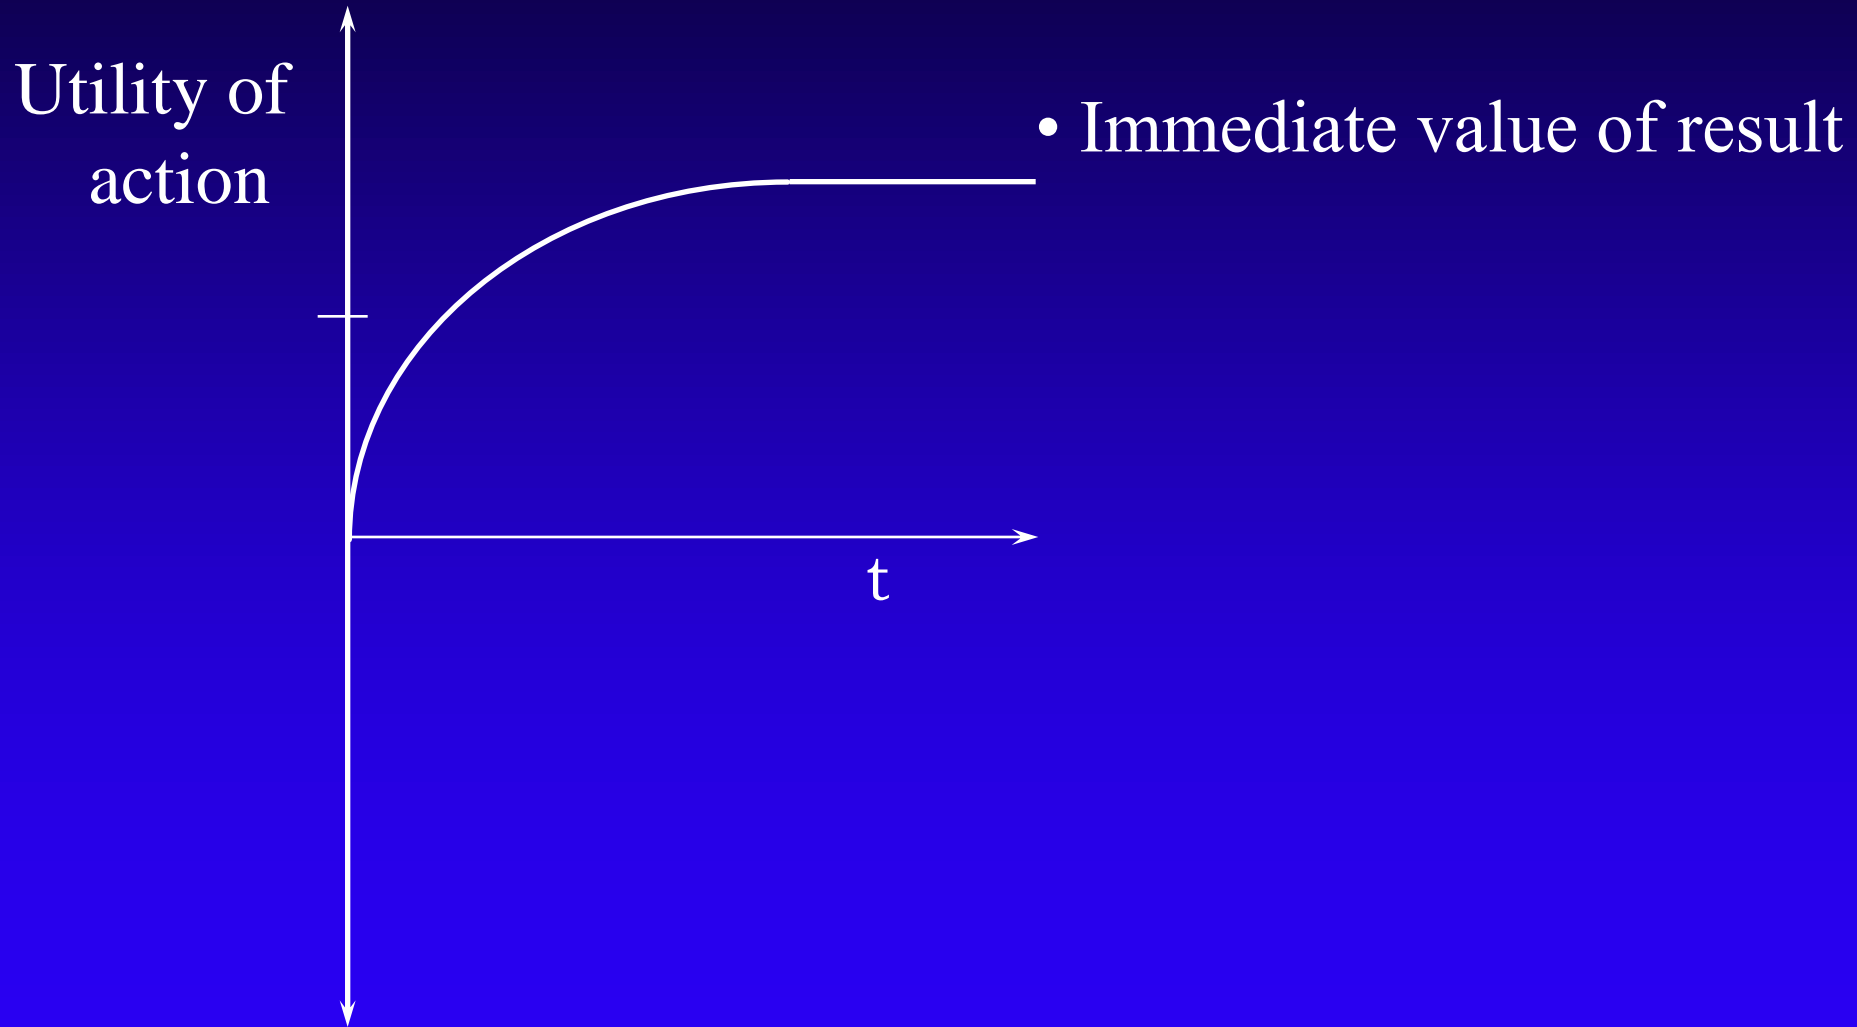


Insights

# Economics of Flexible Computation




# Economics of Flexible Computation

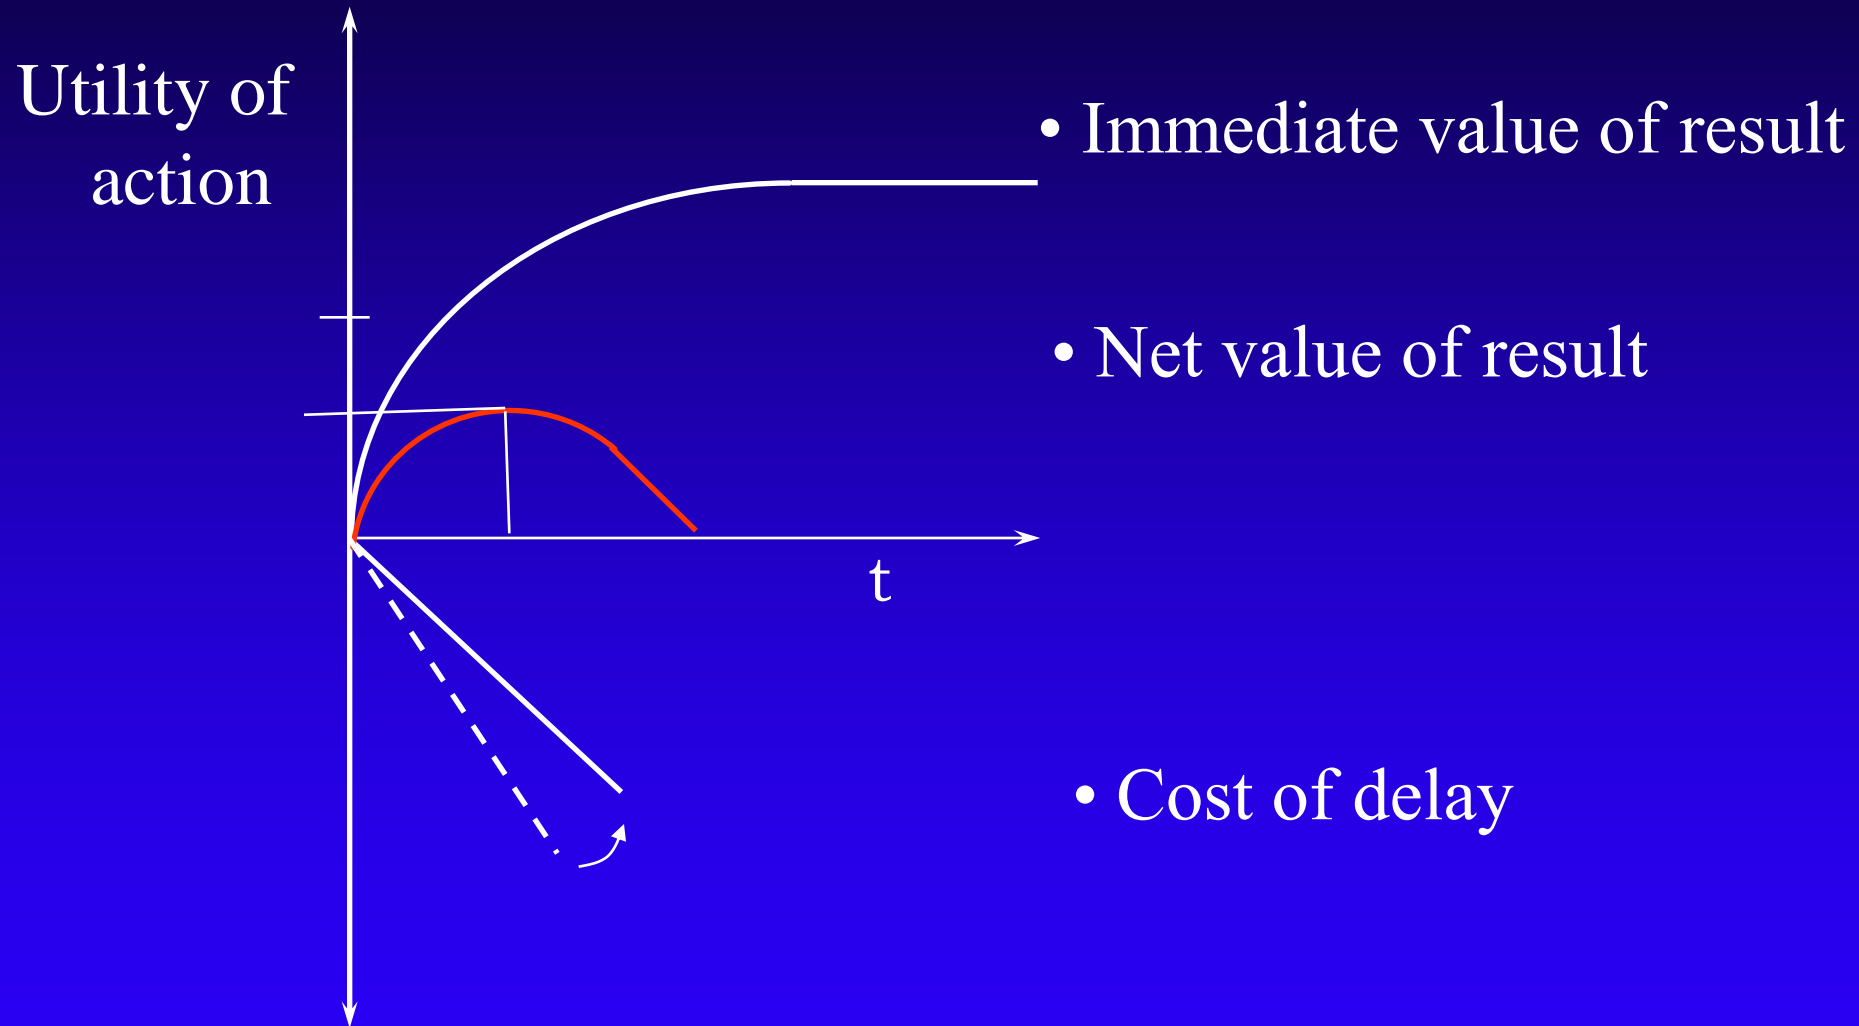




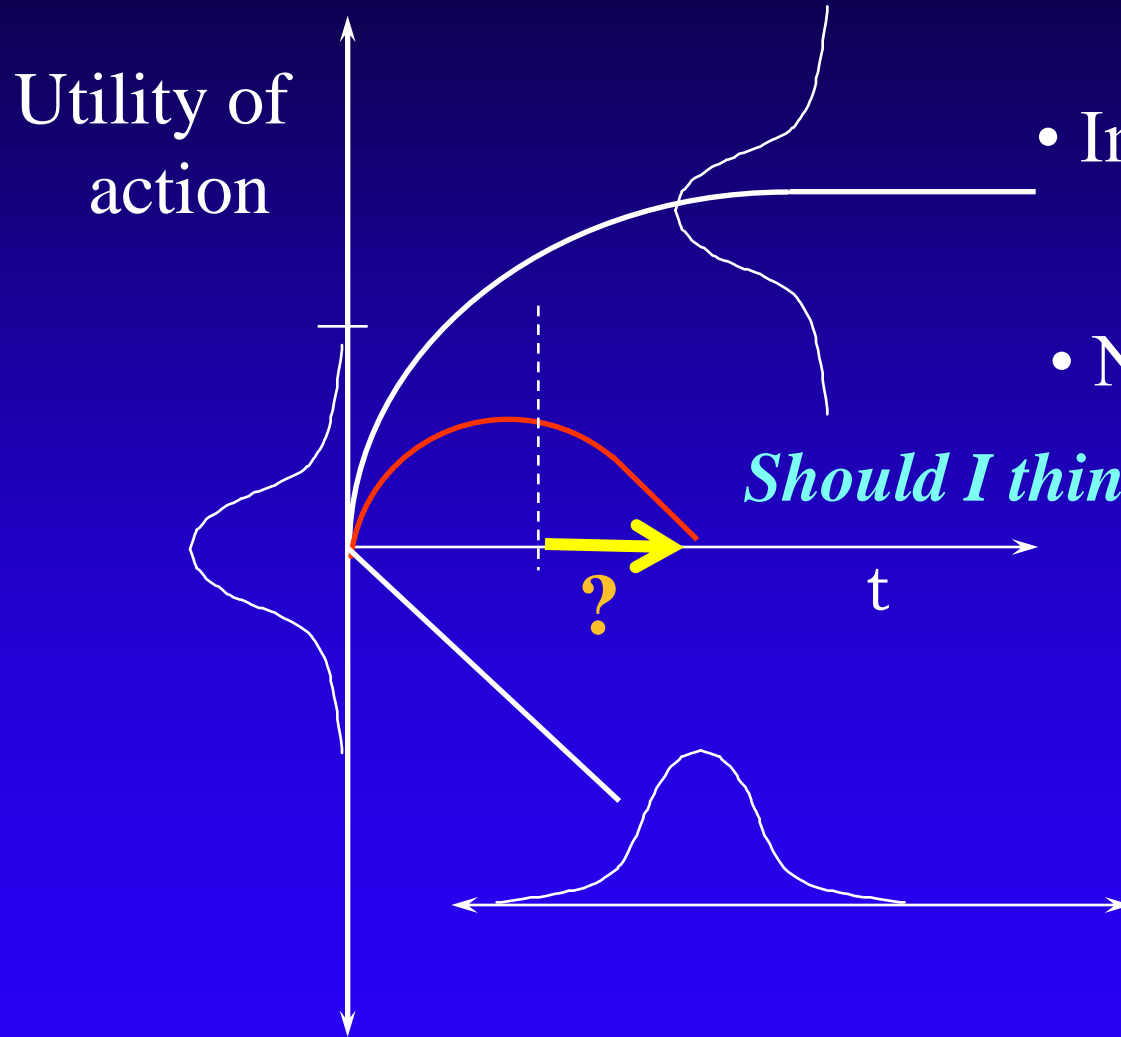
# Economics of Flexible Computation



# Economics of Flexible Computation

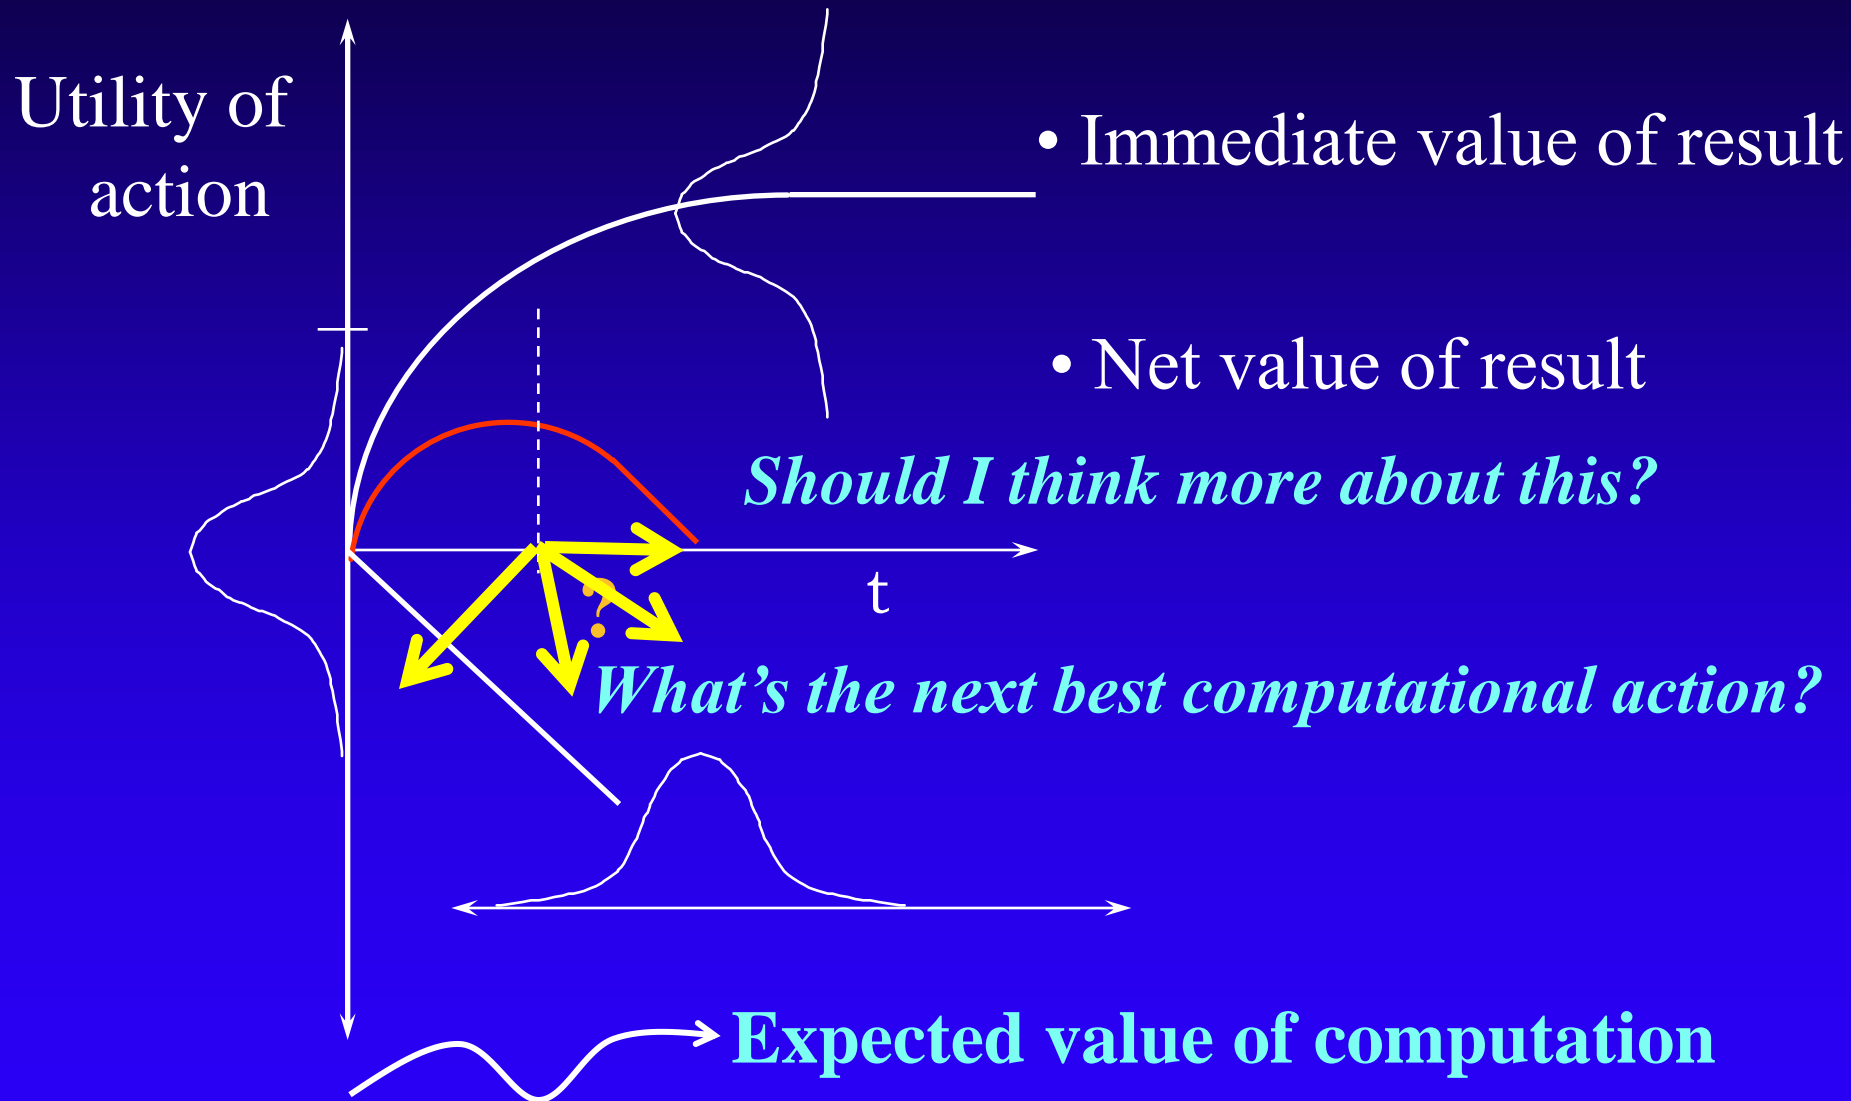


# Flexible Procedures



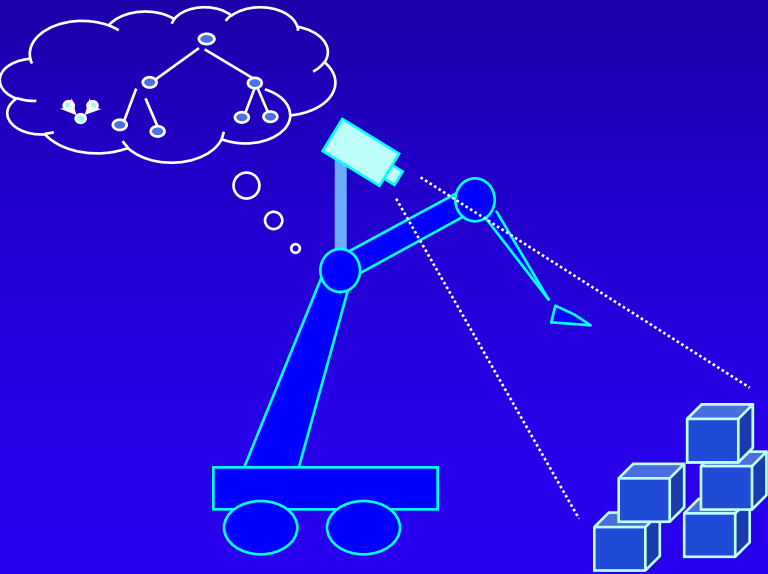
- Immediate value of result
- Net value of result

# Flexible Procedures




# Foundations of Reflection

*Expected value of computation*




# Foundations of Reflection

*Expected value of computation*



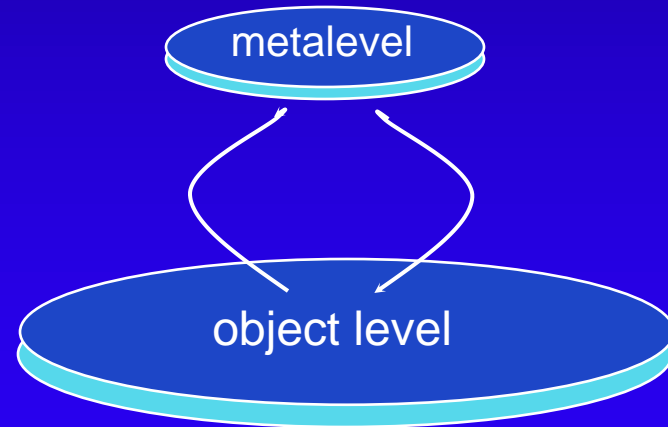
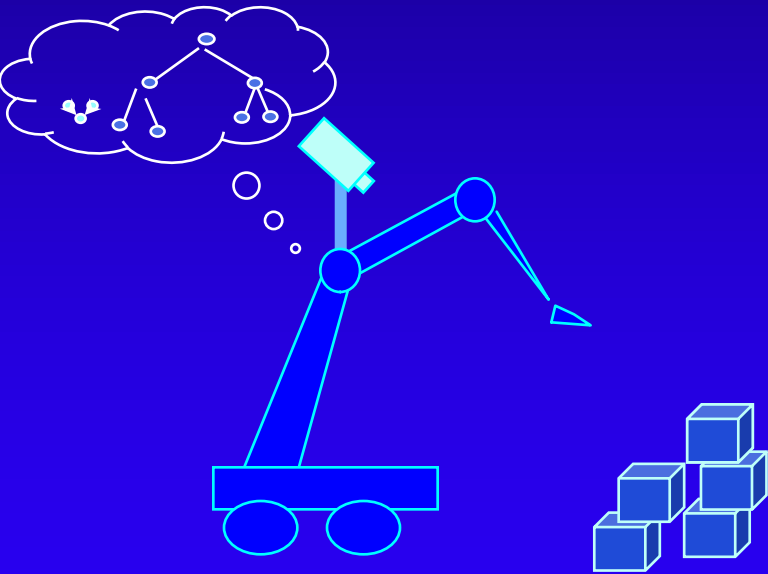
# Foundations of Reflection

*Expected value of computation*



# Foundations of Reflection


*Expected value of computation*





# Real-World: Uncertain performance & Cost of Delay

## Tractable EVC in time-critical medical decisions




# Learning about Instances & Reasoning from a Stream of Problems

e.g., Machine learning & SAT solvers

- Grappling with long tails
- Dynamic restart policy in SAT solvers (Kautz, et al)

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




# Open-World AI



# Toward Situated, Flexible, Long-Lived Open-World Systems

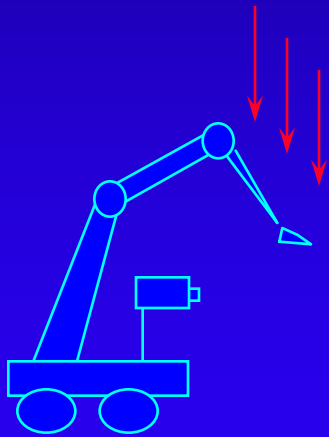
**Flexible adaptation to varying & dynamic situations**

- Streams of observations and challenges over different time frames
  - Broad variation, uncertainty in goals, time criticality, available actions
  - Learning about new objects, predicates, goals, preferences
- ...and about perception & reasoning**



# Flexible adaptation to varying tasks, situations, environments

Situation A



A,E-->MEU

Situation B



Situation C



# Flexible adaptation to varying tasks, situations, environments

Situation A

Situation B

Situation C



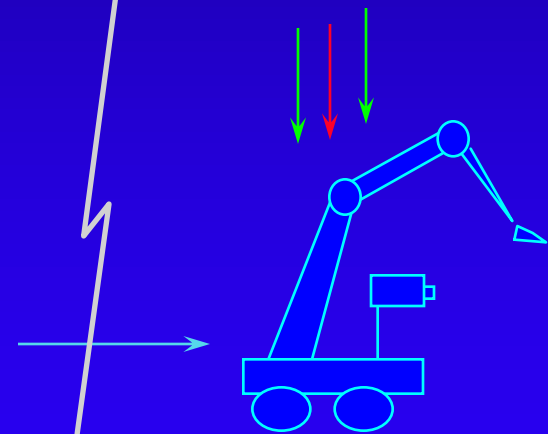


# Flexible adaptation to varying tasks, situations, environments

Situation A

Situation B

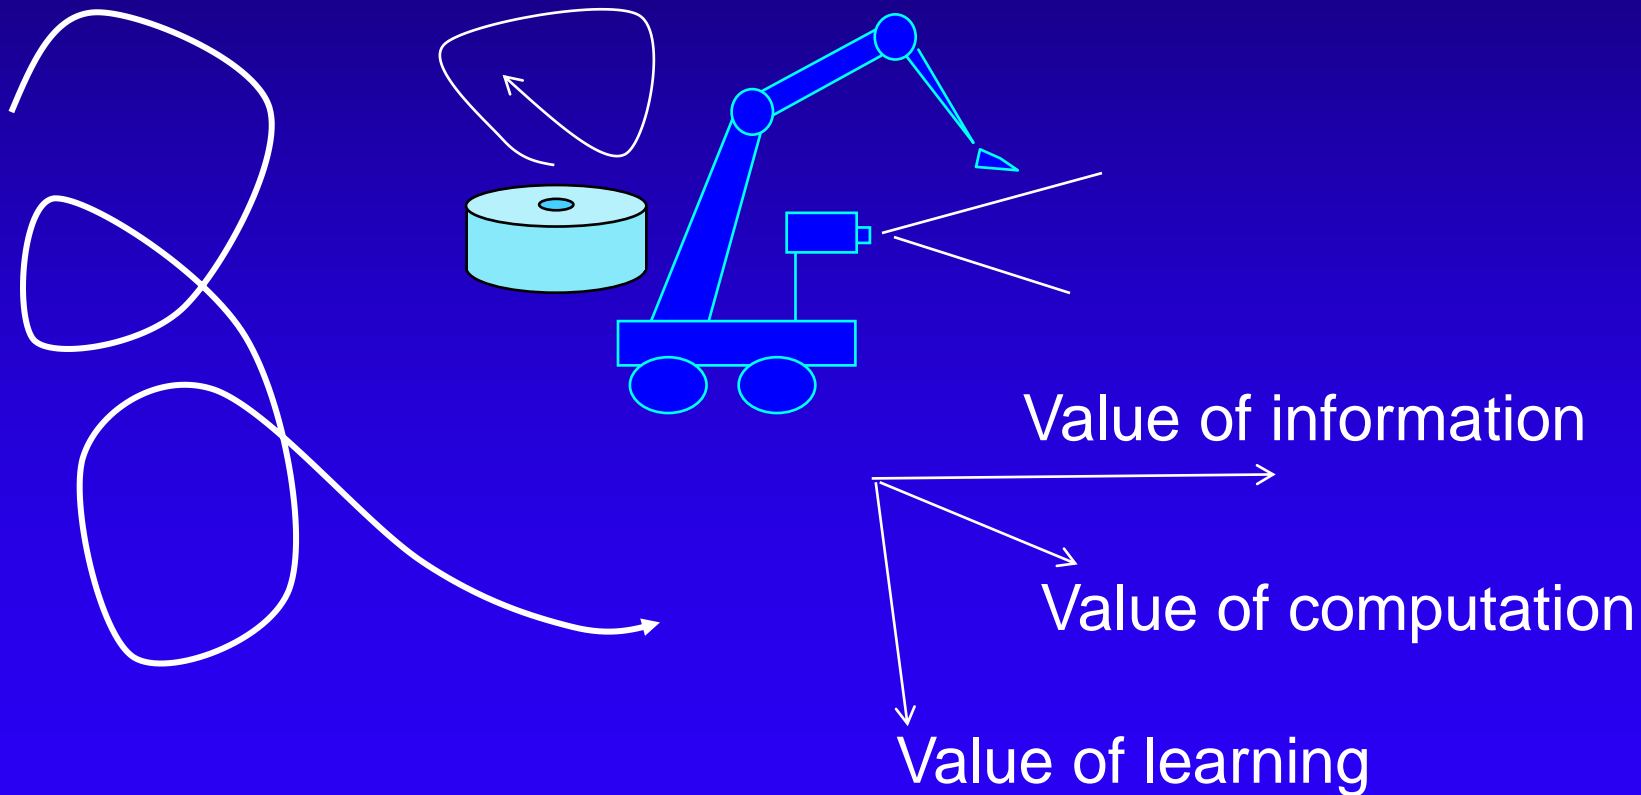
Situation C



A'', E'' --> MEU


# Attentional Challenge: Coordinate Sensing, Reflection, Action, Learning

Standing challenge...



# Challenge: Lifelong Learning

- ◆ Tradeoff local costs of exploration, labeling for long-term gains




# Challenge: Handle Streams of Problems

## ◆ Policies for using all time...incl. *idle* time


---

- How to best use available time to solve future problems?
- How to trade current problem solving for solving future problems?


# Stream of Problem Instances




# Stream of Problem Instances



# Trading Off Present for Future



# Trading Off Present for Future







# **Challenge: Frame – and Framing Problem**

- ◆ **What goals, preferences, ...objects, predicates, relationships should be in a decision model?**
- ◆ **How can tractable, relevant models be constructed automatically?**
- ◆ **How can the system learn about the frame?**


# Automated Framing & Execution of Local Decision Problems

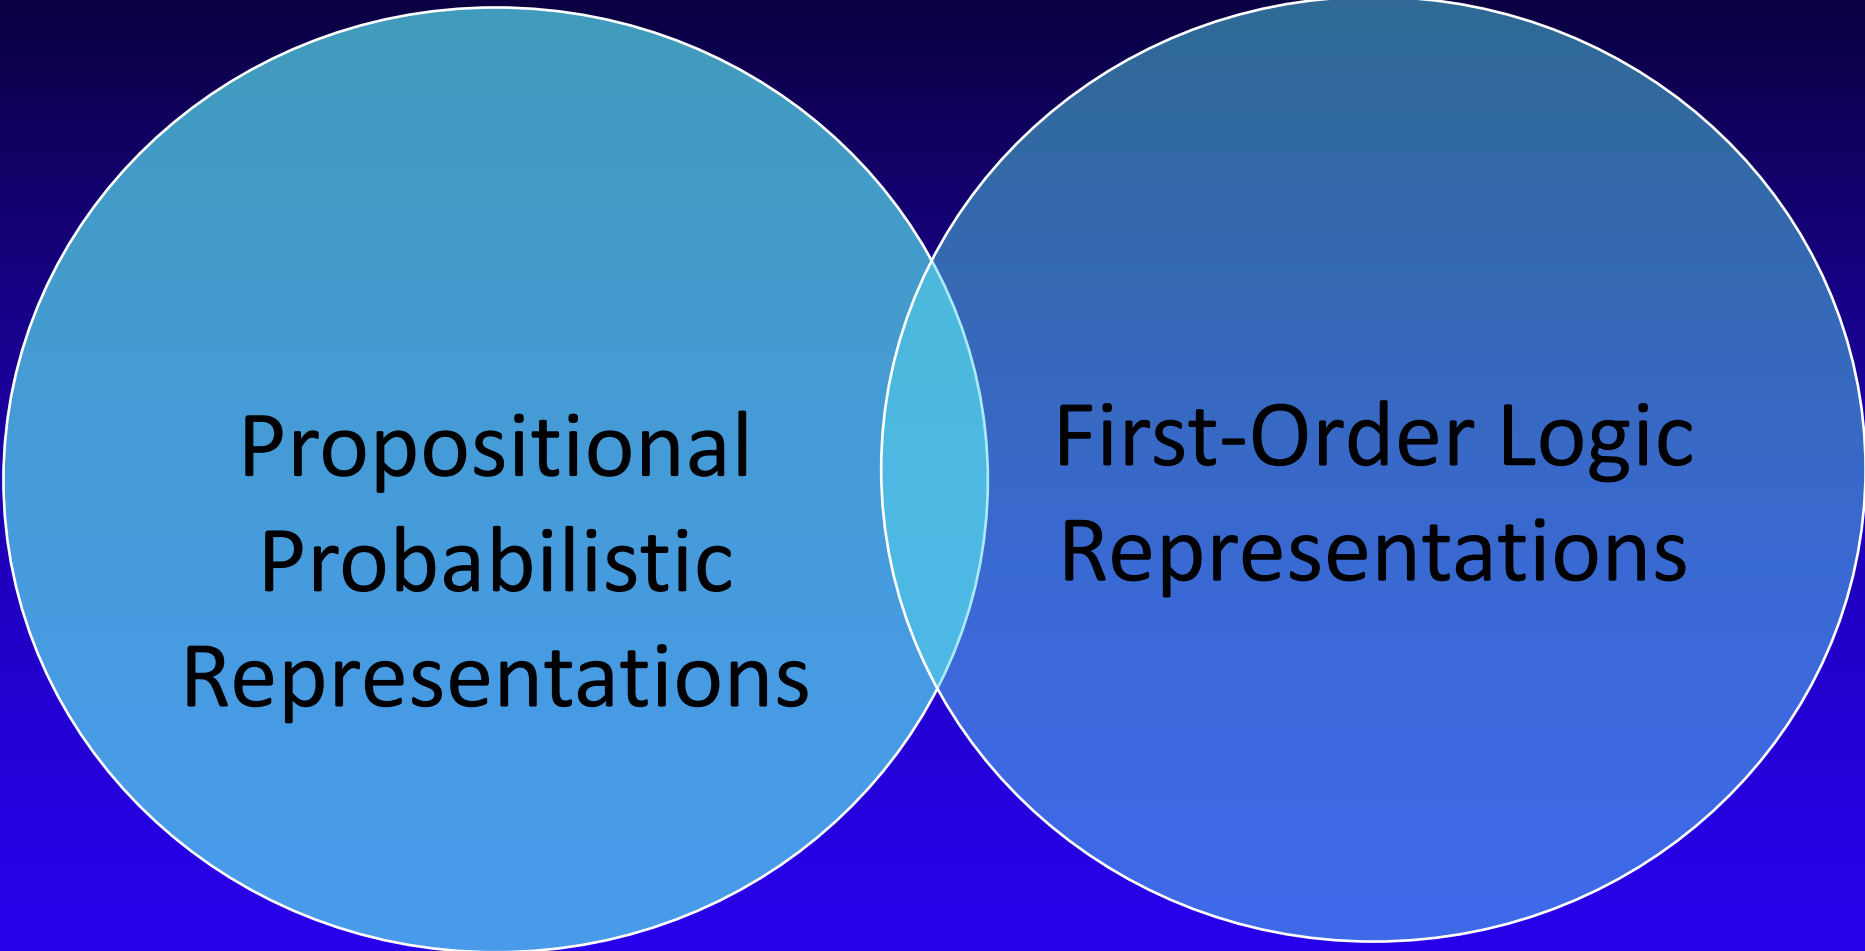


# Automated Framing & Execution of Local Decision Problems



# Automated Framing & Execution of Decision Problems





Propositional  
Probabilistic  
Representations

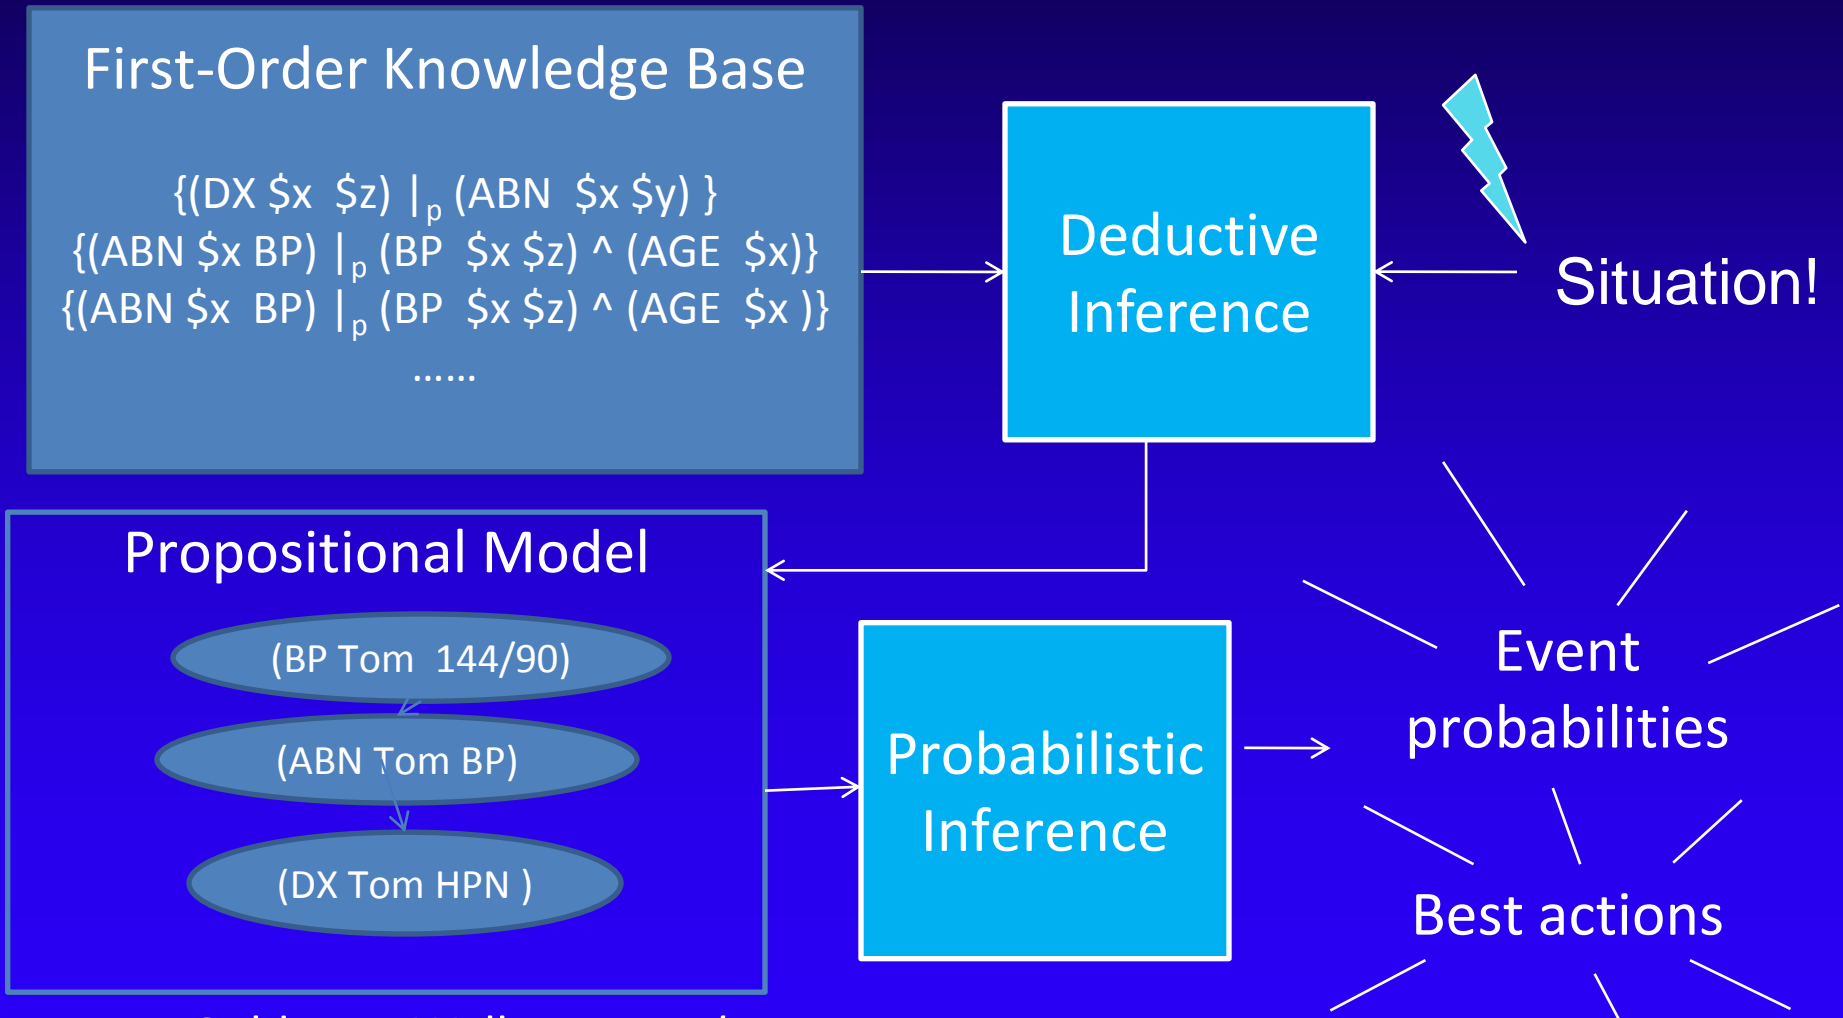
First-Order Logic  
Representations

Propositional  
Probabilistic  
Representations

First-Order Logic  
Representations

# Knowledge-Based Model Construction

## Context-sensitive propositional models from first-order knowledge base





# Learning & First-Order Probabilistic Representations

- ◆ **Generate objects, relations, models**
  - *Plan recognition networks*
  - *Probabilistic relational models*
  - *Markov logic networks*
  - *BLOG: Probabilistic models with unknown objects*



# Reasoning about ...and Expecting the Unknown


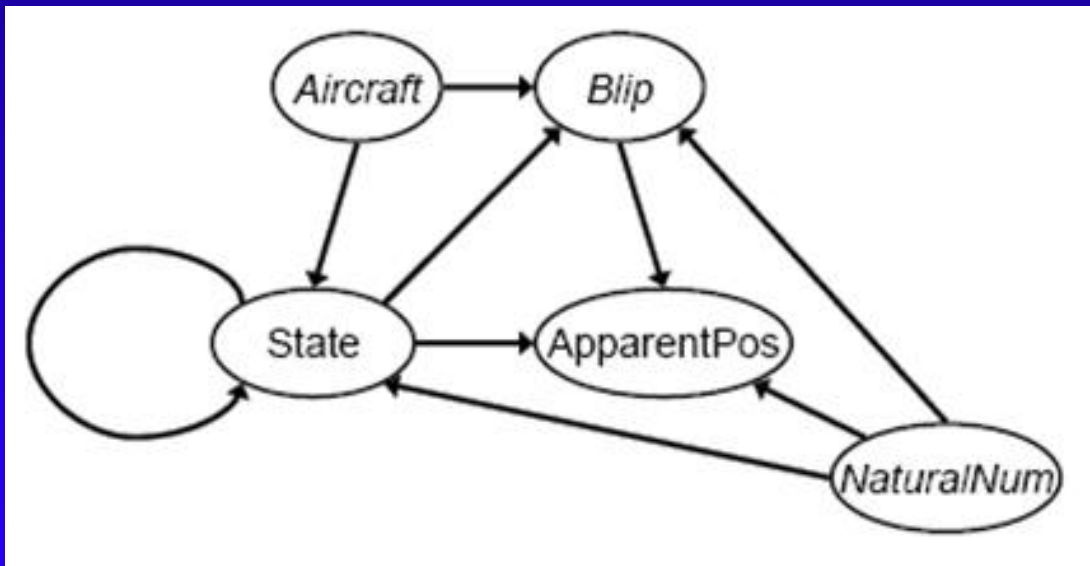
## Probabilistic Models with Unknown Objects



Brian Milch, *et al.*

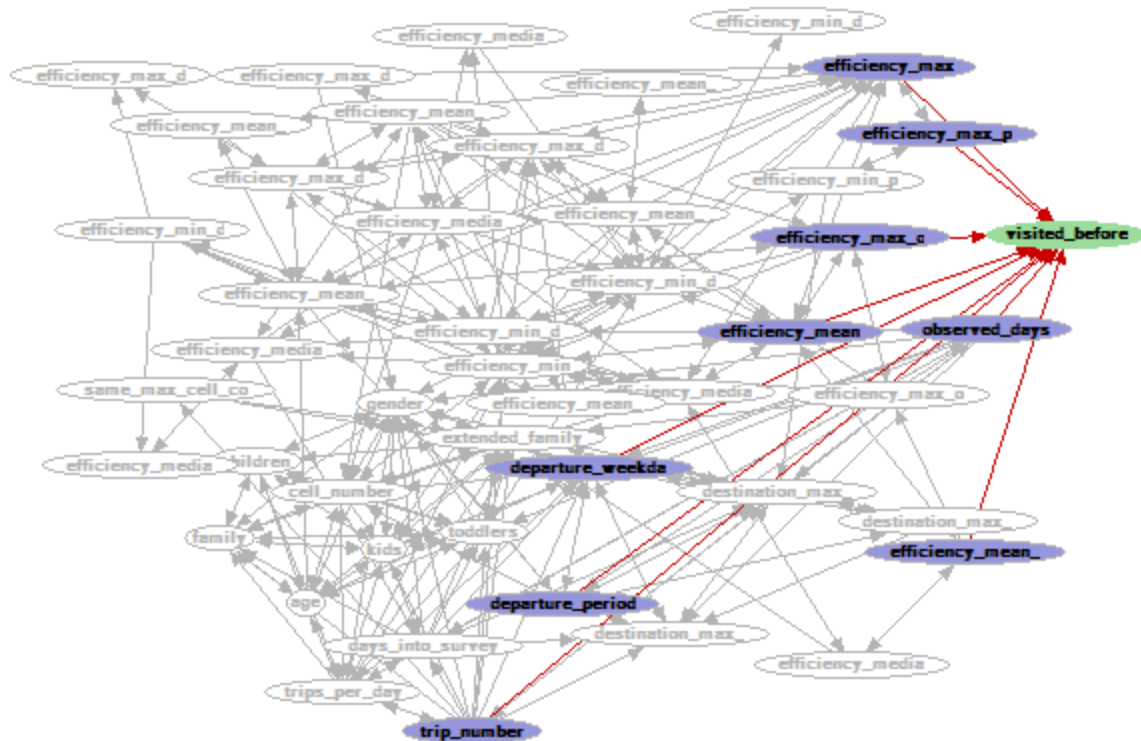
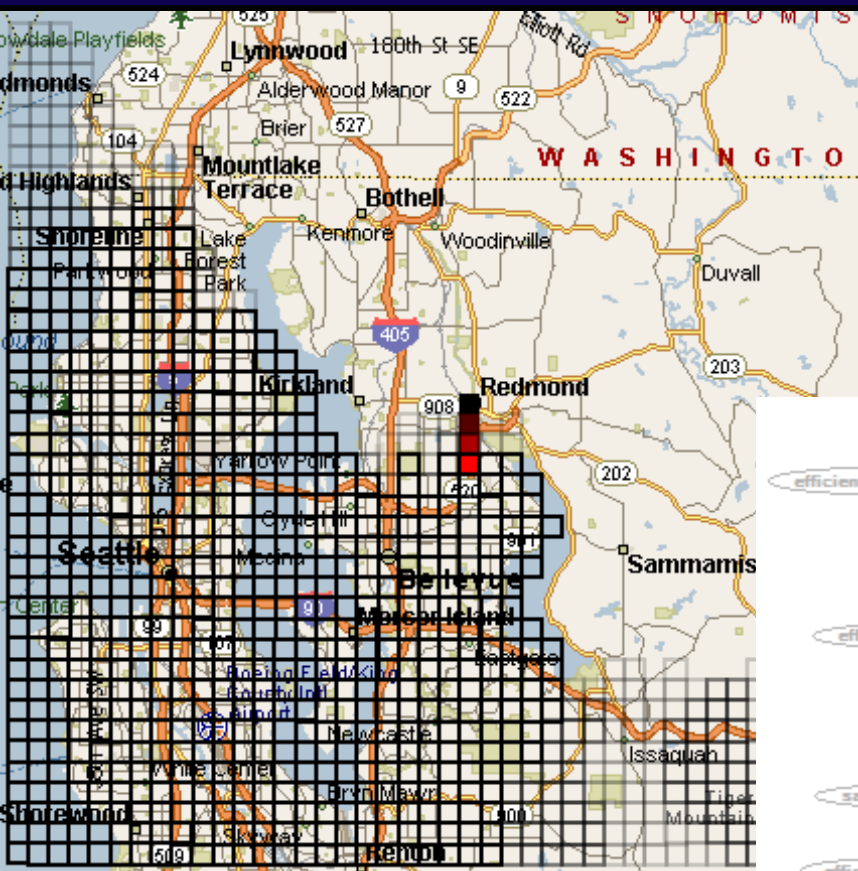
# Reasoning about ...and Expecting the Unknown

## Probabilistic Models with Unknown Objects



Brian Milch, *et al.*

# Reasoning about ...and Expecting the Unknown

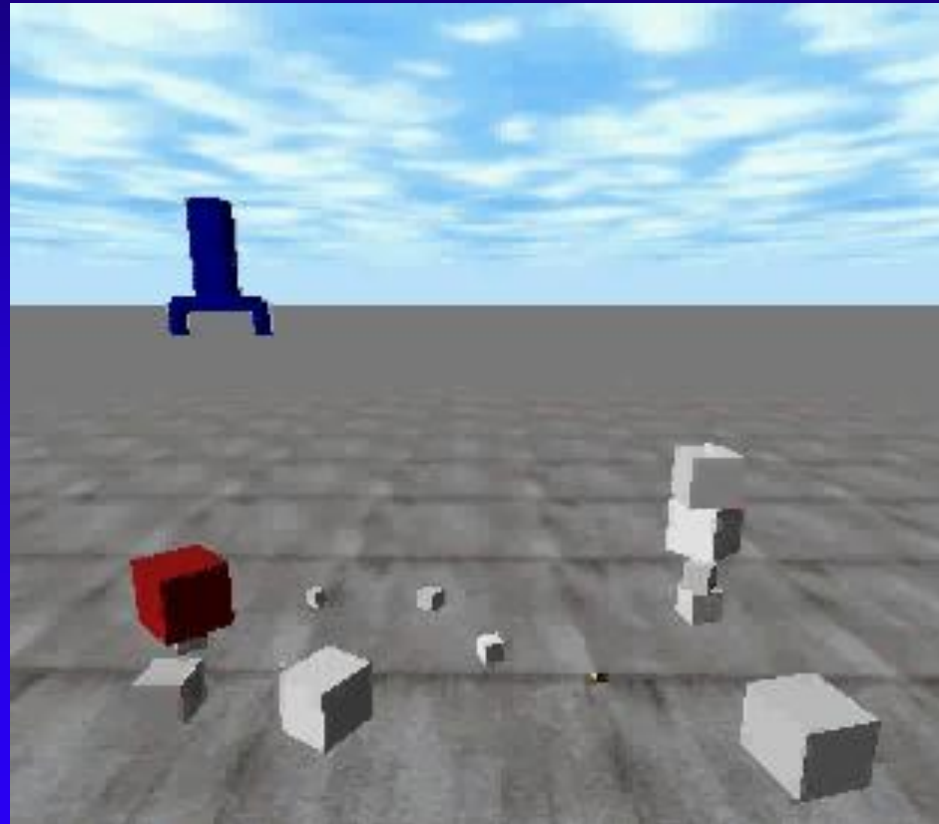


Krumm, H., et al.



# Extending Proficiency in a Messy World

- ◆ Assume *unmodeled* objects, relations, noise
- ◆ Continue to extend models with experiences

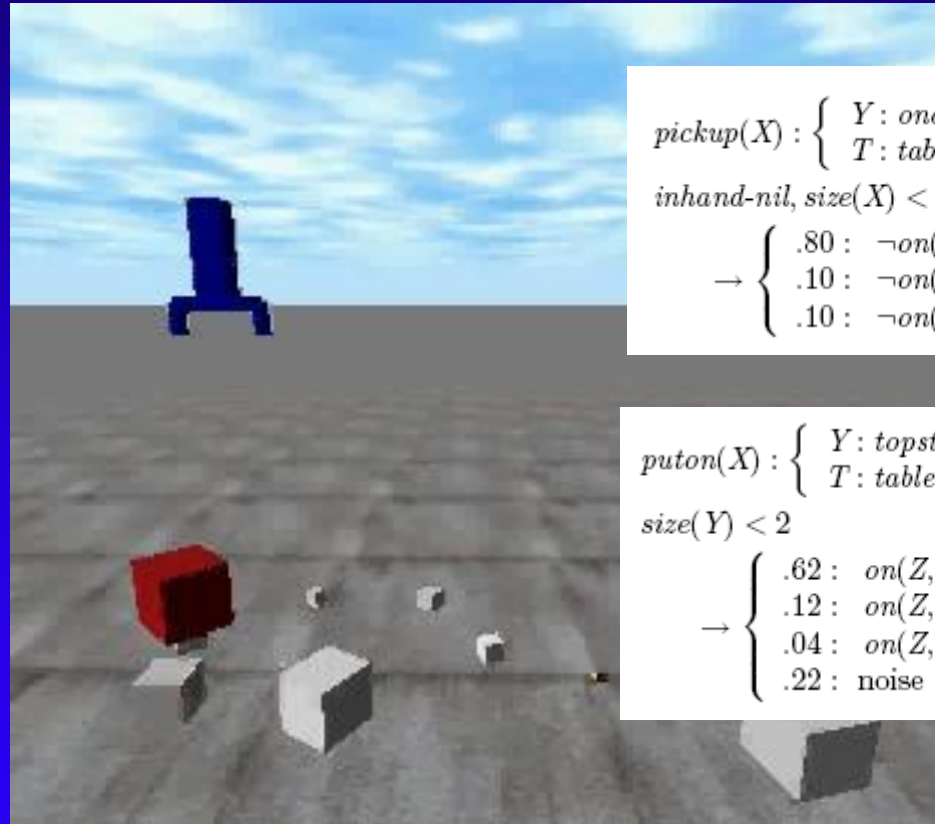


(video)

Pasula, Zettlemoyer, Kaelbling

# Extending Proficiency in an Open World

- ◆ Assume *unmodeled* objects, relations, noise
- ◆ Continue to extend models with experiences



$$\text{pickup}(X) : \left\{ \begin{array}{l} Y : \text{onclear}(X, Y), Z : \text{on}(Y, Z), \\ T : \text{table}(T) \end{array} \right\}$$

$$\text{inhand-nil}, \text{size}(X) < 2$$

$$\rightarrow \left\{ \begin{array}{l} .80 : \neg \text{on}(Y, Z) \\ .10 : \neg \text{on}(X, Y) \\ .10 : \neg \text{on}(X, Y), \text{on}(Y, T), \neg \text{on}(Y, Z) \end{array} \right.$$

$$\text{puton}(X) : \left\{ \begin{array}{l} Y : \text{topstack}(Y, X), Z : \text{inhand}(Z), \\ T : \text{table}(T) \end{array} \right\}$$

$$\text{size}(Y) < 2$$

$$\rightarrow \left\{ \begin{array}{l} .62 : \text{on}(Z, Y) \\ .12 : \text{on}(Z, T) \\ .04 : \text{on}(Z, T), \text{on}(Y, T), \neg \text{on}(Y, X) \\ .22 : \text{noise} \end{array} \right.$$




# Extending Perceptual Proficiency with Shape Prototypes and Deformations



G. Elidan, G. Heitz, and D. Koller

# Extending Perceptual Proficiency with Shape Prototypes and Deformations

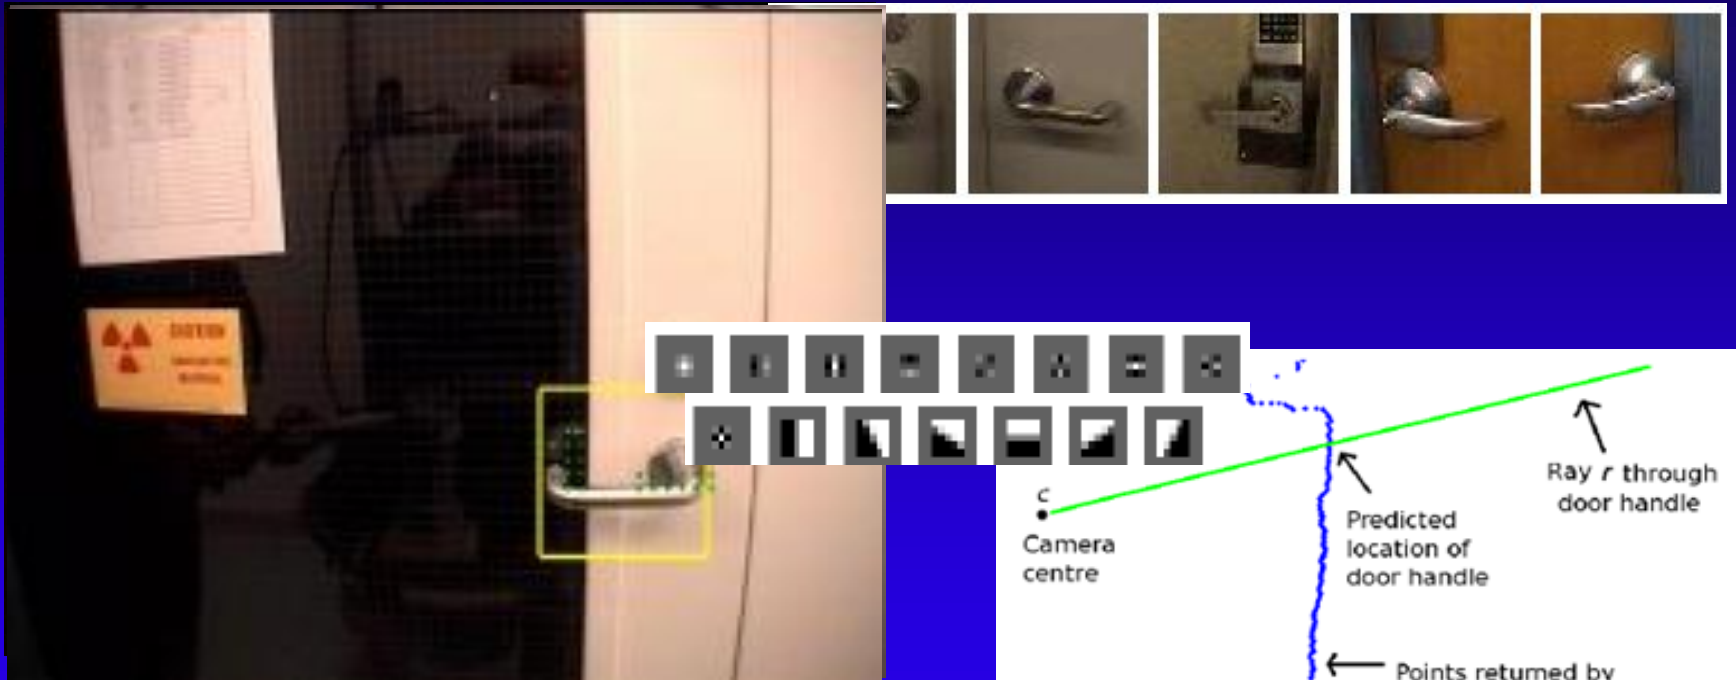


G. Elidan, G. Heitz, and D. Koller



# Proficiency via Transfer in Open Worlds

- ◆ Transferring prior knowledge to new situations



Klingbeil, Saxena, Ng, *et al.*

# Proficiency via Transfer in Open Worlds

- ◆ Transferring prior knowledge to new situations



(video)

Klingbeil, Saxena, Ng, *et al.*

# Prospering in the Open World

*To know that you do not know is the best.*

Lao Tzu

# Prospering in the Open World

*To know that you do not know is the best.*

Lao Tzu

- ◆ **Modeling model competencies, limitations, extensions**
- ◆ **Context-sensitive failures & successes**
- ◆ **Models of anomaly & surprise**
- ◆ **Value of prototypes, analogy, transfer**
- ◆ **Learning objects, predicates, preferences, goals in noisy environments**


# Value of Open-World Challenges

- ◆ **AAAI / CVPR Semantic Robot Vision Challenge**
  - Robots must perform a scavenger hunt in a *previously-unknown* indoor environment.



-e.g.,

- scientific calculator
- Ritter Sport Marzipan
- DVD "Shrek"
- DVD "Gladiator"
- CD "Hey Eugene"
- electric iron



Web access

# Value of Open-World Challenges

## 2004 DARPA Challenge: Sandstorm

- Closed world vs open world model
  - Model of anomaly, surprise
  - Thinking out of the box?



# Where there's Smoke...



(video)

# Stepping into the Open World

- ◆ Key technical challenges
- ◆ AI moving into the world
- ◆ AI research community



# **AI Moving into the World: Trends & Directions**

- ◆ **Robust services in dynamic settings**
- ◆ **Human-computer collaboration**
- ◆ **Integrative intelligence**
- ◆ **Sciences**
- ◆ **Harnessing web**

# Robust Services in Dynamic Settings

- ◆ Context-sensitive competencies & policies

Example: Traffic prediction & routing



# Robust Services in Dynamic Settings


- ◆ Context-sensitive competencies & policies

Example: Traffic prediction & routing









P u g e t  
S o u n d

Seattle-Bainbridge Island Ferry  
Seattle-Bremerton Ferry

# ClearFlow

- ◆ **Fielded in 72 cities in North America (April 2008)**
- ◆ **Roads speeds assigned to ~60 million streets across North America *every few minutes***

Pittsburgh	Baltimore	Denver	Oklahoma City
Louisville	Jacksonville	Harrisburg	Lincoln
Philadelphia	Boston	Cleveland	Phoenix
Norfolk	Greensboro	Wilkes-Barre	Omaha
Dallas/Ft. Worth	Washington, D.C.	Portland	Providence
Columbus	Nashville	Buffalo	Little Rock
Detroit	Miami	Sacramento	Salt Lake City
Charlotte	West Palm Beach	Dayton	Mobile
Houston	Tampa	Seattle	Cincinnati
Hartford	New Orleans	Fresno	Portsmouth-Manchester
Los Angeles	Orlando	Atlanta	San Antonio
Raleigh-Durham	Tucson	Grand Rapids	Rochester
New York	San Diego	Birmingham	Milwaukee
Richmond	Albuquerque	Toledo	Syracuse
San Francisco	Minneapolis	Indianapolis	Austin
Tulsa	Colorado Springs	Greenville	Spokane
Chicago	St. Louis	Las Vegas	Kansas City
Albany	Allentown	Memphis	Toronto




# Chicago



Default

# Chicago

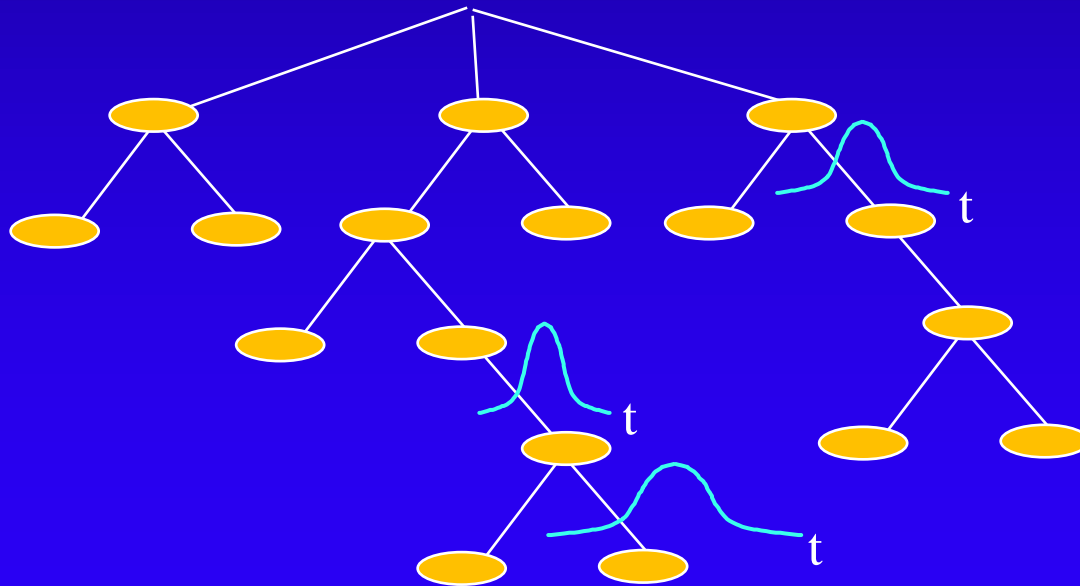


Clearflow



# Planning while the Sands are Shifting

- Lab: Temporal models forecast future speeds and uncertainties
- Path planning with changing situation
- Variances, robustness, flexibility  
... finding paths with contingencies




# Human-Computer Collaboration in Open Worlds



# Grounding

Converging on shared references,  
beliefs, and intentions




# Grounding

Converging on shared references,  
beliefs, and intentions



# Grounding

Converging on shared references,  
beliefs, and intentions




# Grounding

Converging on shared references, beliefs, and intentions

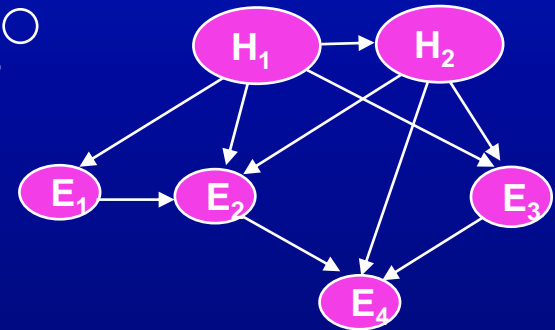





# Grounding

Converging on shared references, beliefs, and intentions




# Grounding on Beliefs






# Grounding on Beliefs




# Grounding on Beliefs



- ◆ Debiassing human judgment
- ◆ Tutoring, advising, education
- ◆ Human expectation & surprise
- ◆ Ideal display & alerting


# Mixed-Initiative Collaboration

- ◆ Interleaving of contributions from machine and human to jointly solve problems.
  - Problem recognition, decomposition, coordination



# Mixed-Initiative Collaboration

- ◆ Interleaving of contributions from machine and human to jointly solve problems.
  - Problem recognition, decomposition, coordination



# Mixed-Initiative Collaboration

- ◆ Interleaving of contributions from machine and human to jointly solve problems.

The image illustrates a mixed-initiative collaboration between a human and a genie. The human, Robert Croft, sends an email to Eric Horvitz (the genie) about planning a session. The genie responds with a question: "Should I go ahead and schedule that?". The human then sends an appointment request for Friday at 3 PM. The genie responds with a suggestion: "You will be busy then with the Lumiere project meeting... How about trying Friday at 3 PM.". The human then sends an appointment request for Friday at 3 PM. The genie responds with a reminder and a note about a conflict with the Lumiere project meeting.

Planning for session - Message (Plain T... - [X]

File Edit View Insert Format Tools Actions Help

You replied on Tuesday, September 22, 1998 9:41 PM.

From: Robert Croft Sent: Tue 9/22/98 8:25 PM  
To: Eric Horvitz  
Cc:  
Subject: Planning for session

I'd like to catch up on plans for the conf. session on decision making. How about speaking by phone sometime around 2pm on Fri? Looking forward to chatting.  
-Robert

Should I go ahead and schedule that?

You replied on Tuesday, September 22, 1998 9:41 PM.

From: Robert Croft Sent: Tue 9/22/98 8:25 PM  
To: Eric Horvitz  
Cc:  
Subject: Planning for session

I'd like to catch up on plans for the conf. session on decision making. How about speaking by phone sometime around 2pm on Fri? Looking forward to chatting.  
-Robert

-- Genie is listening -- for commands.

You will be busy then with the Lumiere project meeting... How about trying Friday at 3 PM.

Planning for session - Appointment - [X]

File Edit View Insert Format Tools Actions Help

Appointment Attendee Availability Online

Adjacent to another appointment on your Calendar.

Subject: Planning for session

Location: [v] This is an online meeting

Start time: Fri 9/25/98 3:00 PM All day event

End time: Fri 9/25/98 4:00 PM

Reminder: [v] 15 minutes Show time as: Busy

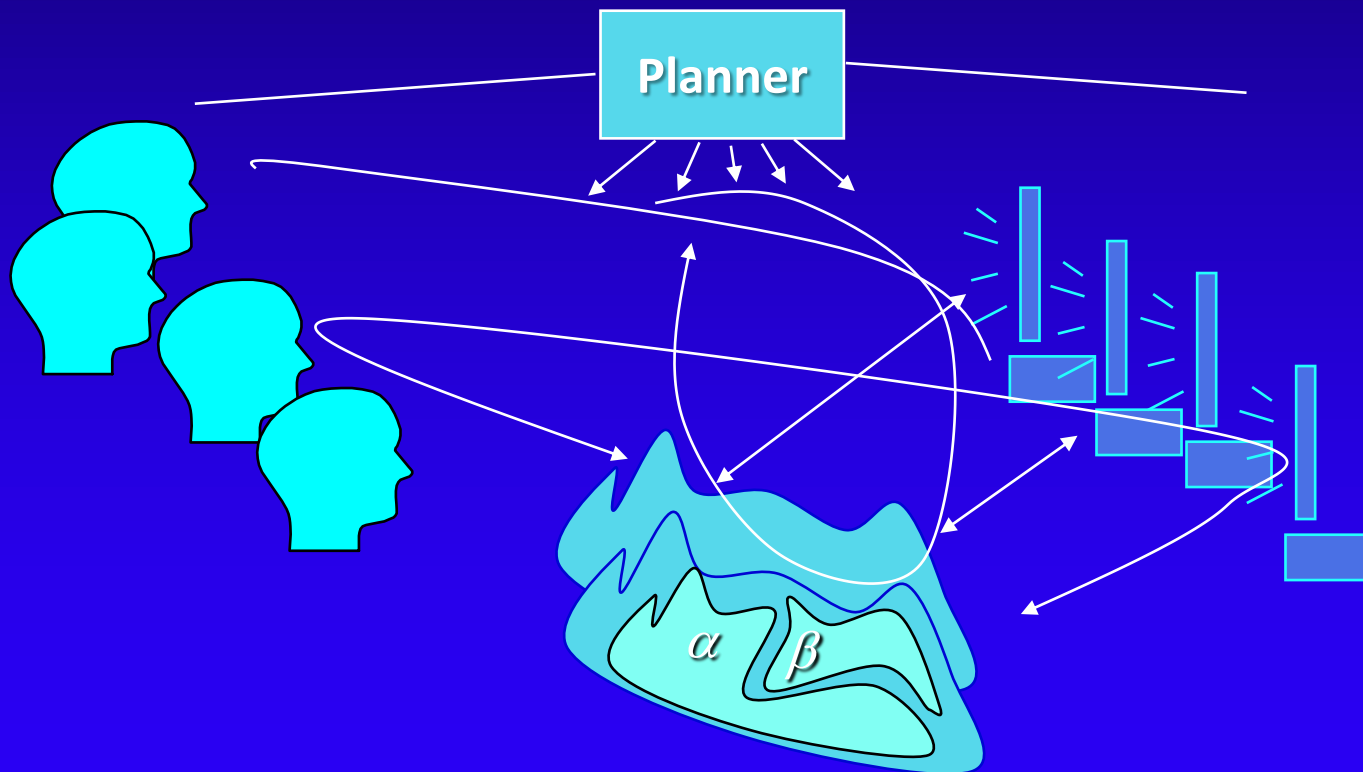
(Appointment moved from 9/25/98 2:00 PM due to conflict with the Lumiere project meeting)

Eric,  
I'd like to catch up on plans for the conf. session on decision making. How about speaking by phone sometime around 2pm on Fri? Looking forward to chatting.  
-Robert

Categories... Private

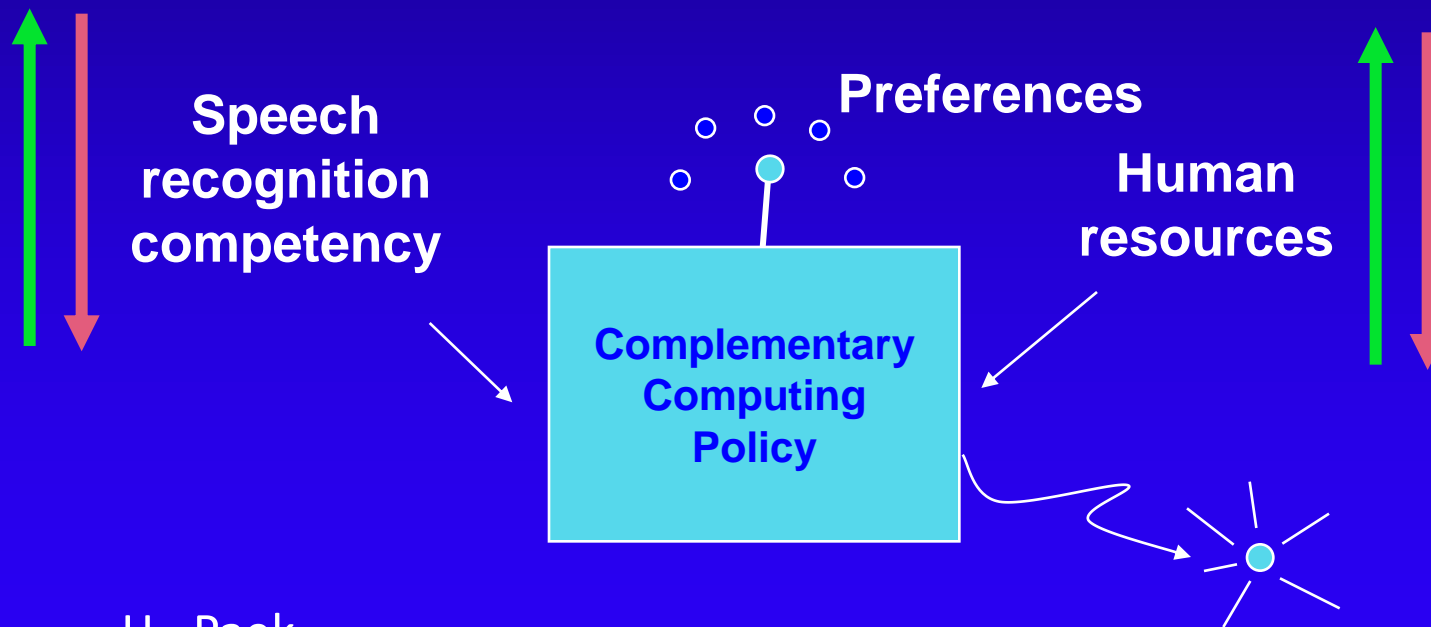
# Complementary Computing

- ◆ Consider systems of people & computation
- ◆ Policies for coordinating contributions
- ◆ *Task markets* for humans, computers, sensors, effectors



# Example: Adaptive Policies in an Automated Reception System

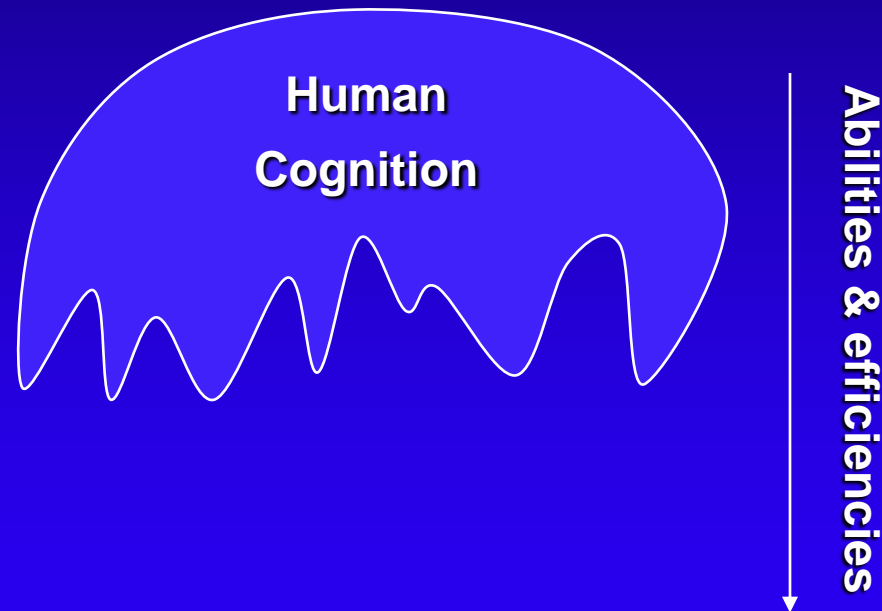
- ◆ Changing dialog competencies
- ◆ Changing load on staff



# Learning about Human Cognition for Augmenting Human Abilities

20<sup>th</sup> Century cognitive psychology:

Characterizable limitations & bottlenecks






# Learning about Human Cognition for Augmenting Human Abilities

20<sup>th</sup> Century cognitive psychology:

Characterizable limitations & bottlenecks

*Promise of  
sensing &  
reasoning*

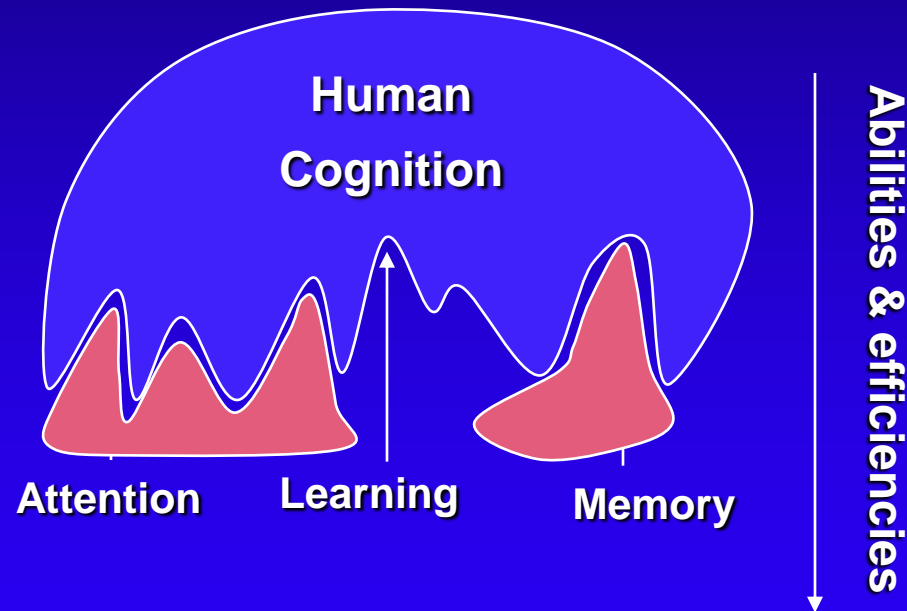


# Learning about Human Cognition for Augmenting Human Abilities

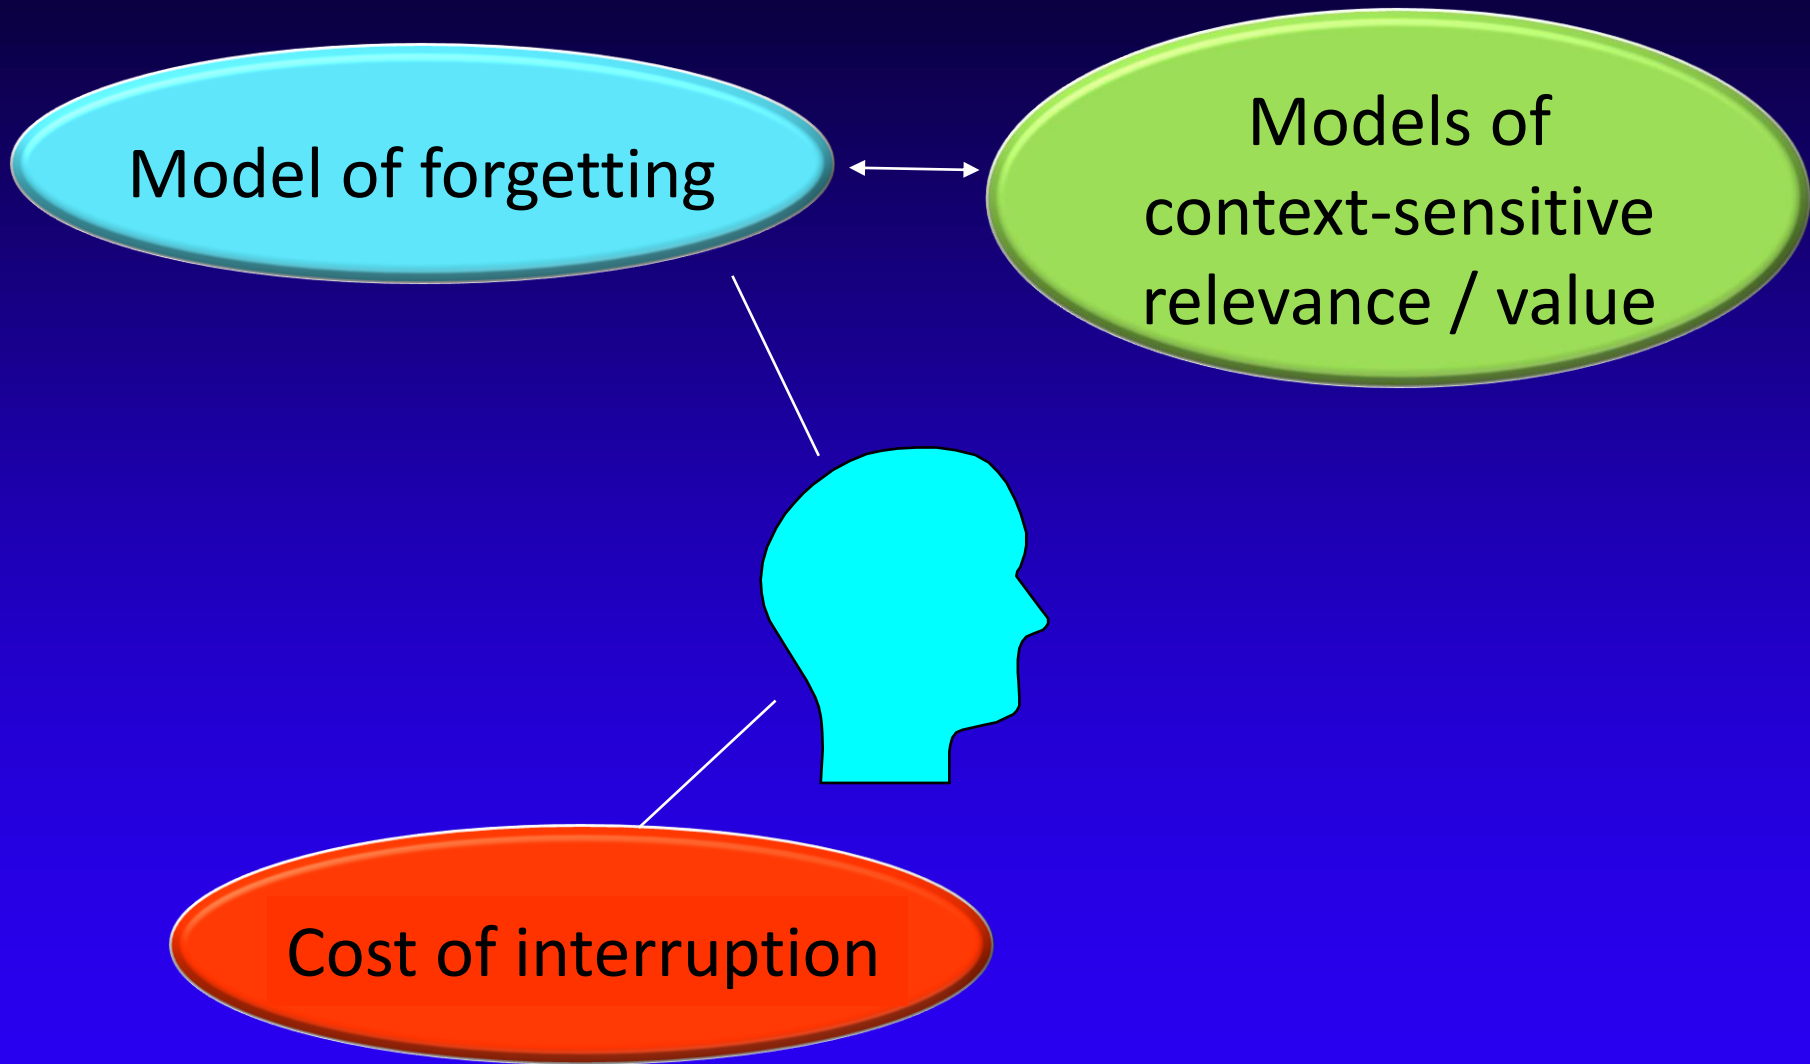
20<sup>th</sup> Century cognitive psychology:

Characterizable limitations & bottlenecks

*Promise of  
sensing &  
reasoning*

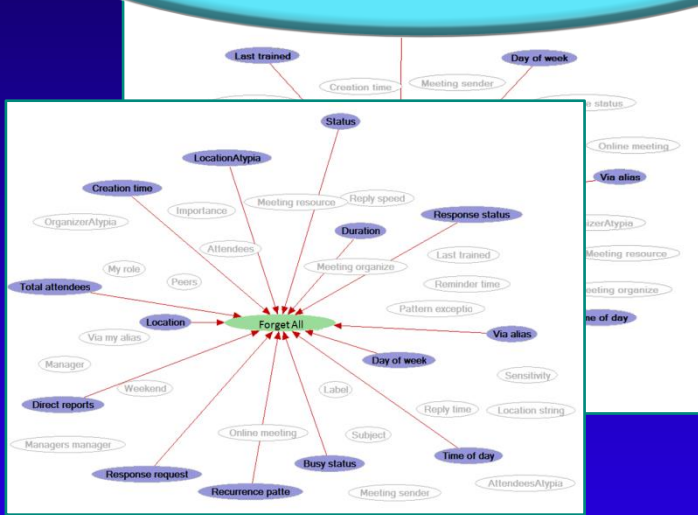


# Models of Memory in Reminding

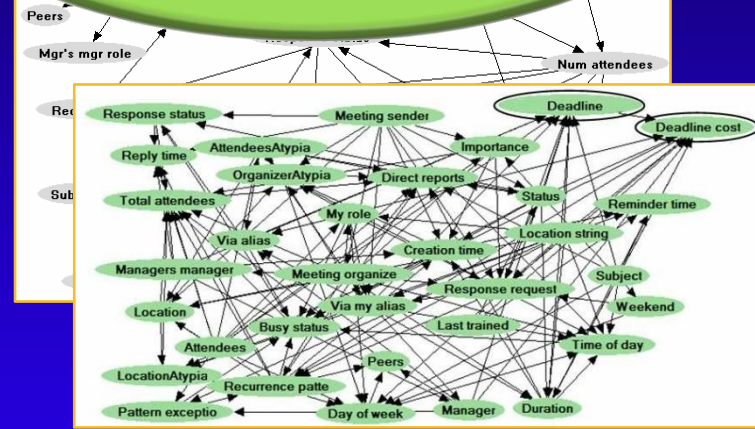


# Models of Memory in Reminding

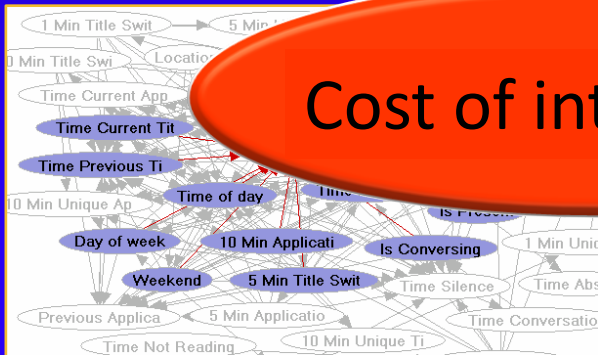
Model of forgetting



Models of context-sensitive relevance / value



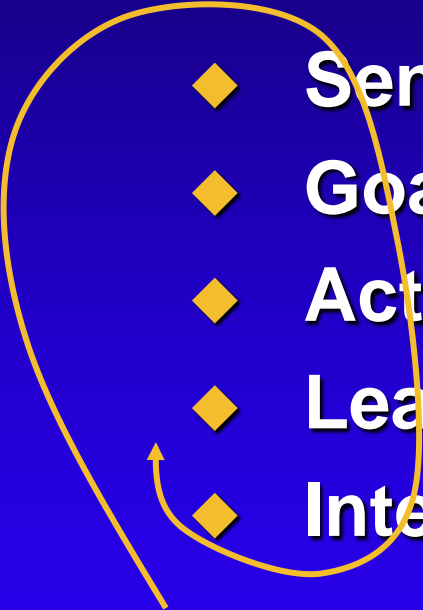
Cost of interruption



Reminders at ideal times

# Integrative Intelligence

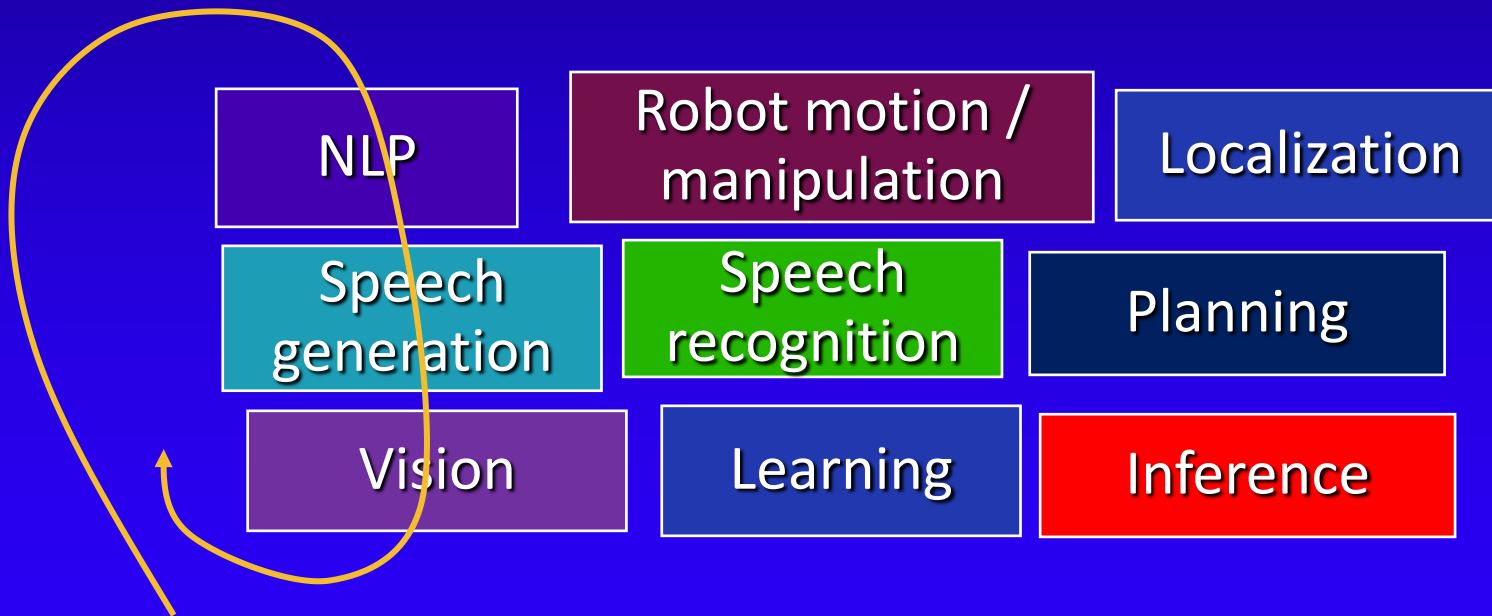
Opportunity: More comprehensive, “integrative intelligences,” probing open world

- ◆ Sensing & symbols
  - ◆ Goals & preferences
  - ◆ Action execution, monitoring
  - ◆ Learning, semi- and unsupervised
  - ◆ Interaction of components
- 

# Integrative Intelligence

Opportunity: More comprehensive, “integrative intelligences,” probing open world

- ◆ Weaving together components that have been developed separately



# Integrative Intelligence

- ◆ Vision, manipulation, navigation, learning

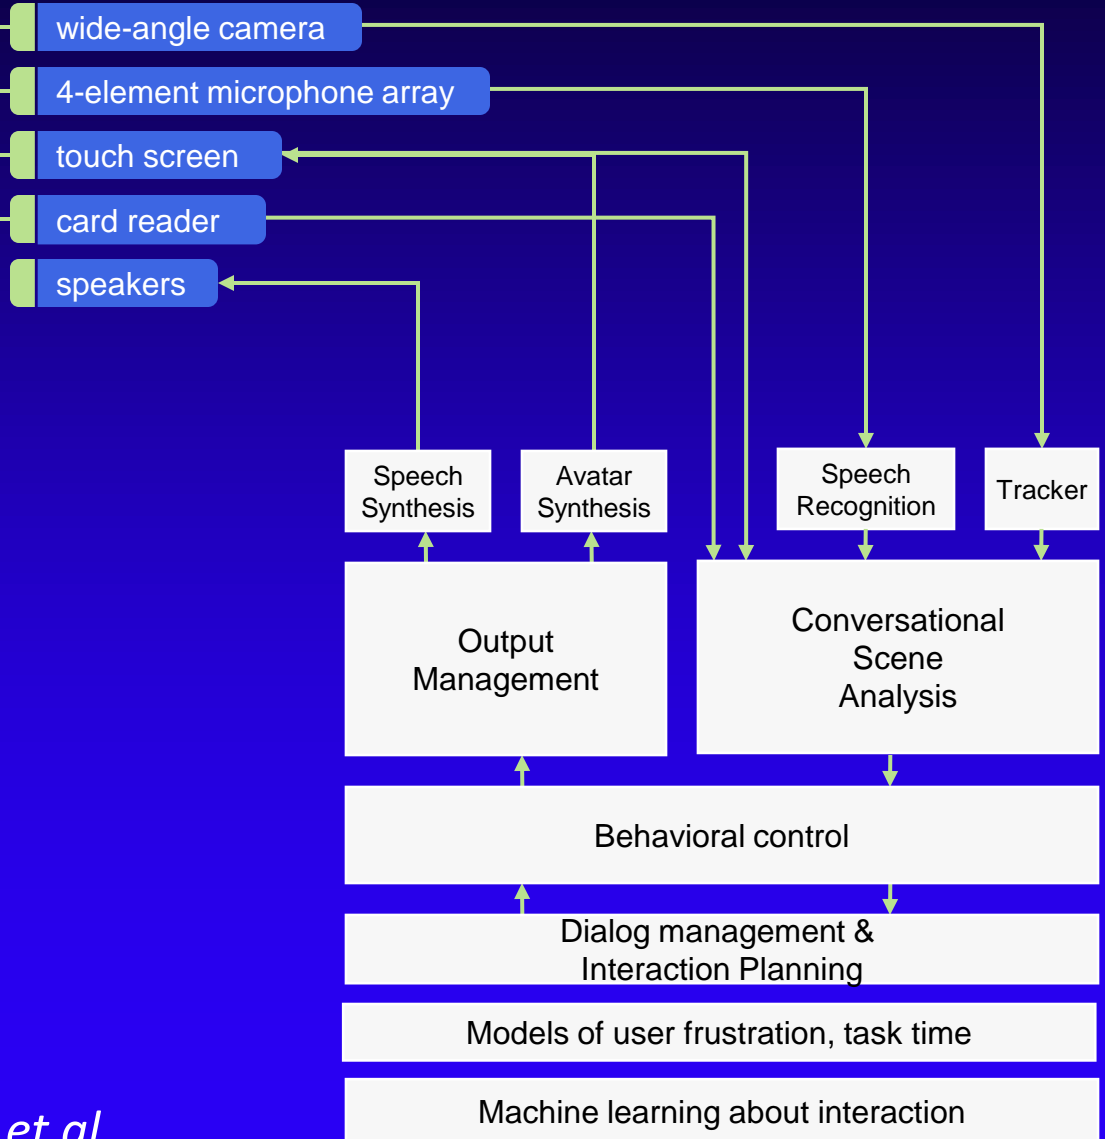


(video)

*STAIR, Ng, et al.*



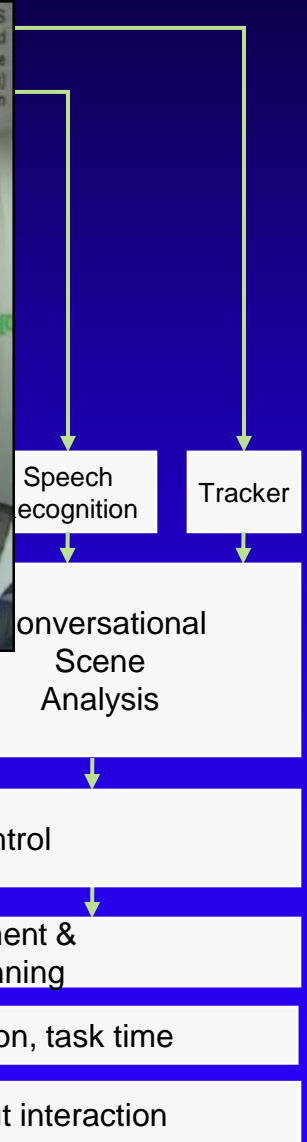
# Integrative Intelligence



*Receptionist: Bohus, H., et al.*



# Integrative Intelligence




*Receptionist: Bohus, H., et al.*


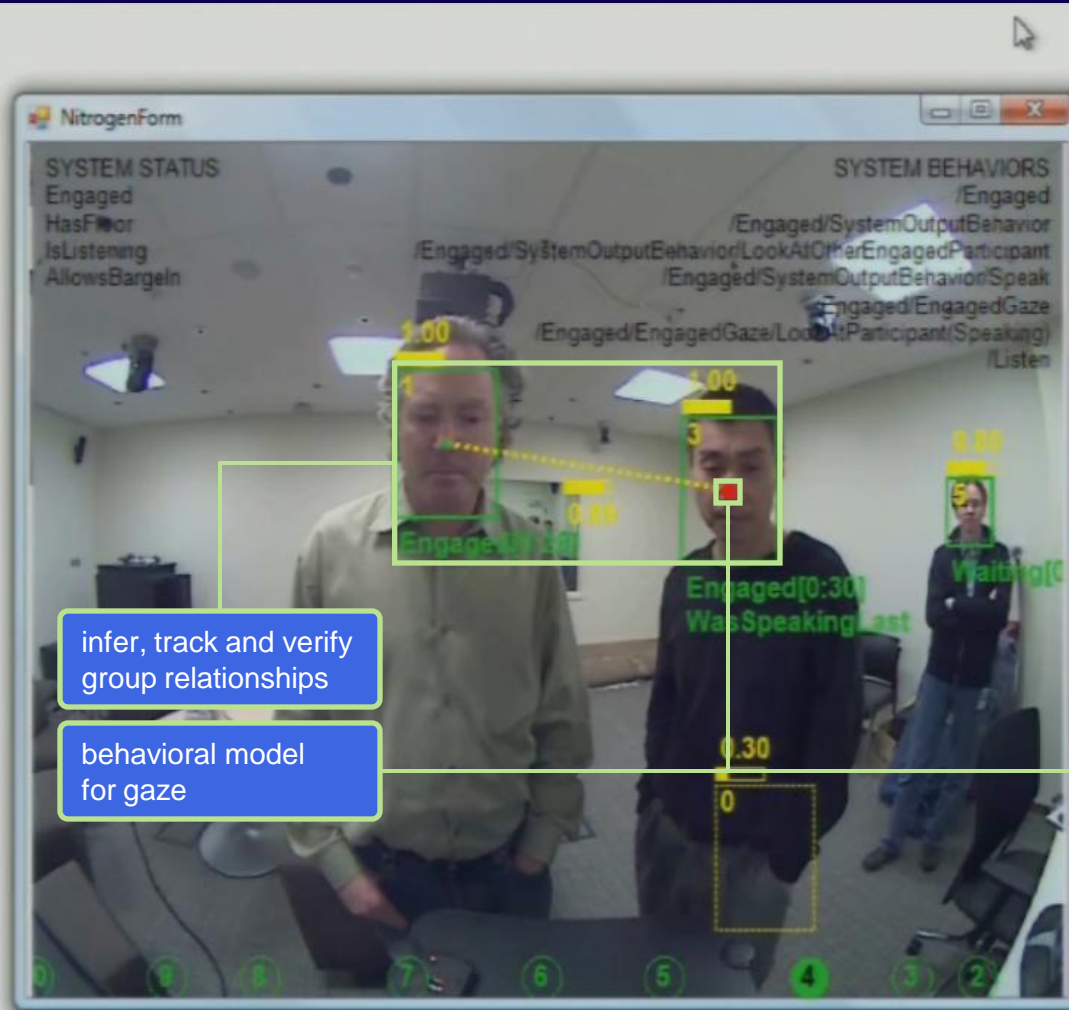
# Integrative Intelligence



*Receptionist: Bohus, H., et al.*

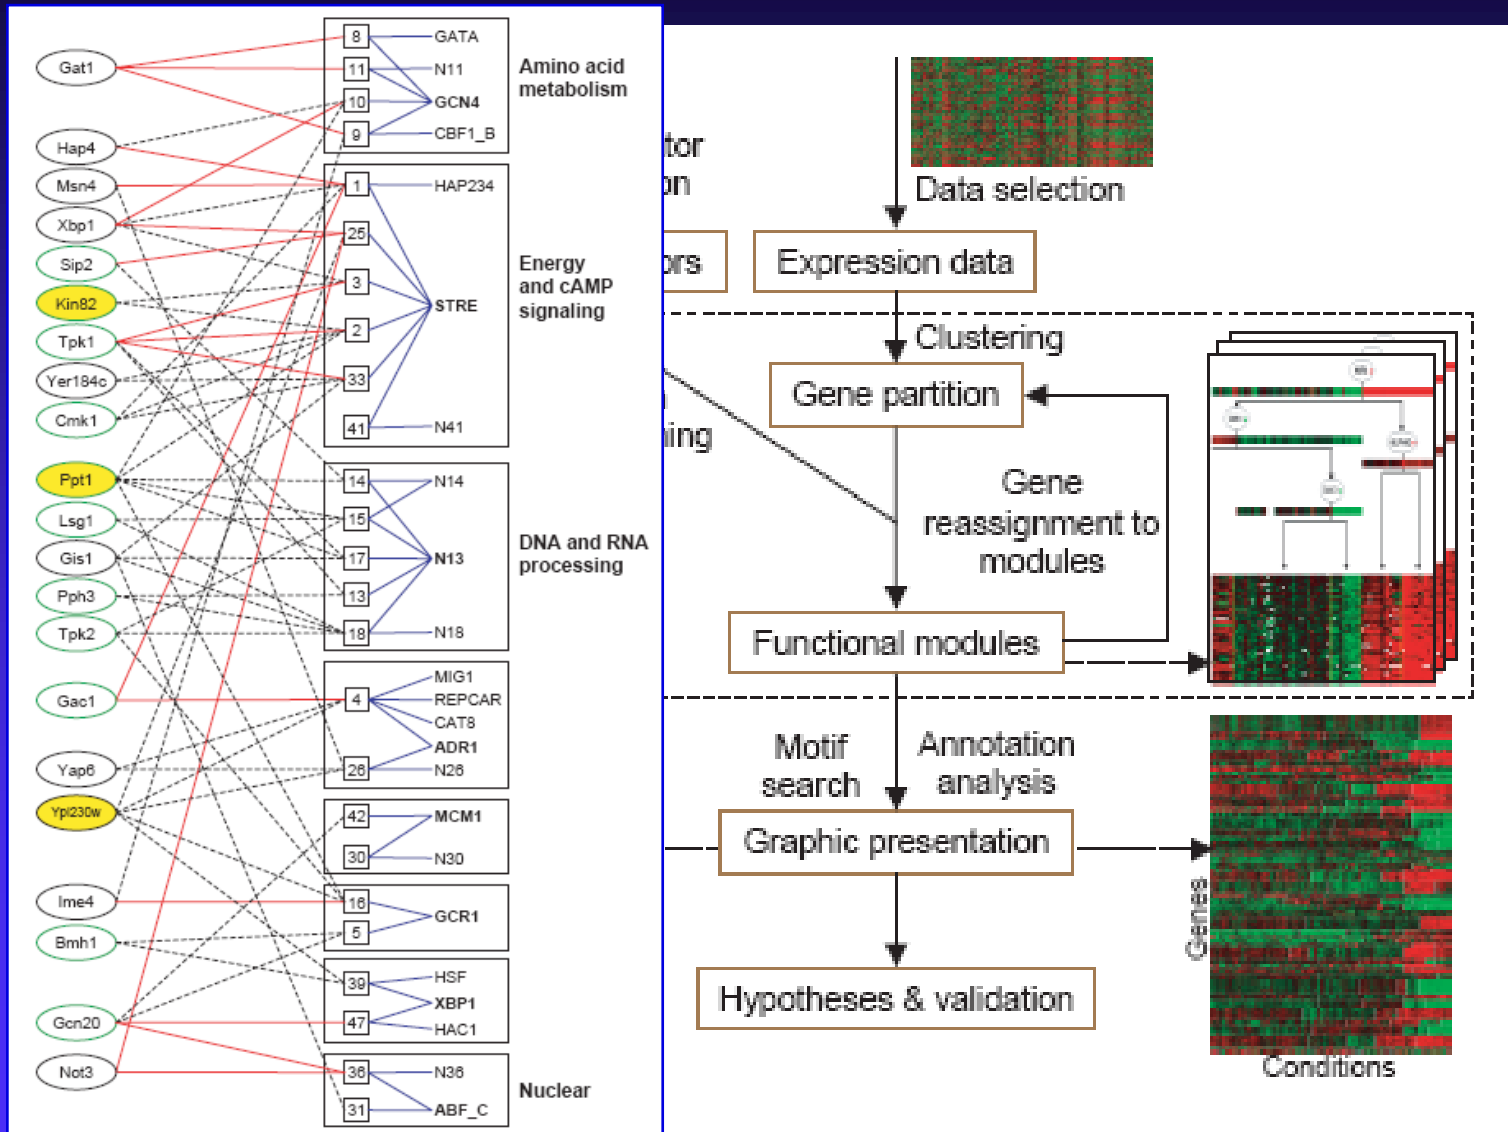


# Integrative Intelligence






# Learning about Structure & Function



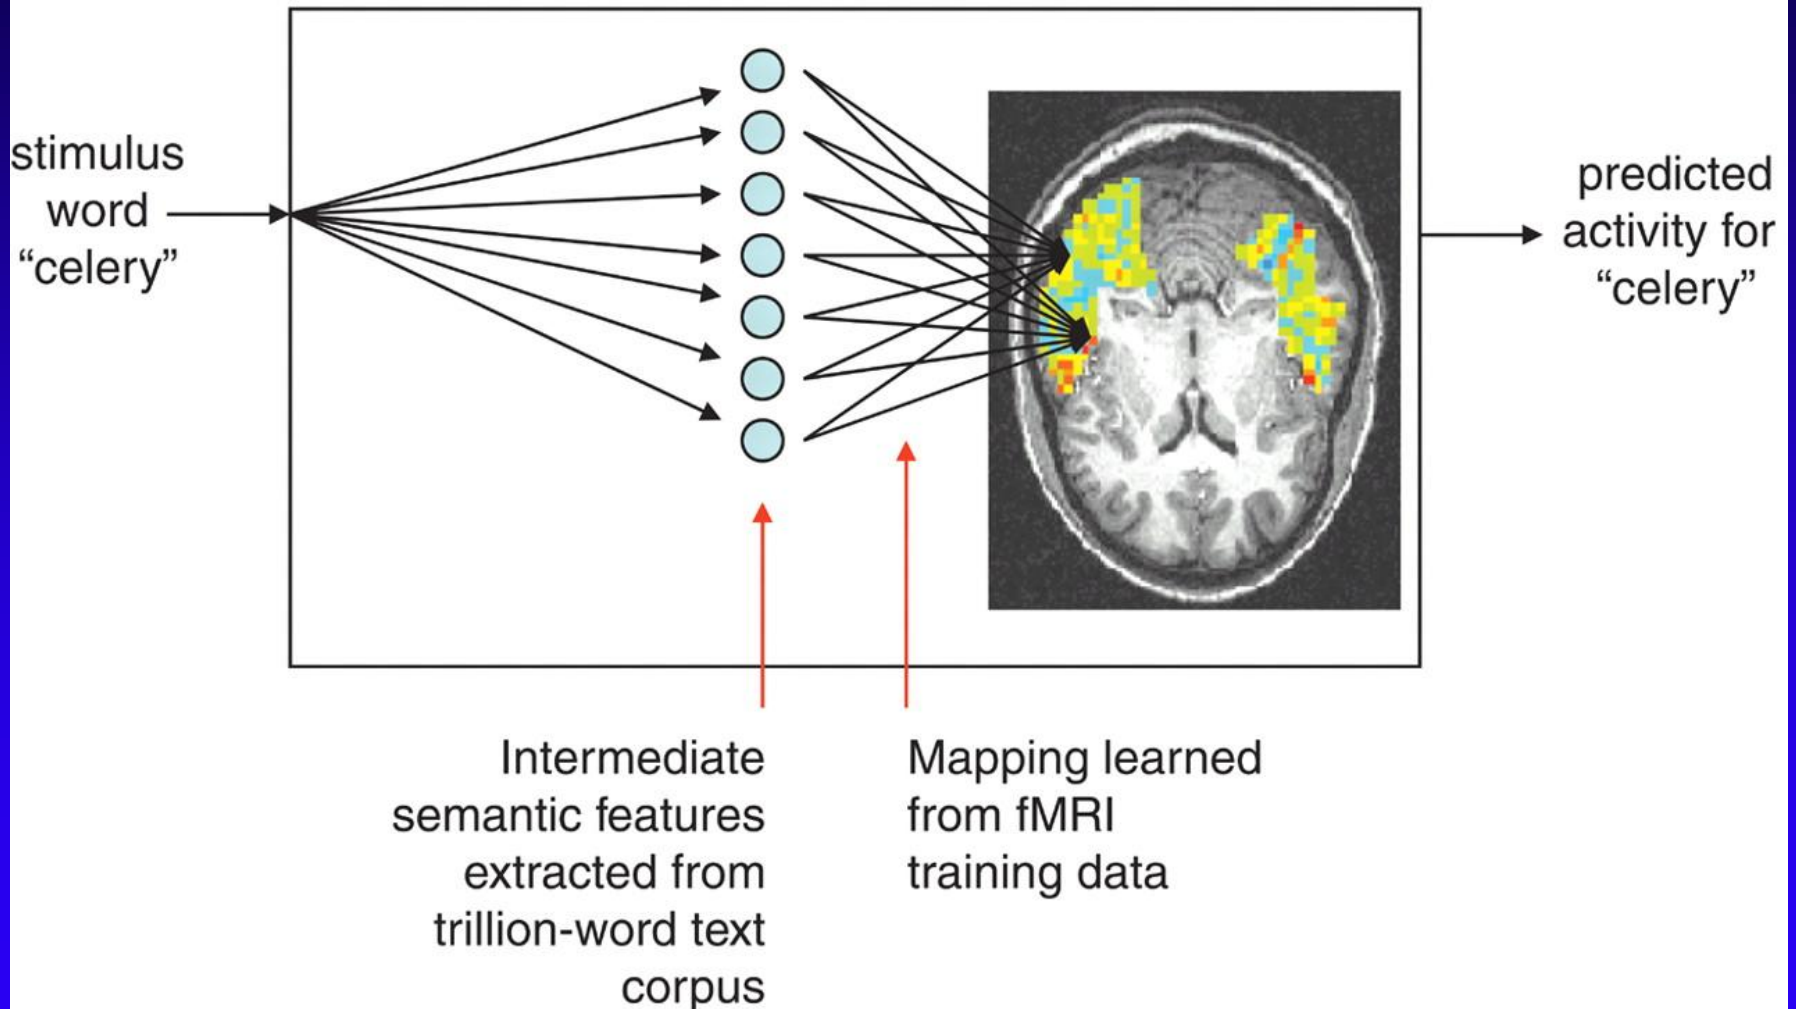


# Insights From & About Neurobiology

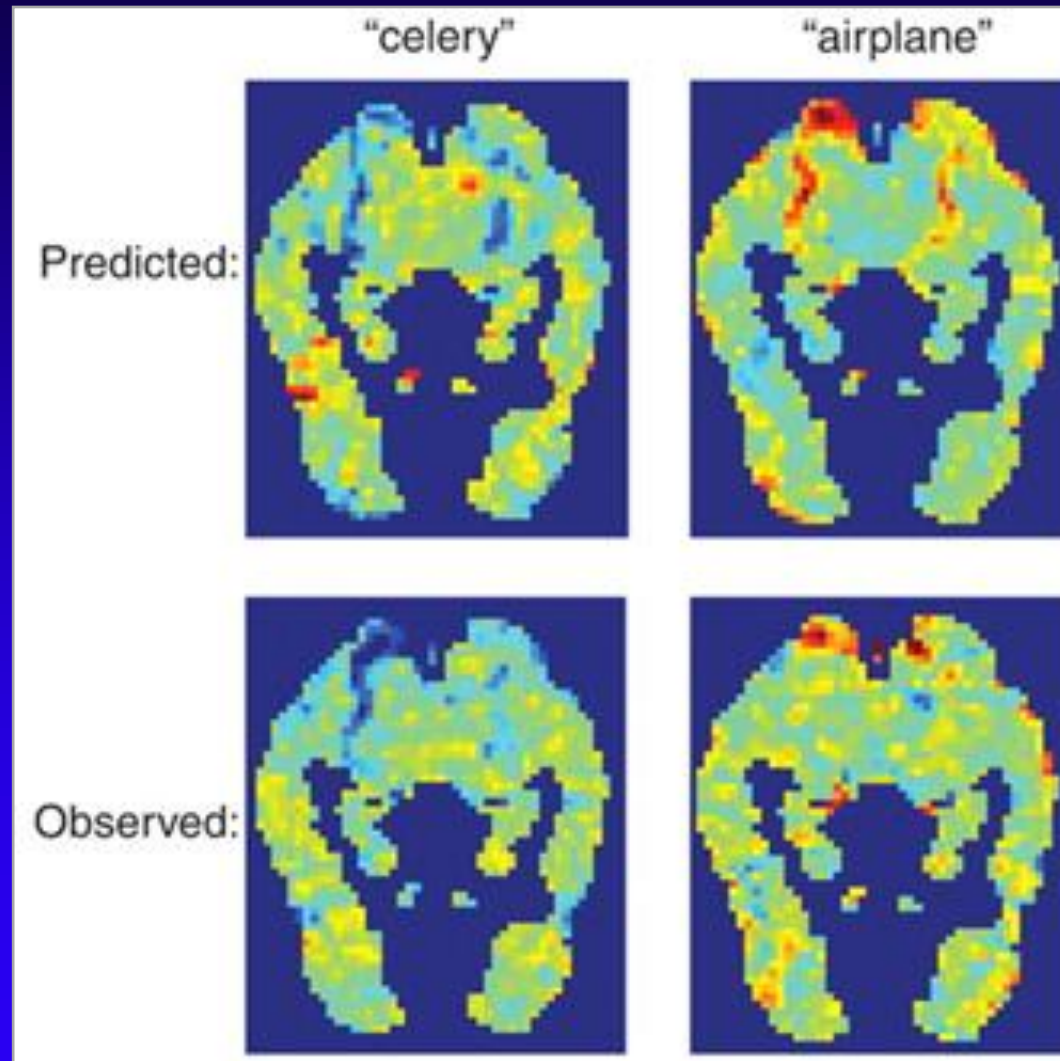


# A Study in Representation

## Predictive model




# A Study in Representation






# A Study in Representation



Mitchell, *et al.*

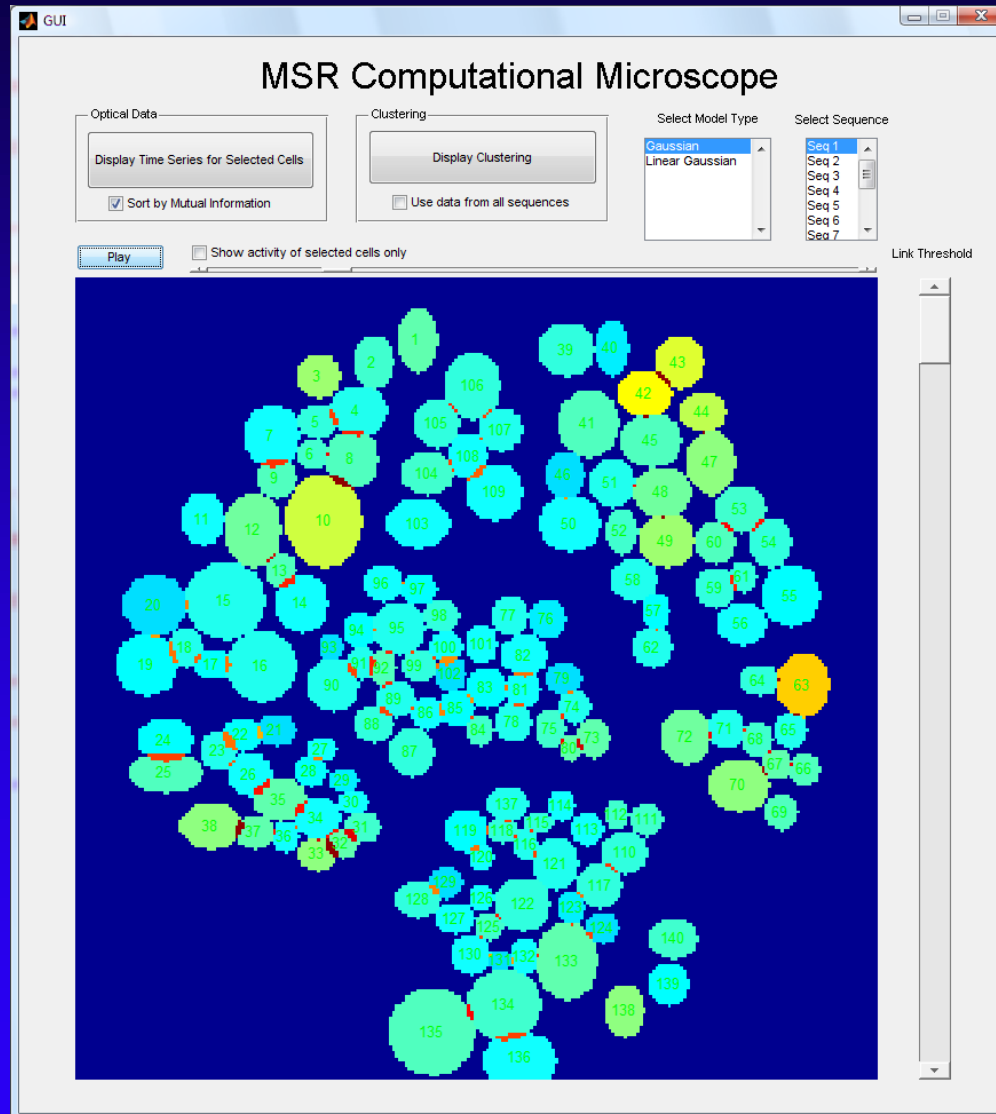
# Coming Era of Neuroinformatics



(video)

Reid, *et al.*


# Toward a Computational Microscope




(video)

Kapoor, H., *et al.*

# Learning to Harness the Web



# Learning to Harness the Web



*University of Washington*  
Department of Computer Science & Engineering

KnowItAll

▷ CSE Home

▷ About Us ▷ Search

## Demonstrations

[TextRunner](#)

[Opine](#)

## People

[Oren Etzioni](#)

[Stephen Soderland](#)

[Matt Broadhead](#)

[Michele Banko](#)

[Michael Cafarella](#)

[Doug Downey](#)

[Ana-Maria Popescu](#)

[Mausam](#)

[Marcus Sammer](#)

[Michael Schmitz](#)

[Alex Yates](#)

- **How can a computer accumulate a massive body of knowledge?**
- **What will Web search engines look like in ten years?**

To address these questions, the **KnowItAll** project has been developing a variety of domain systems that extract information from the Web in an autonomous, scalable manner.


The KnowItAll project has been sponsored in part by federal research grants from the [National Science Foundation](#) and the [Office of Naval Research](#).

## Demos:

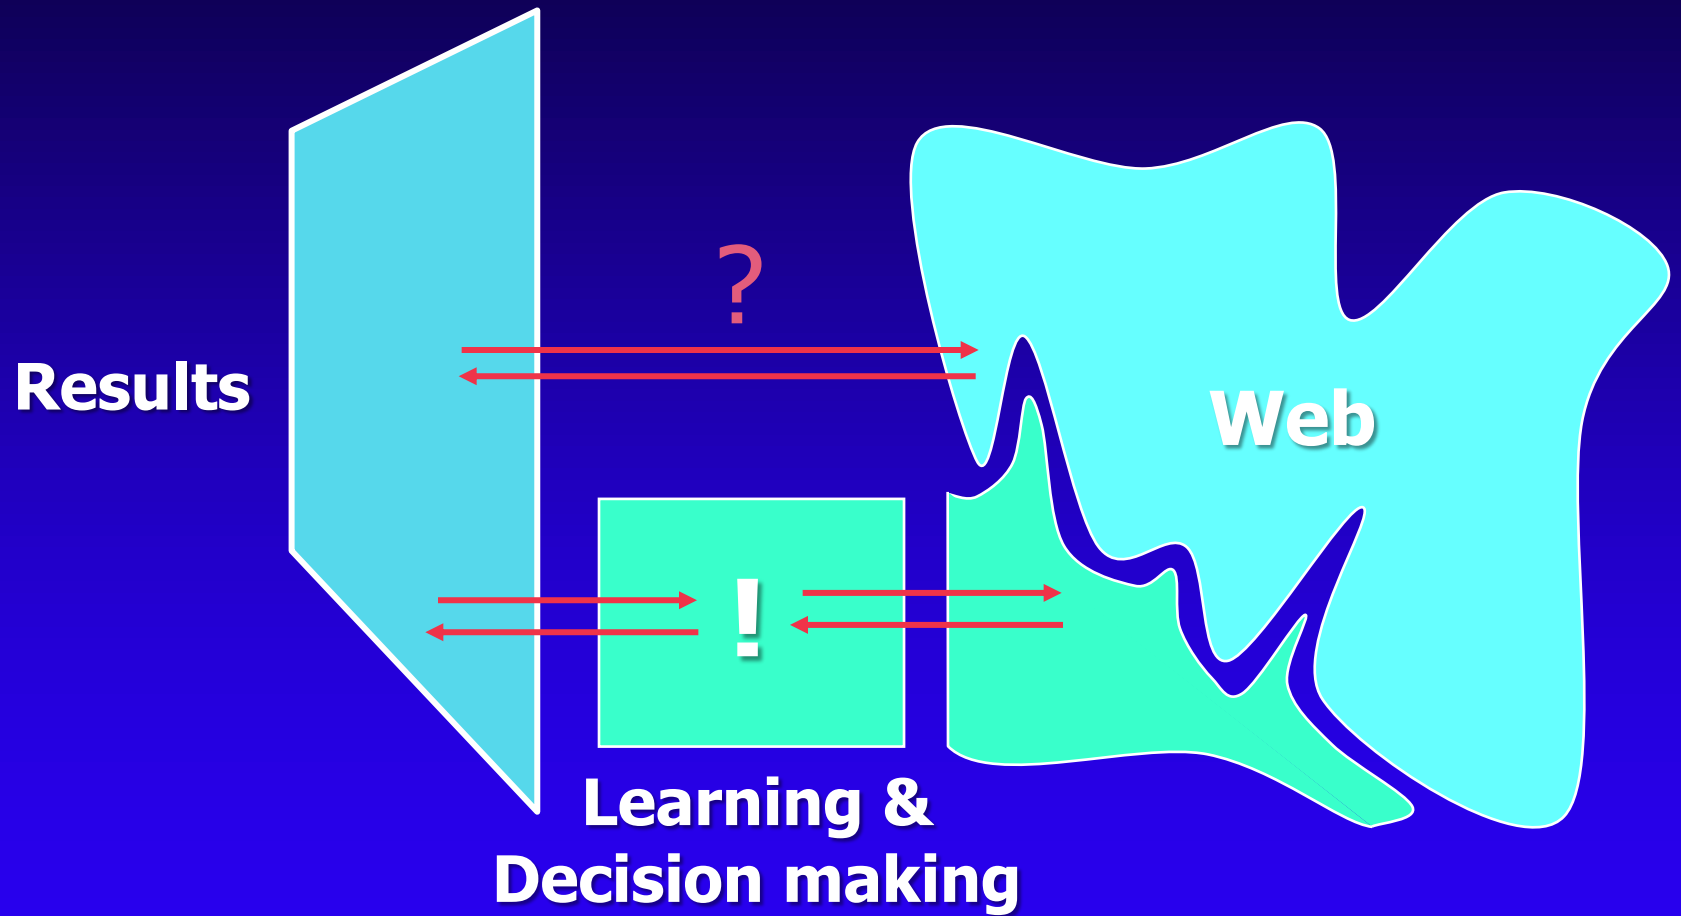
- [TextRunner](#): TextRunner searches 77,652,885 tuples extracted on the topic of nutrition and ranks the results by probability.
- [Opine](#): Opine is a review mining system which uses KnowItAll-type techniques in order to extract product features and customer opinions from product reviews.



# Learning to Harness the Web



# Learning to Harness the Web



# Learning to Harness the Web

## AskMSR: Automatic Question Answering

Question

Submit Query

Possible Answers for "*Who is Bill Gates married to?*":

**Melinda French (53%)**

**Microsoft Corporation (16%)**

**two children (8%)**

**French They (7%)**

**Hill (7%)**

**wife (7%)**



# Learning to Harness the Web

## AskMSR: Automatic Question Answering

Question

Submit Query

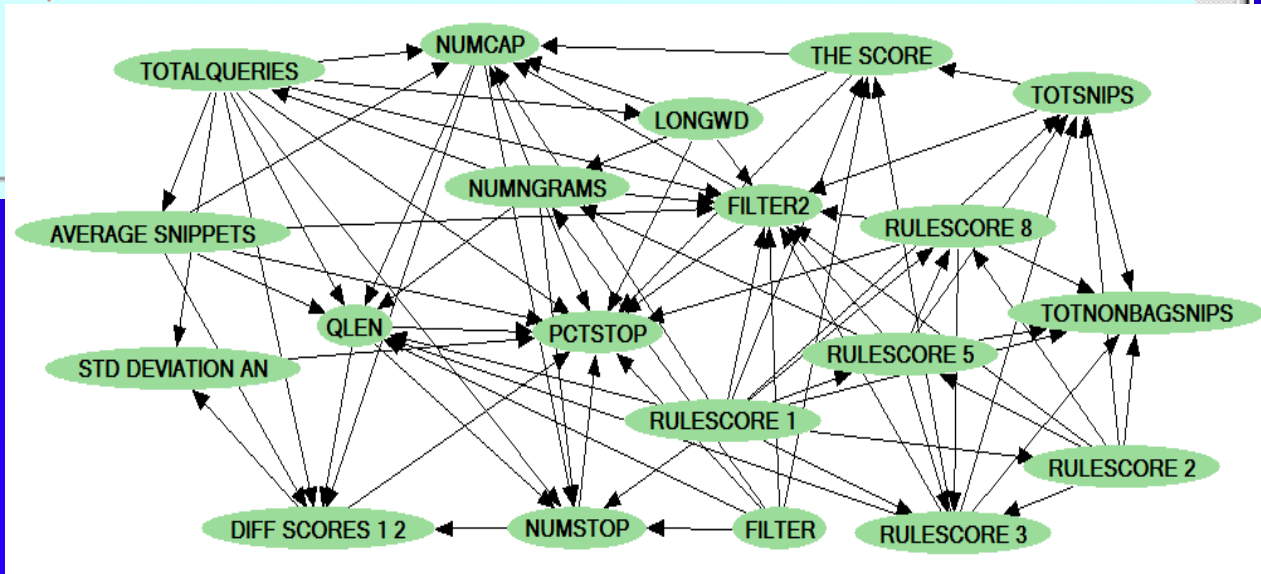
Possible Answers for "Who is Bill Gates married to?":

**Melinda French (53%)**

**Microsoft Corporation (16%)**

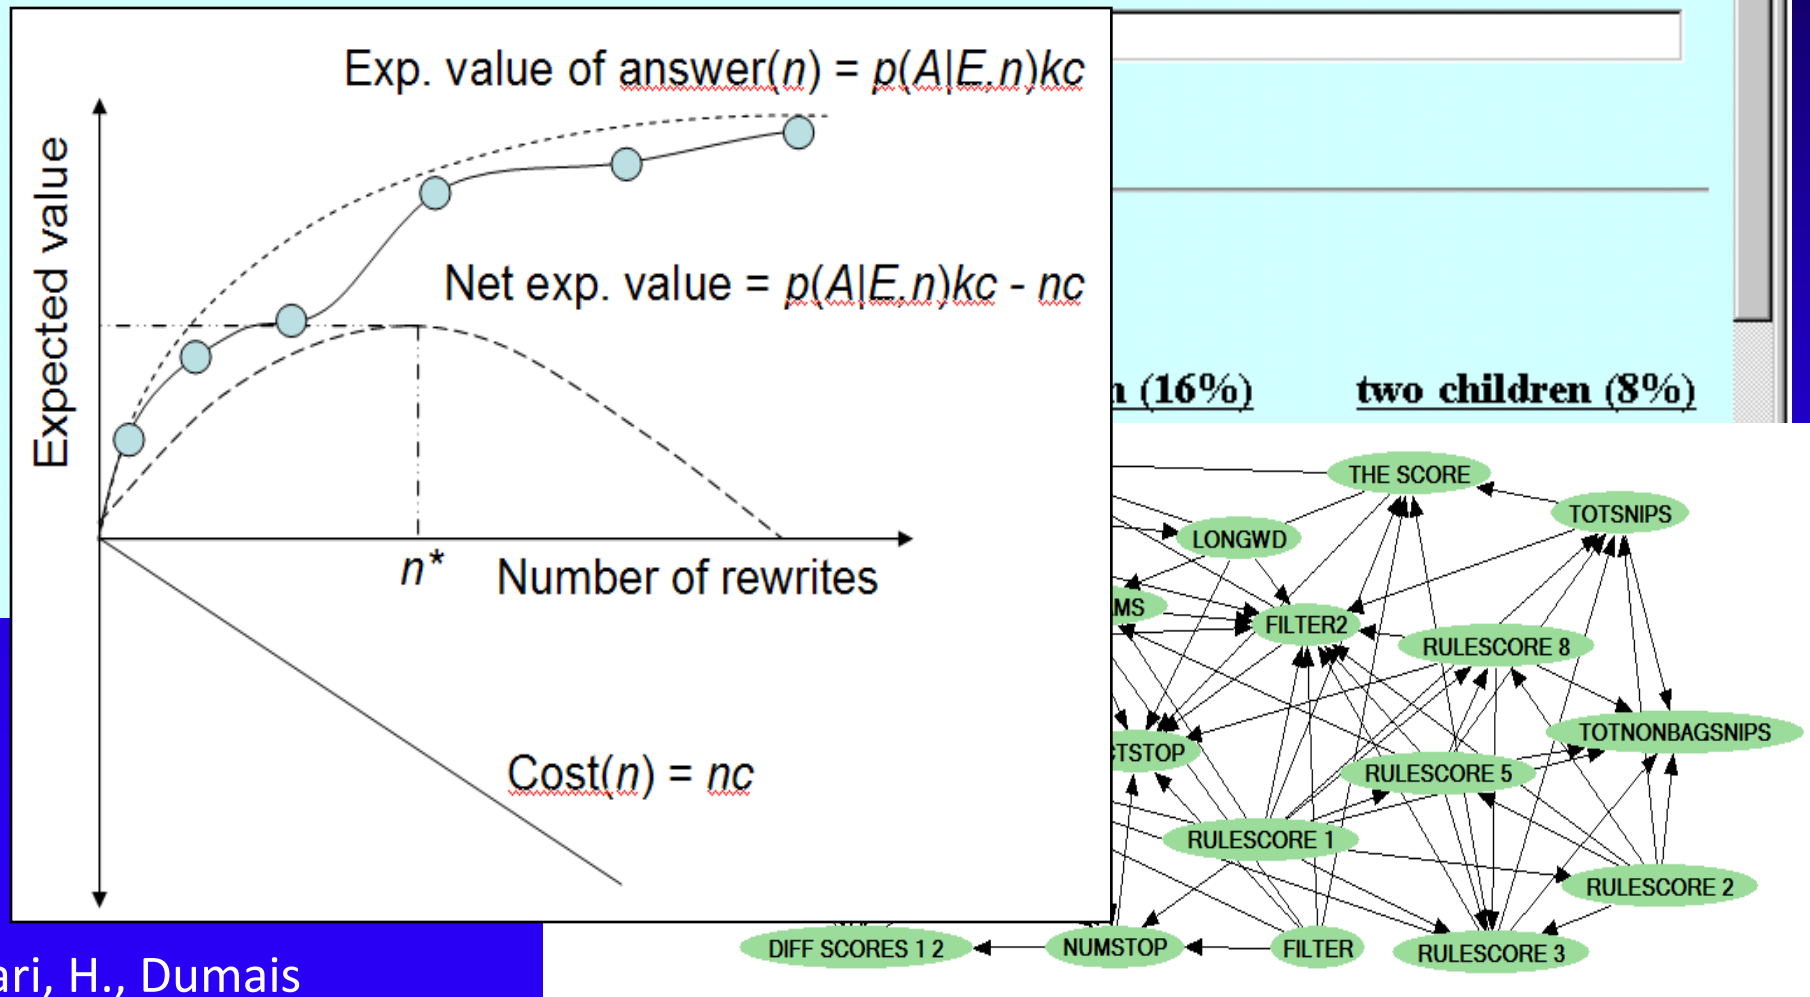
**two children (8%)**

**French They (7%)**



# Learning to Harness the Web

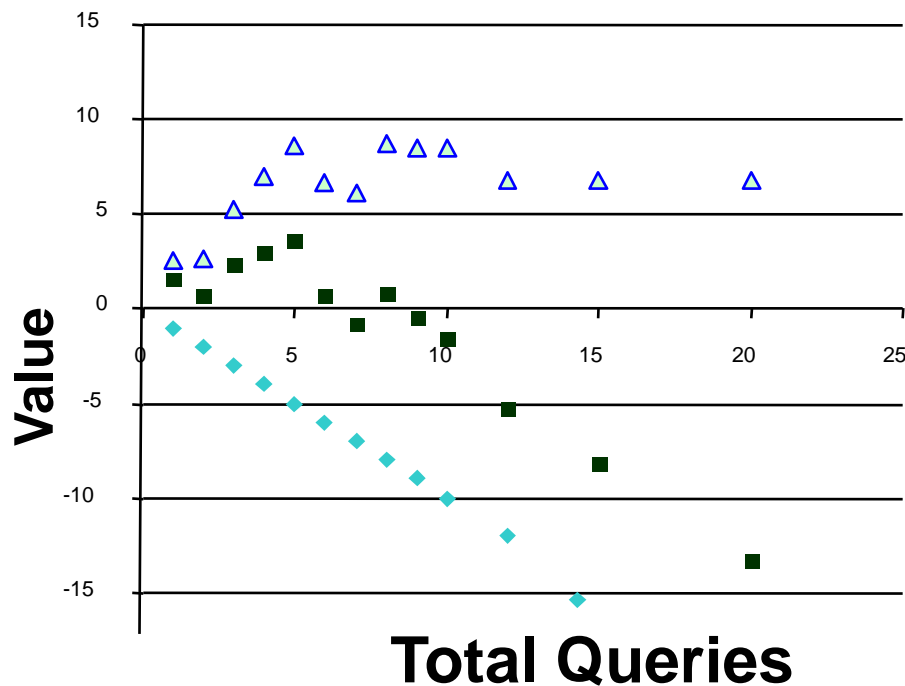
## AskMSR: Automatic Question Answering



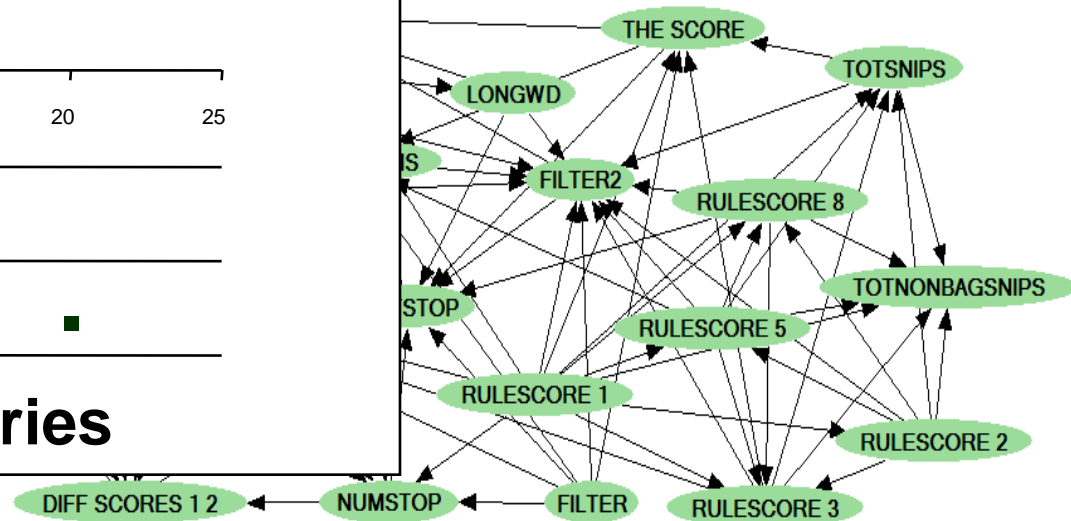
# Learning to Harness the Web

## AskMSR: Automatic Question Answering

“Where is the Orinoco River?”



(16%) two children (8%)



# AI in the Open World: Responsibilities

- ◆ Social value & quality of life
- ◆ Privacy, democracy, freedom
- ◆ Long-term AI futures



# AI in the Open World: Responsibilities

## AAAI Presidential Panel on *Long-Term AI Futures*

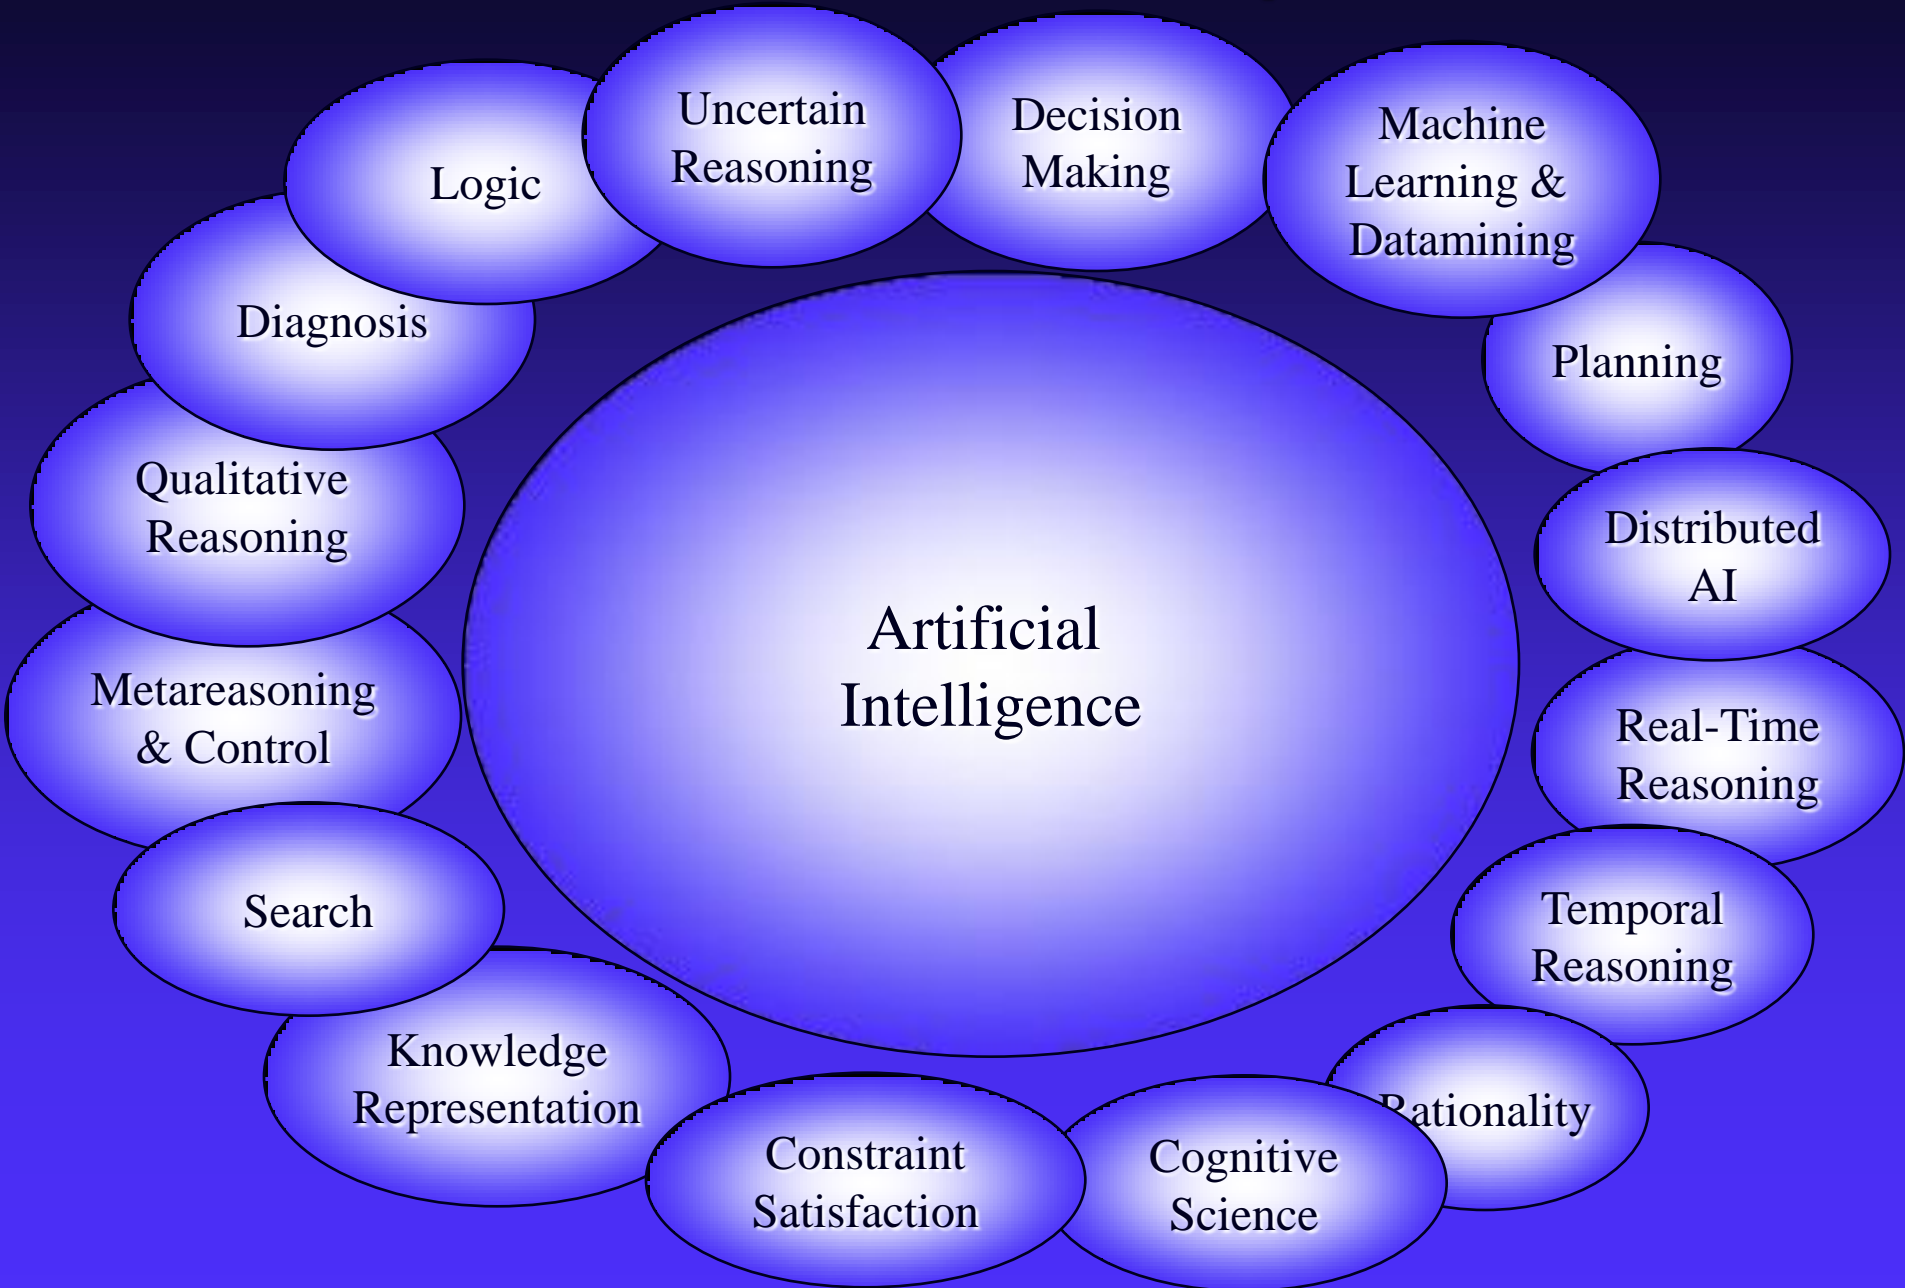
“...Deliberation will include reflection about concerns about long-term outcomes, and, if warranted, on potential recommendations for guiding research and on creating policies that might constrain or bias the behaviors of autonomous and semi-autonomous systems so as to address the concerns...”

# Stepping into the Open World

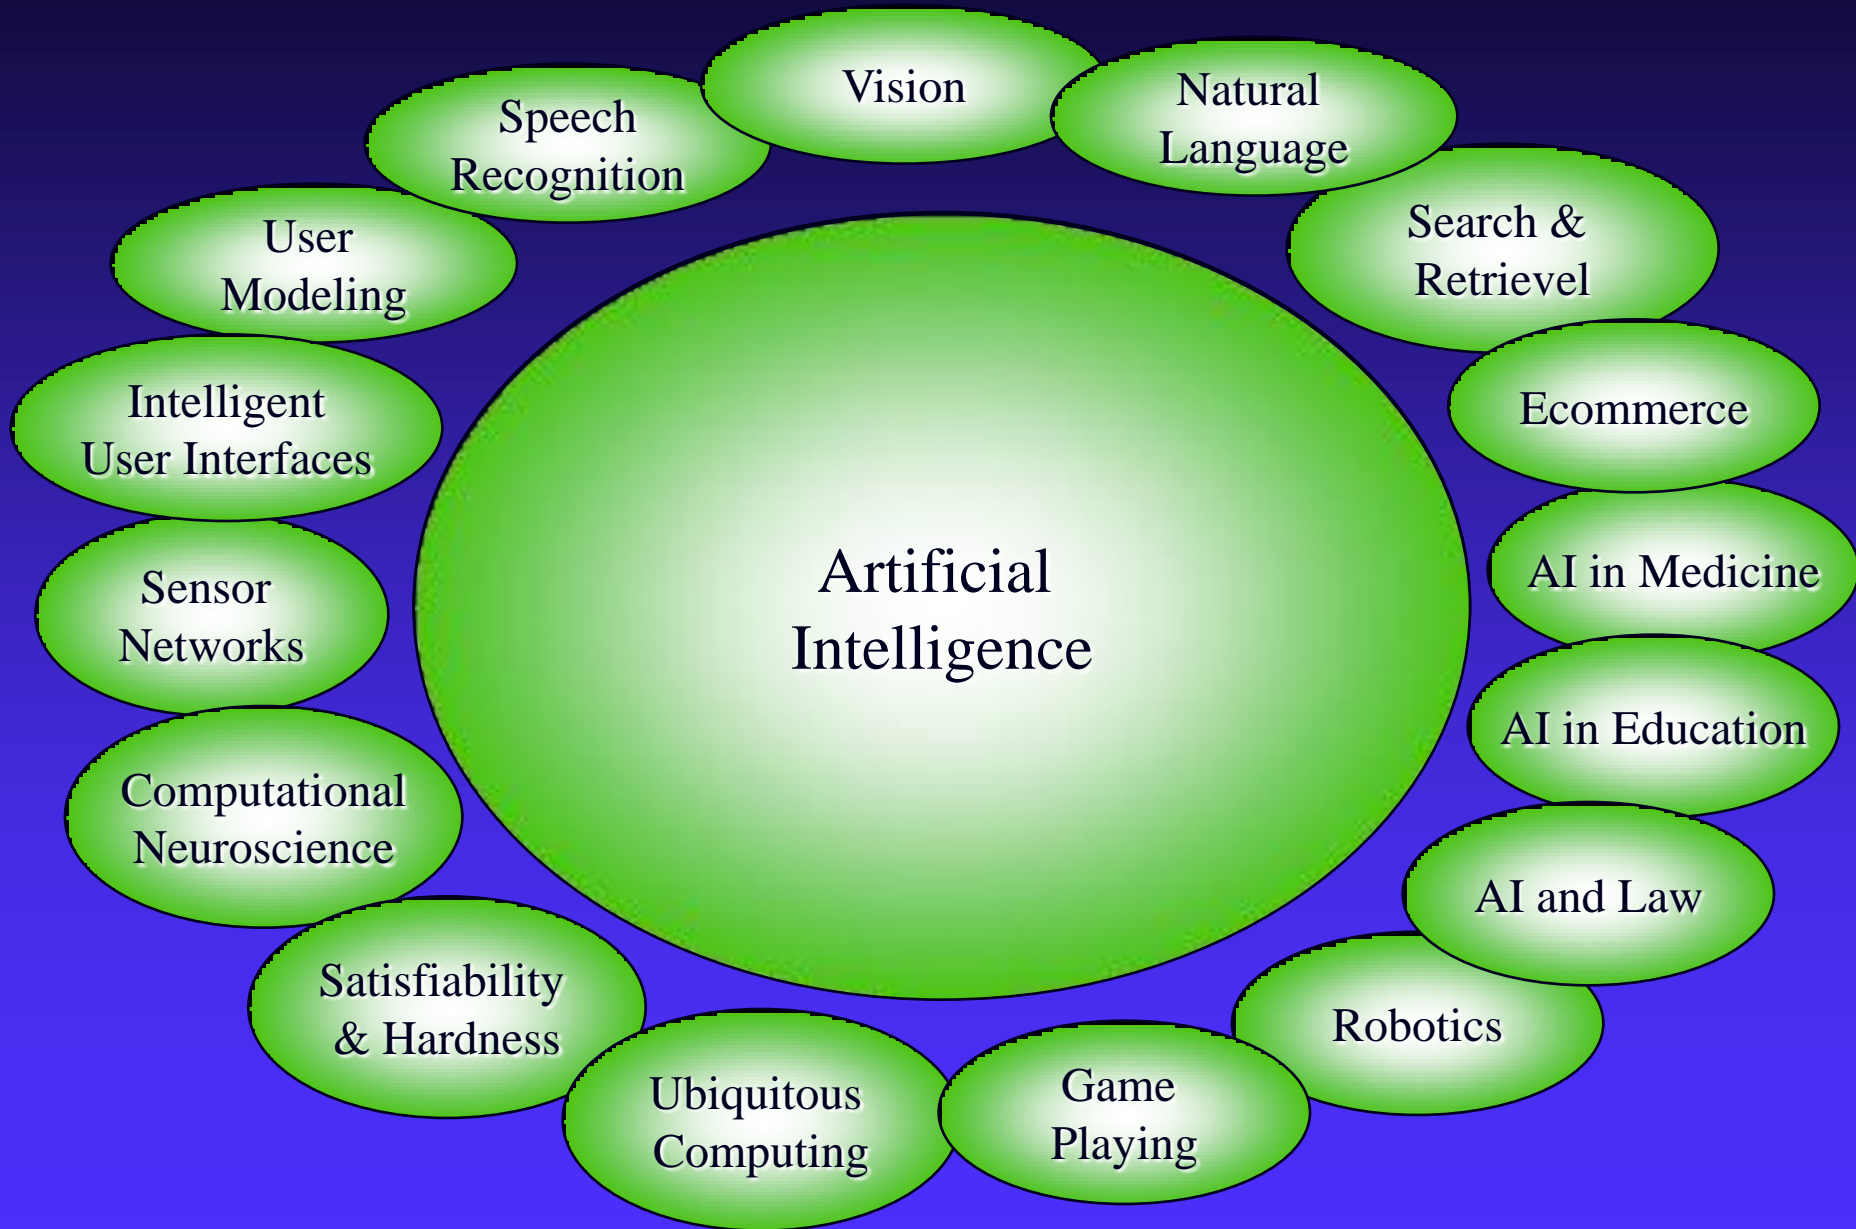
- ◆ **Key technical challenges**
- ◆ **AI moving into the world**
- ◆ **AI research community**



# Evolution of Subdisciplines

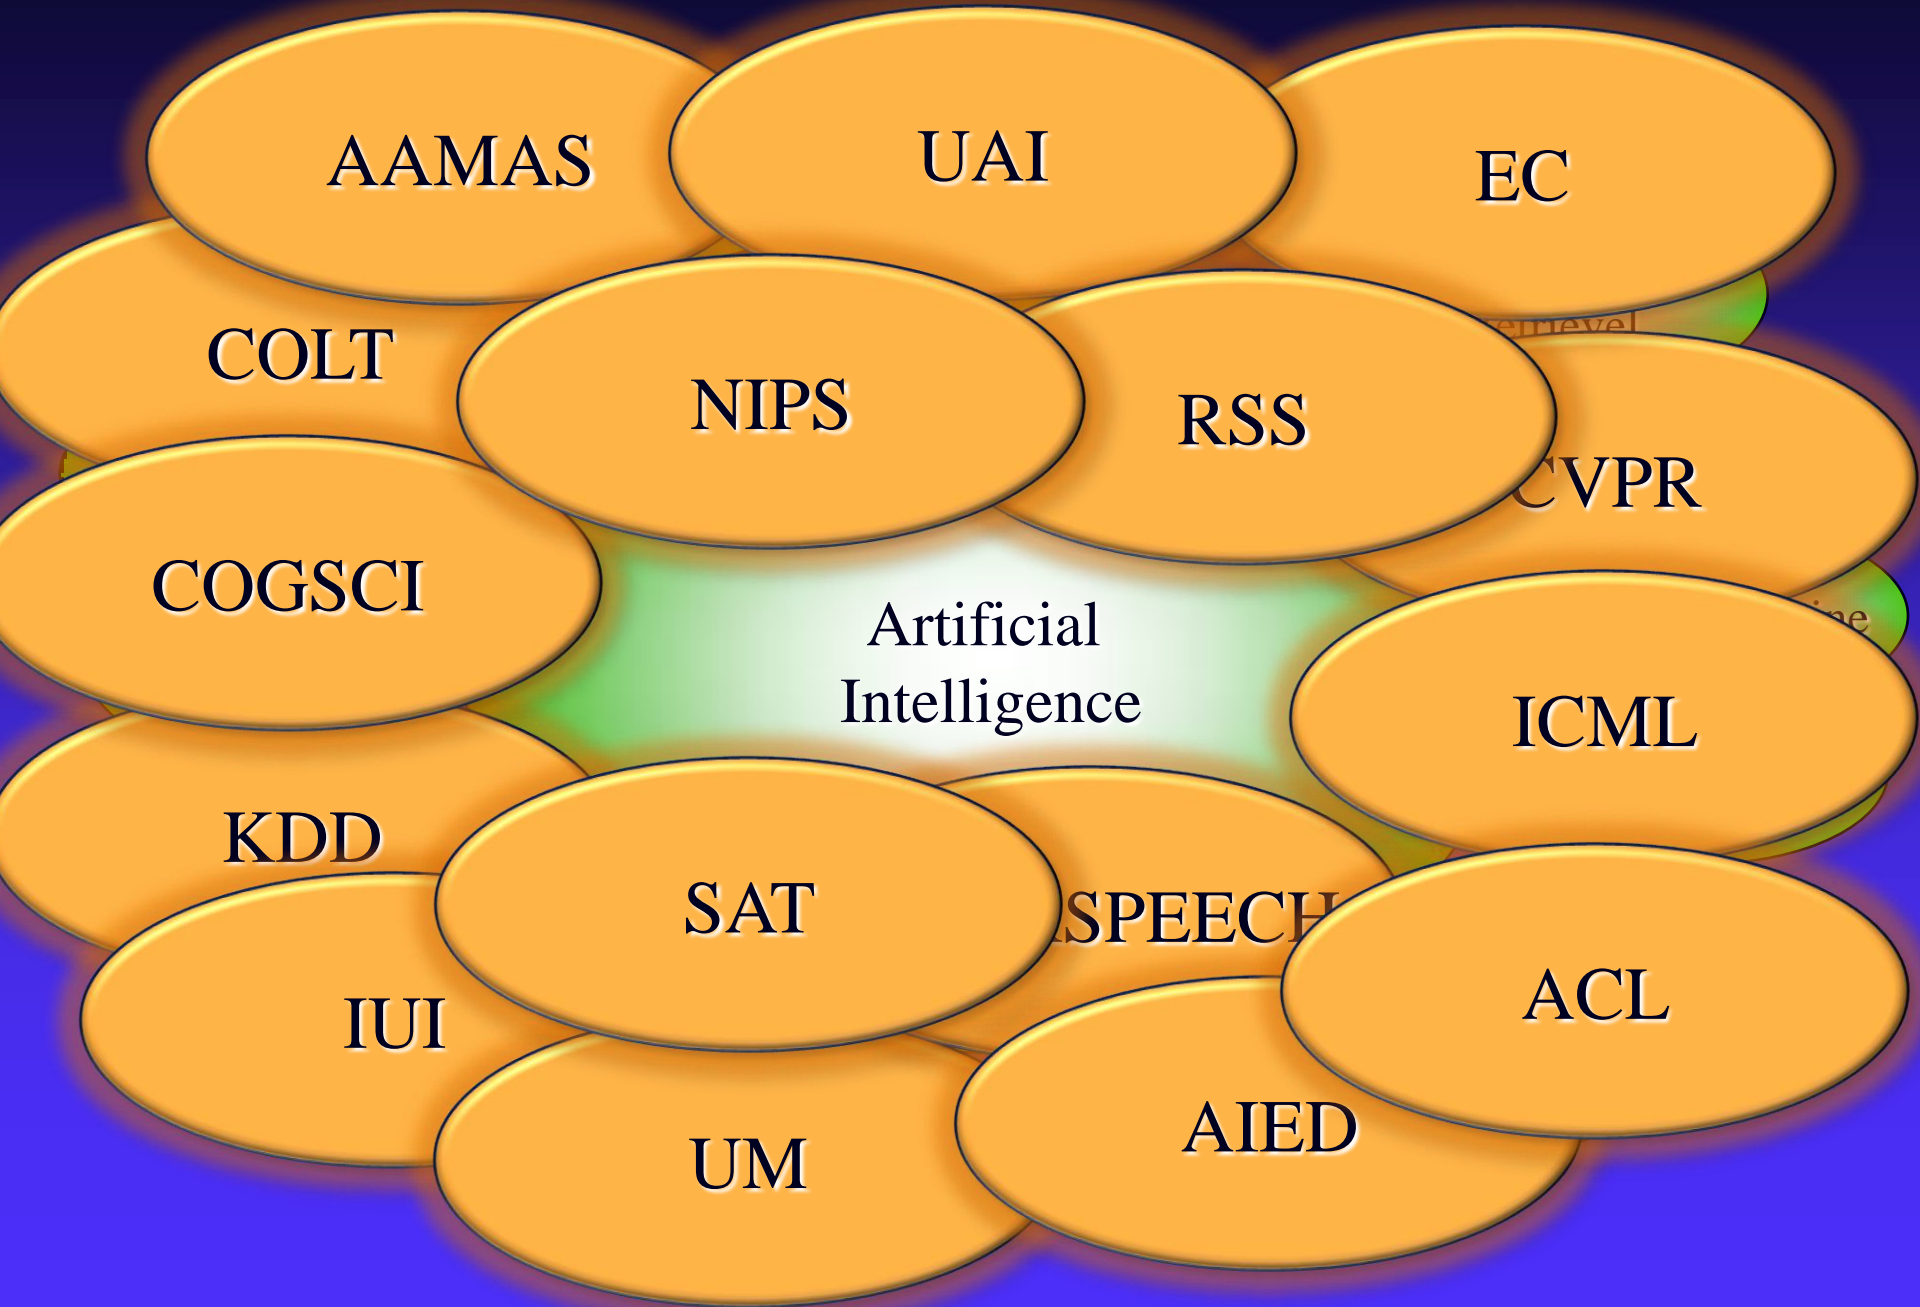


# Evolution of Specialty Areas





# Evolution of Communities



# Cooperation, Coordination, and Innovation in the Open World

*On the Nature of the Organization*

Herb Simon, 1947

‘Three bricklayers were asked what they were doing.

“Laying bricks,”

“Building a wall,”

“Helping to build a great cathedral”

...were their respective answers.’

# Bricks, Arches, Domes ...and Spandrels





# Bricks, Arches, Domes ...and Spandrels



# In the Distance, Through the Mist



# In the Distance, Through the Mist





