



16.410 and 16.412: Principles of Autonomy and Decision Making

Prof Brian Williams, Prof Emilio Frazzoli and Sertac Karaman
September, 8th, 2010



Assignments

- Homework:
 - Class signup, return at end of class;
 - Problem Set #1: Java warm up
Out Today,
Due next Wednesday, September 15th
- Reading:
 - Today: *[AIMA] Ch. 2, [JINS] Ch. 1-3,5*
 - Monday: search *[AIMA] Ch. 3.1-4*



Outline

1. Trends in Computing
2. Examples of autonomous systems - Williams
3. Autonomous systems architectures
4. Principles = modeling + inference + search
5. More examples – Frazzoli
6. Course logistics and schedule
7. Projects and programming

Brian Williams, Fall 10

3

3



Human Brain

- 100 Billion neurons
- On average, connected to 1 K others
- Neurons are slow. Firing rates < 100 Hz.
- Can be classified into
 - Sensory
 - Motor
 - Central
 - (reasoning, problem solving, language..)

Images of brain activity removed due to copyright restrictions.

Brian Williams, Fall 10

4

4



Trends in Biological and Machine Evolution - Moravec



- 1 neuron = 1000 instructions/sec
- Human brain then processes 10^{14} IPS
- 1 synapse = 1 byte of information
- Human brain then has 10^{14} bytes of storage
- In 2000, we had 10^9 IPS and 10^9 bytes on a desktop machine
- In 25 years, assuming Moore's law we obtain human level computing power



Outline



1. Trends in Computing
2. [Examples of autonomous systems - Williams](#)
3. Autonomous systems architectures
4. Principles = modeling + inference + search
5. More examples - Frazzoli
6. Course logistics and schedule
7. Projects and programming



Williams Research: Model-based Programming of Autonomous Systems



Robust, mission-directed agents:

1. Self-repairing agents
2. Agents that are agile
3. Science explorers

1. Self-Repairing Agents

- 7 year cruise
- ~ 150 - 300 ground operators
- ~ 1 billion \$
- 7 years to build

Affordable Missions

↓

- 150 million \$
- 2 year build
- 0 ground ops

Cassini Maps Titan


© Source unknown. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <http://ocw.mit.edu/fairuse>.



Remote Agent on Deep Space One



Goals



1. Commanded by giving goals
2. Reasoned from commonsense models
3. Closed loop on goals

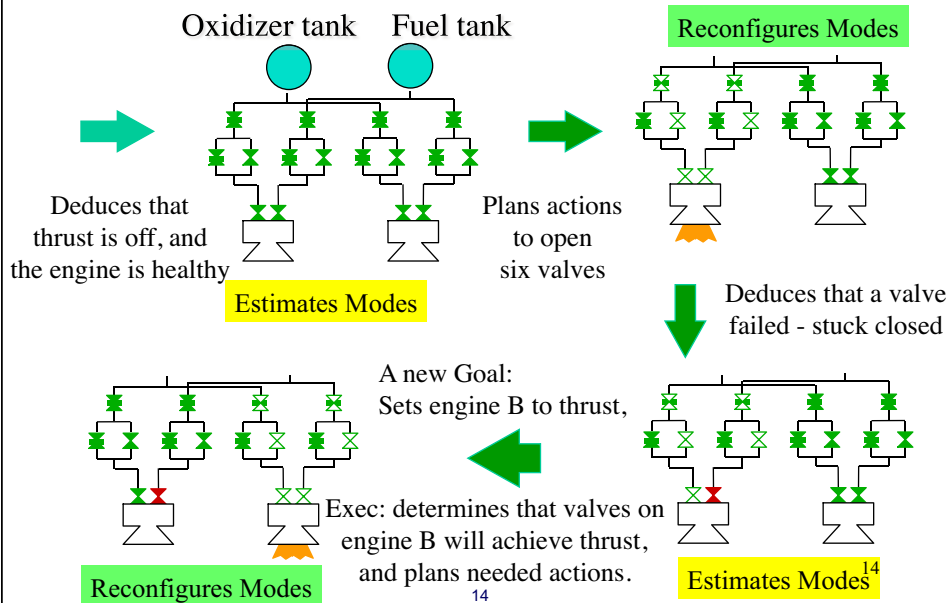
Brian Williams, Fall 10

[Williams & Nayak, AAI 95; Muscettola et al, AIJ 00]

13




A Goal sets engine A state to thrusting, and the agent . . .



14

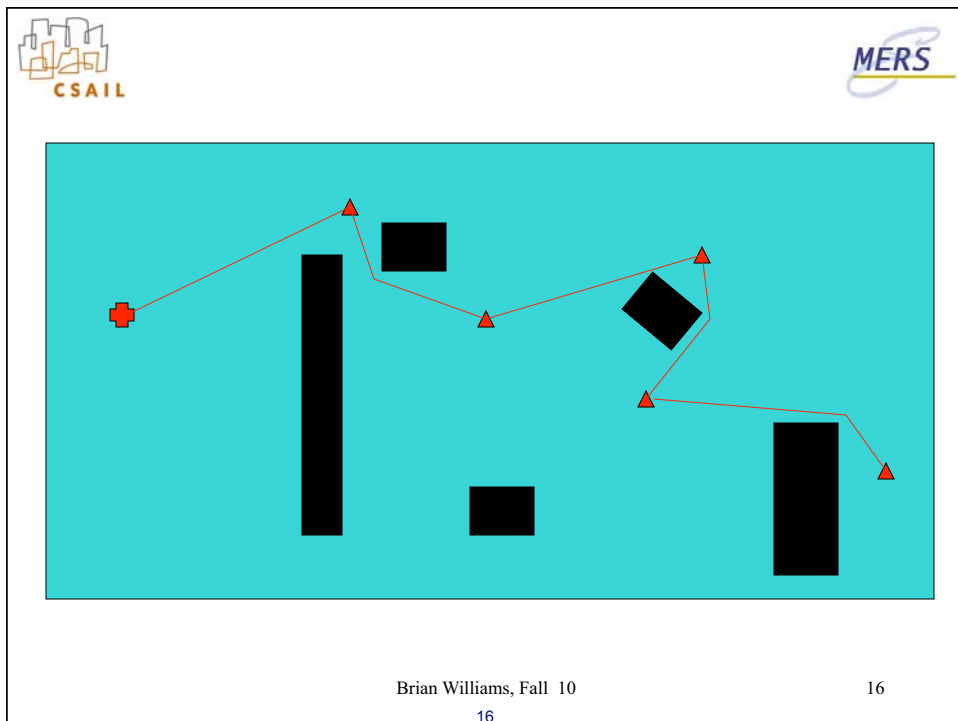
2. Agile Agents



Brian Williams, Fall 10

15

Image credit: NASA.



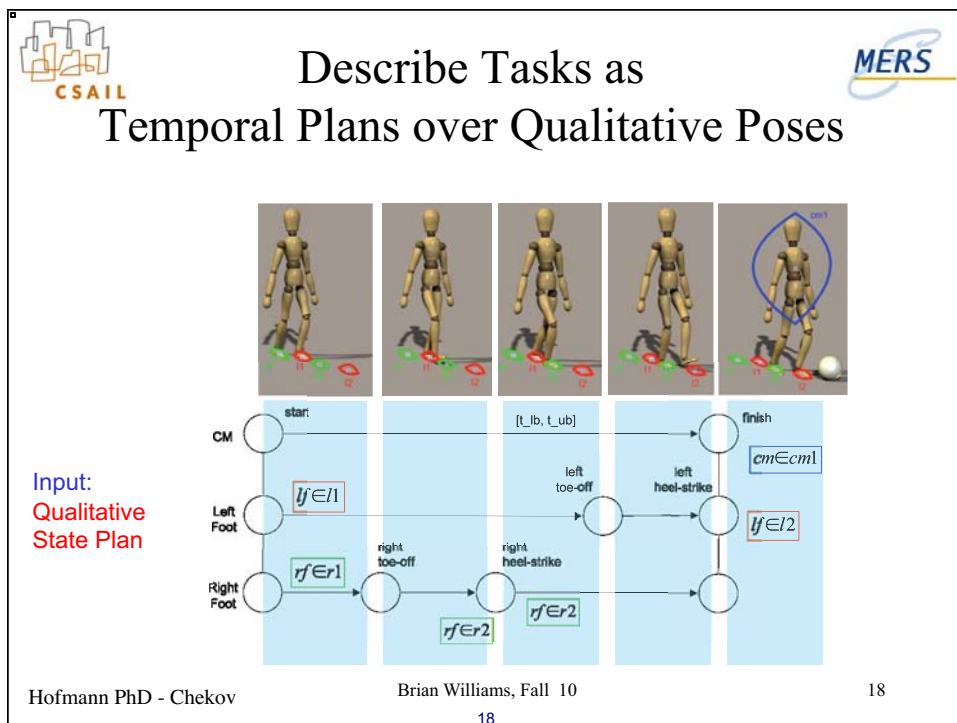
Brian Williams, Fall 10

16

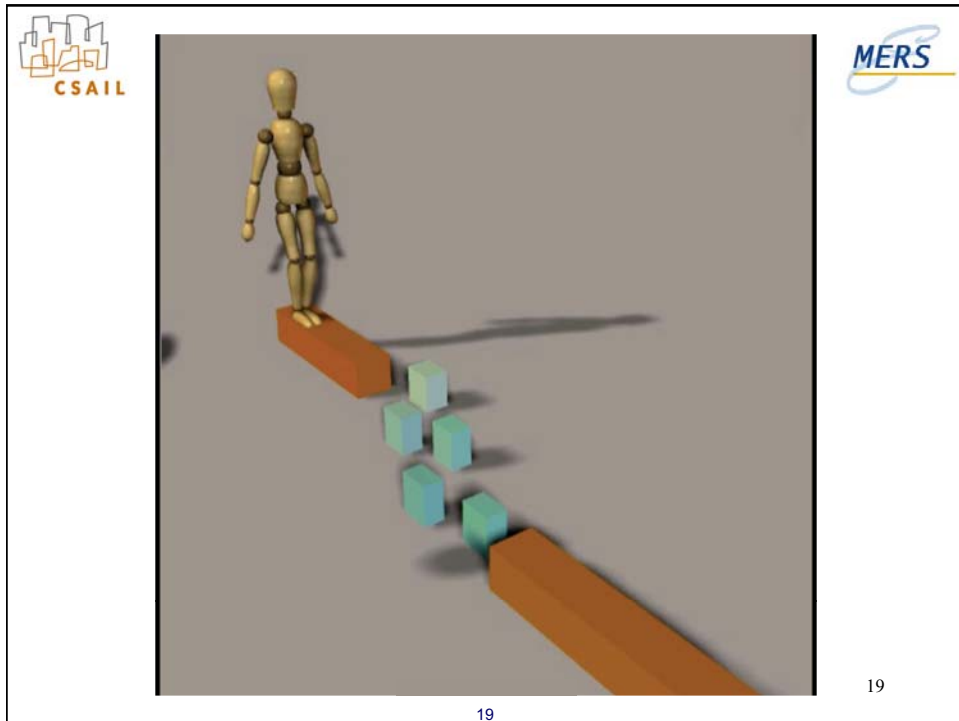
16

Image of hilly terrain removed due to copyright restrictions.

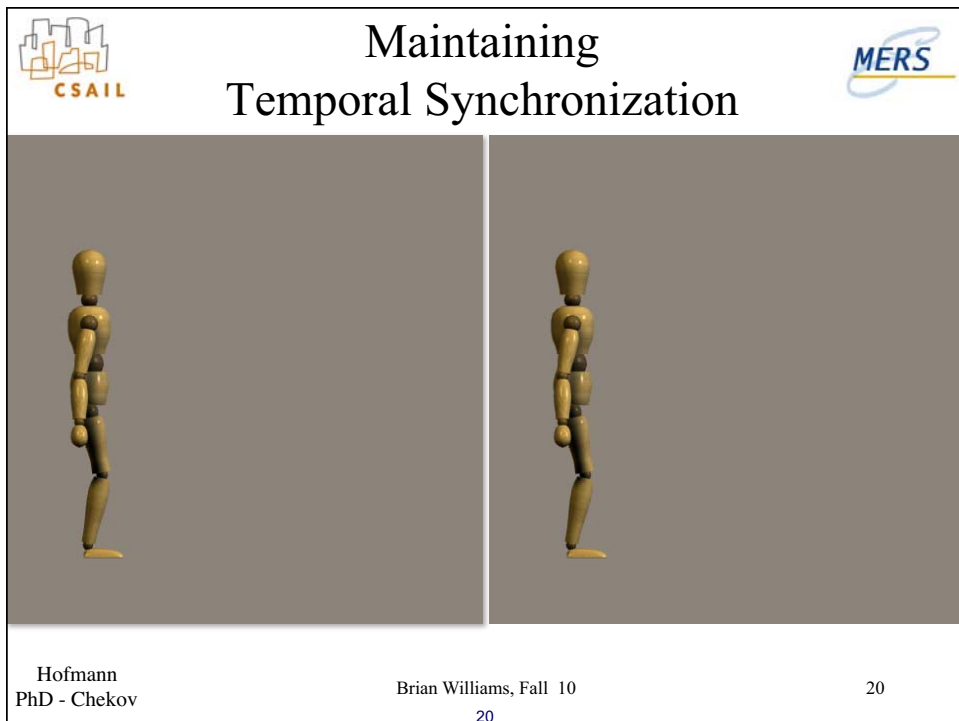
Courtesy JPL




Courtesy of Andreas Hofmann. Used with permission.




Courtesy of Andreas Hofmann. Used with permission.




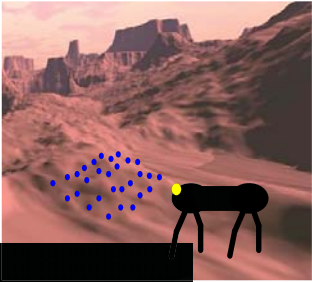
Courtesy of Andreas Hofmann. Used with permission.

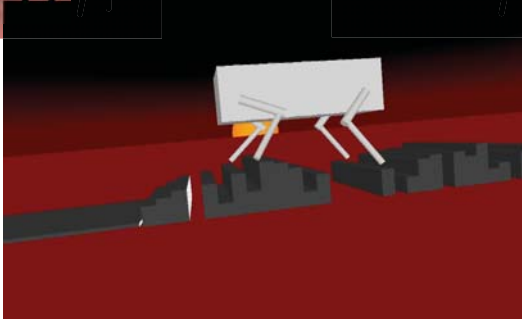


CDIO Capstone: Moretta










Hofmann
PhD - Chekov


Brian Williams, Fall 10

21

Courtesy of Andreas Hofmann. Used with permission.



3. Science Explorers

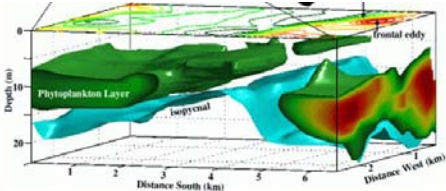


MBARI Dorado-class AUV:


- 6000m rated
- 20 hour operation
- Multibeam Sonars
- 3+ knots speed

Challenges:

- Long mission duration
- Limited communication
- GPS unavailable
- Uncertainty
 - tides and currents
 - estimation error



Graph © source unknown. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <http://ocw.mit.edu/fairuse>.



Brian Williams, Fall 10

22



Robust, Goal-directed Deep Sea Exploration



00:00 Go to x_1, y_1
00:20 Go to x_2, y_2
00:40 Go to x_3, y_3
...
04:10 Go to x_n, y_n

Command script

Commands

Plant



Brian Williams, Fall 10

Leaute & Williams, AAAI 05

23



Robust, Goal-directed Deep Sea Exploration



"Remain in mapping region for at least 100s, then remain in bloom region for at least 50s, then return to pickup region. Avoid obstacles at all times"

Qualitative State Plan

Model-based Executive

Observations

Commands

Plant





Optimal

Brian Williams, Fall 10

Leaute & Williams, AAAI 05

24



Blackmore PhD


25

25

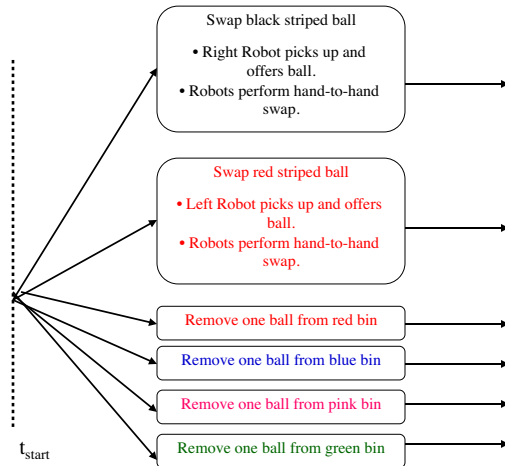

Courtesy of Lars Blackmore. Used with permission.

Image of NASA Athlete rover removed due to copyright restrictions.

Multi-Robot Teamwork



Agents choose and schedule activities




Brian Williams, Fall 10

27



Multi-Robot Teamwork



- Multi-modal
- Multi-agent
- Multi-task
- Multi-robot

Mit Massachusetts Institute of Technology

Brian Williams, Fall 10

28






Image of Vecna's Bear Robot removed due to copyright restrictions.



Outline

1. Trends in Computing
2. Examples of autonomous systems
3. **Autonomous systems architectures**
4. Principles = modeling + inference + search
5. More examples
6. Course logistics and schedule
7. Projects and programming

Agent Paradigms

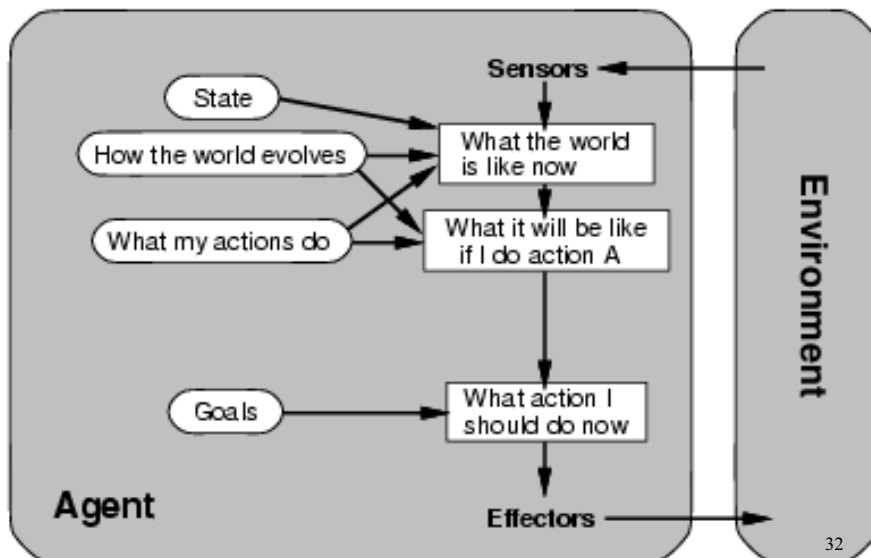


Brian Williams, Fall 10

31

31

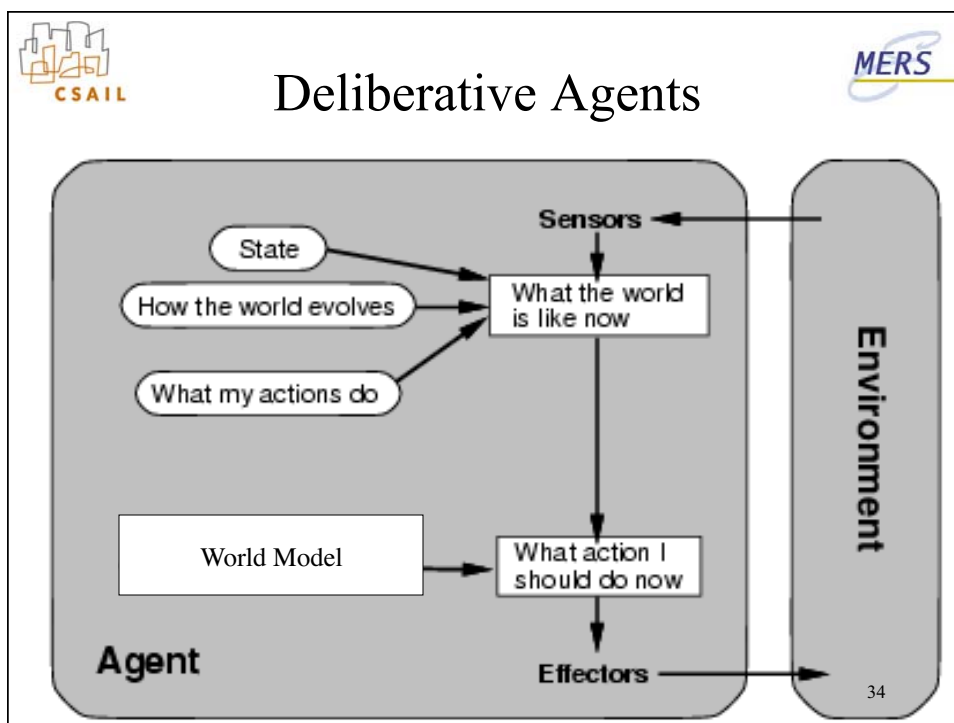
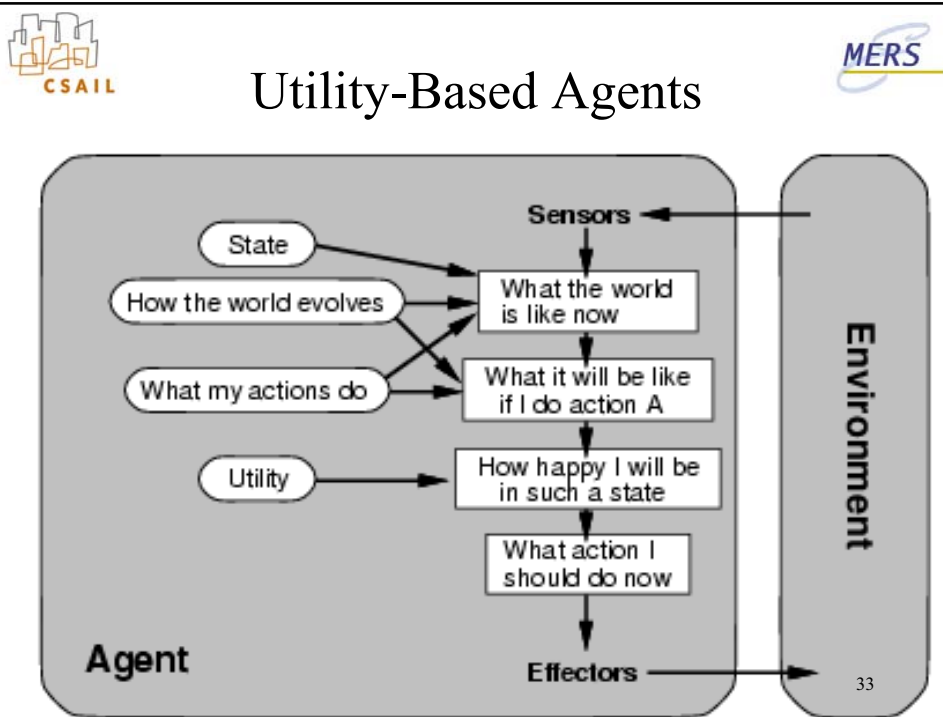
Goal-Oriented Agents



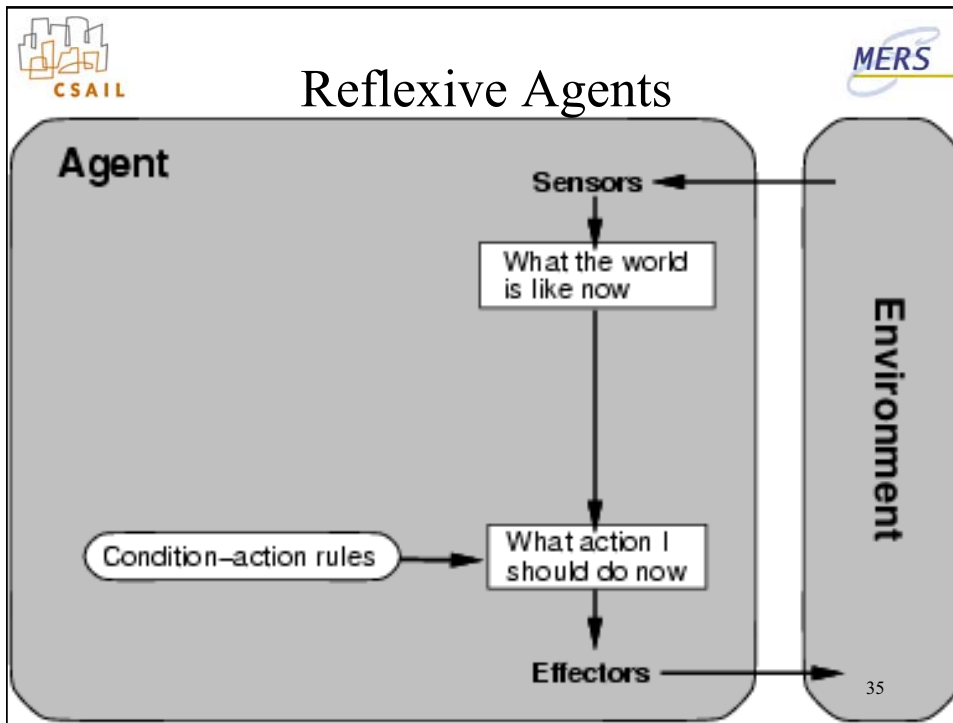
32

32

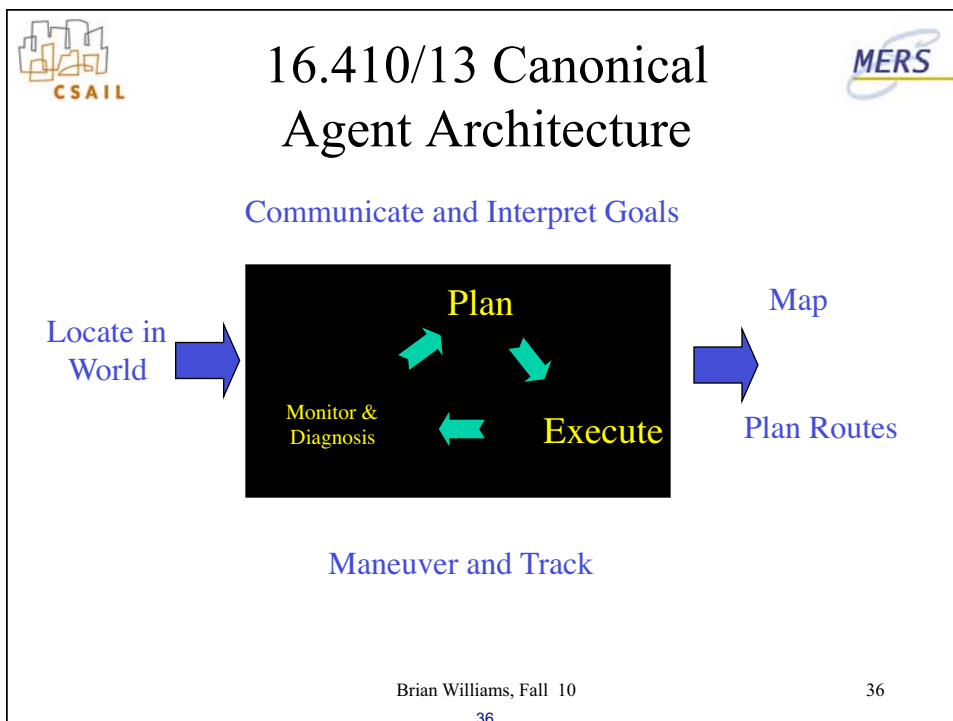
Courtesy of Stuart Russell and Peter Norvig. Used with permission.



Courtesy of Stuart Russell and Peter Norvig. Used with permission.



Courtesy of Stuart Russell and Peter Norvig. Used with permission.





Outline

1. Trends in Computing
2. Examples of autonomous systems
3. Autonomous systems architectures
4. Principles = modeling + inference + search
5. More examples
6. Course logistics and schedule
7. Projects and programming

Brian Williams, Fall 10

37

37



Modeling Example: Robust, Goal-directed Deep Sea Exploration

*“Remain in **mapping region** for at least 100s, then remain in **bloom region** for at least 50s, then return to **pickup region**.
Avoid obstacles at all times”*

Qualitative State Plan

Model-based Executive

Observations

Commands

Plant



Optimal


Brian Williams, Fall 10

Leaute & Williams, AAAI 05

38

Modeling Goal Behavior using Qualitative State Plans

A qualitative state plan is a model-based program that is unconditional, timed, and hybrid and provides flexibility in state and time.




“Remain in bloom region for between 50 and 70 seconds. Afterwards, remain in mapping region for between 40s and 50s. End in the pickup region. Avoid obstacles at all times. Complete the mission within 300s”

Approach: Frame as Model-Predictive Control using Mixed Logic or Integer / Linear Programming.

Leaute & Williams, AAAI 05

Brian Williams, Fall 10



39

Modeling Goal-directed Planning as a Mathematical Program

$$\min_{\mathbf{U}} J(\mathbf{X}, \mathbf{U}) + H(x_T)$$

s.t. ↑ Cost function (e.g. fuel consumption)

Dynamics
(Discrete time) $\forall_{0 \leq t \leq T-1} x_{t+1} = Ax_t + Bu_t$

Constraints $\bigwedge_{t=0}^T \bigwedge_{i=0}^N \bigvee_{j=0}^M h_t^{iT} x_t \leq g_t^{ij}$

Mixed Logic or Integer

$\mathbf{X} = [x_0 \cdots x_t]$ State vector (e.g. position of vehicle)

$\mathbf{U} = [u_0 \cdots u_{t-1}]$ Control inputs

Brian Williams, Fall 10

40



Specify Building Blocks using Declarative Programming



- Mathematical Programming
- Constraint Programming
- Logic Programming
- Agent Programming
 - Model-based programming
 - Timed concurrent constraint programming
 - Golog
 - Temporal logic programming

Brian Williams, Fall 10

41

41



Solve Declarative Programs using:



Search + Inference

Search: try taking the subway, or
try taking the bus.

Inference: It takes 35 minutes to get to MIT,
20 min subway + 15 min walking.

Brian Williams, Fall 10

42

42



Outline

1. Trends in Computing
2. Examples of autonomous systems
3. Autonomous systems architectures
4. Principles = modeling + inference + search
5. More examples
6. Course logistics and schedule
7. **Projects and programming**

MIT OpenCourseWare
<http://ocw.mit.edu>

16.410 / 16.413 Principles of Autonomy and Decision Making
Fall 2010

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.