Applications of AI

Alan Mackworth

UBC CS 322 - Intro 3

January 7, 2013

Textbook §1.5 - 1.6

Today's Lecture

- Recap from last lecture
- Further Representational Dimensions
- Applications of AI

Representation and Reasoning (R&R) System

problem ⟹ representation ⟹ computation⟹ representation ⟹ solution

- A representation language that allows description of
  - The environment and
  - Problems (questions/tasks) to be solved

- Computational reasoning procedures to
  - Compute a solution to a problem
  - E.g., an answer/sequence of actions

- How should an agent act given the current state of its environment and its goals?
- How should the environment be represented in order to help an agent to reason effectively?

Main Representational Dimensions Considered

Domains can be classified by the following dimensions:

- 1. Uncertainty
  - Deterministic vs. stochastic domains
- 2. How many actions does the agent need to perform?
  - Static vs. sequential domains

An important design choice is:

- 3. Representation scheme
  - Explicit states vs. propositions vs. relations

Features vs. States, another example

T_{11}: student 1 takes course 1
T_{12}: student 1 takes course 2
T_{21}: student 2 takes course 1
T_{22}: student 2 takes course 2

Does student 2 take course 2?

- Feature-based: Is T_{22} true?
- State-based: are we in one of the red states?

	T_{11}	T_{12}	T_{21}	T_{22}
S_0	0	0	0	0
S_1	0	0	0	1
S_2	0	0	1	0
S_3	0	0	1	1
S_4	0	1	0	0
S_5	0	1	0	1
S_6	0	1	1	0
S_7	0	1	1	1
S_8	1	0	0	0
S_9	1	0	0	1
S_{10}	1	0	1	0
S_{11}	1	0	1	1
S_{12}	1	1	0	0
S_{13}	1	1	0	1
S_{14}	1	1	1	0
S_{15}	1	1	1	1

Course overview

Dimen-sions / Course Modules	Deterministic vs. Stochastic	Static vs. Sequential	States vs. Features vs. Relations
1. Search	Deterministic	Static	States
2. CSPs	Deterministic	Static	Features
3. Planning	Deterministic	Sequential	States or Features
4. Logic	Deterministic	Static	Relations
5. Uncertainty	Stochastic	Static	Features
6. Decision Theory	Stochastic	Sequential	Features

Example problems:

"find path in known map"

"are deliveries feasible?"

"what order to do things in to finish jobs fastest?"

"HasCoffee(Person) if InRoom(Person, Room) ∧ DeliveredCoffee(Room)"

"probability of slipping"

"given that I may slip and the utilities of being late and of crashing, should I take a detour?"

Today's Lecture

- Recap from last lecture
- Further Representational Dimensions
- Applications of AI

Further Dimensions of Representational Complexity

We've already discussed:

1. Deterministic versus stochastic domains
2. Static vs. Sequential domains
3. Explicit state or features or relations

Some other important dimensions of complexity:

4. Flat vs. hierarchical representation
5. Knowledge given vs.
   knowledge learned from experience
6. Goals vs. complex preferences
7. Single-agent vs. multi-agent
8. Perfect rationality vs. bounded rationality

4. Flat vs. hierarchical

- Should we model the whole world on the same level of abstraction?
  - Single level of abstraction: flat
  - Multiple levels of abstraction: hierarchical

- *Example: Planning a trip from here to a resort in Cancun*
  - *Going to the airport*
    - *Take a cab*
      - *Call a cab*
        - *Lookup number*
        - *Dial number*
      - *Ride in the cab*
      - *Pay for the cab*
    - *Check in*
    - *….*

- Delivery robot: Plan on level of cities, districts, buildings, …
- This course: mainly flat representations
  - Hierarchical representations required for scaling up

5. Knowledge given vs. knowledge learned from experience

- The agent is provided with a model of the world once and for all OR
- The agent can learn how the world works based on experience
  - in this case, the agent often still does start out with some prior knowledge (no *tabula rasa*!)

- Delivery robot: Known/learned map, prob. of slipping, …
- This course: mostly knowledge given
  - Learning: CPSC 340 and CPSC 422

6. Goals vs. (complex) preferences

- An agent may have a goal that it wants to achieve
  - E.g., there is some state or set of states of the world that the agent wants to be in
  - E.g., there is some proposition or set of propositions that the agent wants to make true
- An agent may have preferences
  - E.g., a preference/utility function describes how happy the agent is in each state of the world
  - Agent's task is to reach a state which makes it as happy as possible
- Preferences can be complex
  - E.g., Diagnostic assistant faces multi-objective problem
    - Life expectancy, suffering, risk of side effects, costs, …
  - Delivery robot: “deliver coffee!” vs “mail trumps coffee, but Chris needs coffee quickly, and don’t stand in the way”
- This course: goals and simple preferences
  - Some scalar, e.g. linear combination of competing objectives

7. Single-agent vs. Multiagent domains

- Does the environment include other agents?
- If there are other agents whose actions affect us
  - It can be useful to explicitly model their goals and beliefs, and how they react to our actions

- Other agents can be: cooperative, neutral, competitive, or a bit of each

- Delivery robot: Are there other agents?
  - Should I coordinate with other robots?
  - Are kids out to trick me?
- This course: only single agent scenario
  - Multiagent problems tend to be complex
  - Exception: deterministic 2-player games can be formalized easily

8. Perfect rationality vs. bounded rationality

We've defined rationality as an abstract ideal

- Is the agent able to live up to this ideal?
  - Perfect rationality:
    the agent can derive what the best course of action is
  - Bounded rationality:
    the agent must make good decisions
    based on its perceptual, computational and memory limitations
- Delivery robot:
  - ”Find perfect plan” vs.
  - “Can’t spend an hour thinking (thereby delaying action) to then deliver packages a minute faster than by some standard route”
- This course: mostly perfect rationality
  - But also consider anytime algorithms for optimization problems

Today's Lecture

- Recap from last lecture
- Further Representational Dimensions
- ➡ Applications of AI

Course Map

Dimen-sions / Course Modules	Deterministic vs. Stochastic	Static vs. Sequential	States vs. Features vs. Relations
1. Search	Deterministic	Static	States
2. CSPs	Deterministic	Static	Features
3. Planning	Deterministic	Sequential	States or Features
4. Logic	Deterministic	Static	Relations
5. Uncertainty	Stochastic	Static	Features
6. Decision Theory	Stochastic	Sequential	Features

Search: Checkers

- Early learning work in 1950s by Arthur Samuel at IBM
- Chinook program by Jonathan Schaeffer (UofA)
  - Search to explore the space of possible moves and their consequences
  - 1994: world champion
  - 2007: declared unbeatable

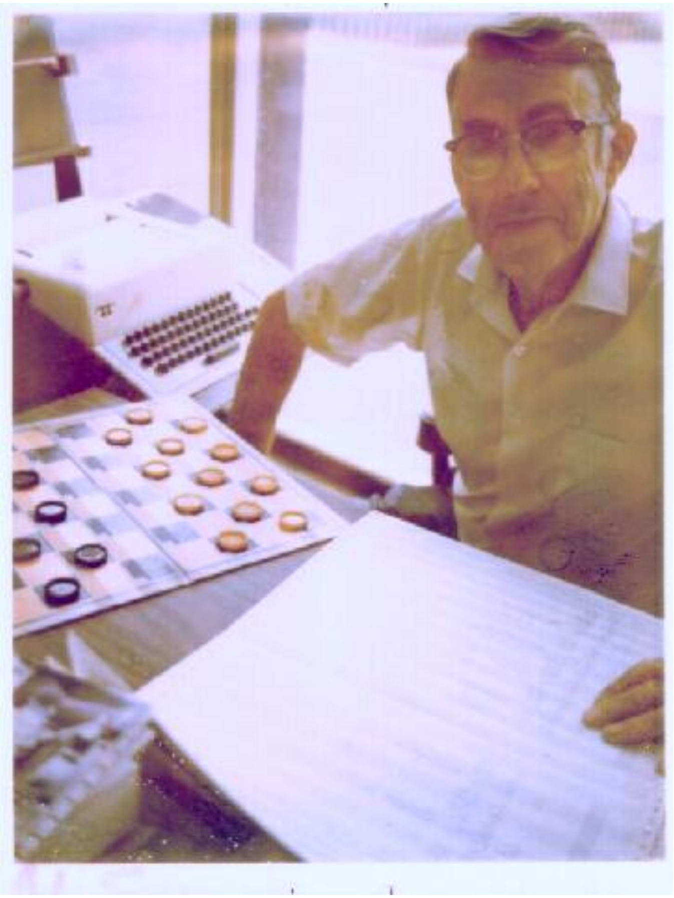

Search: Chess

- In 1997, Gary Kasparov, the chess grandmaster and reigning world champion played against Deep Blue, a program written by researchers at IBM

Source: *IBM Research*

Search: Chess

- Deep Blue won 3 games, lost 2, tied 1

  - 30 CPUs + 480 chess processors
  - Searched 126,000,000 nodes per sec
  - Generated 30 billion positions per move reaching depth 14 routinely

Course Map

Dimen- sions / Course Modules	Deterministic vs. Stochastic	Static vs. Sequential	States vs. Features vs. Relations
1. Search	Deterministic	Static	States
2. CSPs	Deterministic	Static	Features
3. Planning	Deterministic	Sequential	States or Features
4. Logic	Deterministic	Static	Relations
5. Uncertainty	Stochastic	Static	Features
6. Decision Theory	Stochastic	Sequential	Features

CSP: Sudoku

Sudoku rules are extremely easy: Fill all empty squares so that the numbers 1 to 9 appear once in each row, column and 3x3 box.

Sudoku Puzzle

	9	3	6	2	8	1	4	
	6						5	
	3			1			9	
	5		8		2		7	
	4			7			6	
	8						3	
	1	7	5	9	3	4	2	

Sudoku Solution

2	7	1	9	5	4	6	8	3
5	9	3	6	2	8	1	4	7
4	6	8	1	3	7	2	5	9
7	3	6	4	1	5	8	9	2
1	5	9	8	6	2	3	7	4
8	4	2	3	7	9	5	6	1
9	8	5	2	4	1	7	3	6
6	1	7	5	9	3	4	2	8
3	2	4	7	8	6	9	1	5

Constraint optimization problems

- Optimization under side constraints (similar to CSP)
- E.g. mixed integer programming (software: IBM CPLEX)
  - Linear program: max $c^T x$ such that $Ax \leq b$
  - Mixed integer program: additional constraints, $x_i \in Z$ (integers)
  - NP-hard, widely used in operations research and in industry

Transportation/Logistics: SNCF, United Airlines UPS, United States Postal Service, …

Supply chain management software: Oracle, SAP,…

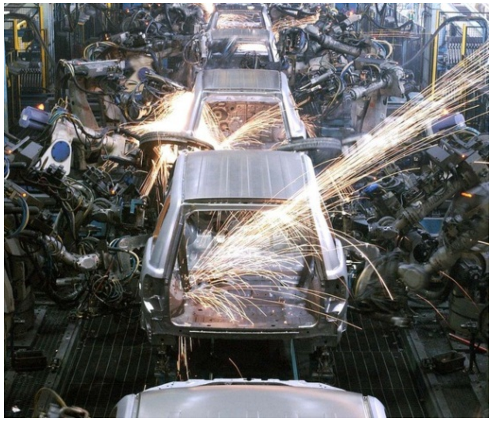

Production planning and optimization: Airbus, Dell, Thyssen, Toyota,Nissan, ...

Course Map

Dimen-sions / Course Modules	Deterministic vs. Stochastic	Static vs. Sequential	States vs. Features vs. Relations
1. Search	Deterministic	Static	States
2. CSPs	Deterministic	Static	Features
3. Planning	Deterministic	Sequential	States or Features
4. Logic	Deterministic	Static	Relations
5. Uncertainty	Stochastic	Static	Features
6. Decision Theory	Stochastic	Sequential	Features

Planning: Spacecraft Control

NASA: Deep Space One spacecraft

- operated autonomously for two days in May, 1999:
  - determined its precise position using stars and asteroids despite a malfunctioning ultraviolet detector
  - planned the necessary course adjustment
  - fired the ion propulsion system to make this adjustment

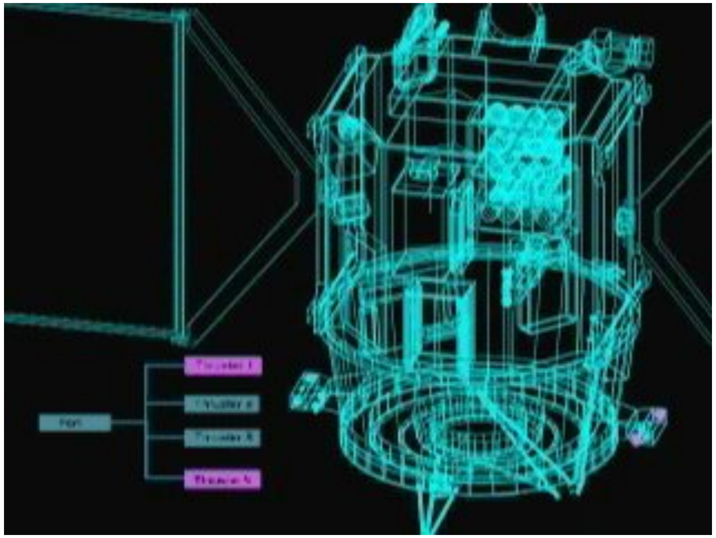

Source: *NASA*

Course Map

Dimen-sions / Course Modules	Deterministic vs. Stochastic	Static vs. Sequential	States vs. Features vs. Relations
1. Search	Deterministic	Static	States
2. CSPs	Deterministic	Static	Features
3. Planning	Deterministic	Sequential	States or Features
4. Logic	Deterministic	Static	Relations
5. Uncertainty	Stochastic	Static	Features
6. Decision Theory	Stochastic	Sequential	Features

Logic: Cyc

- AI project that started 1984 with the objective
  - to codify, in machine-usable form, millions of pieces of knowledge that comprise human common sense

- Logic reasoning procedures, e.g.
  - Every tree is a plant
  - Plants die eventually
  - Therefore, every tree dies eventually

- Criticisms include
  - Difficulty of adding knowledge manually
  - Non-scalability
  - Empirical evaluation – no benchmarks

Logic: Cyc

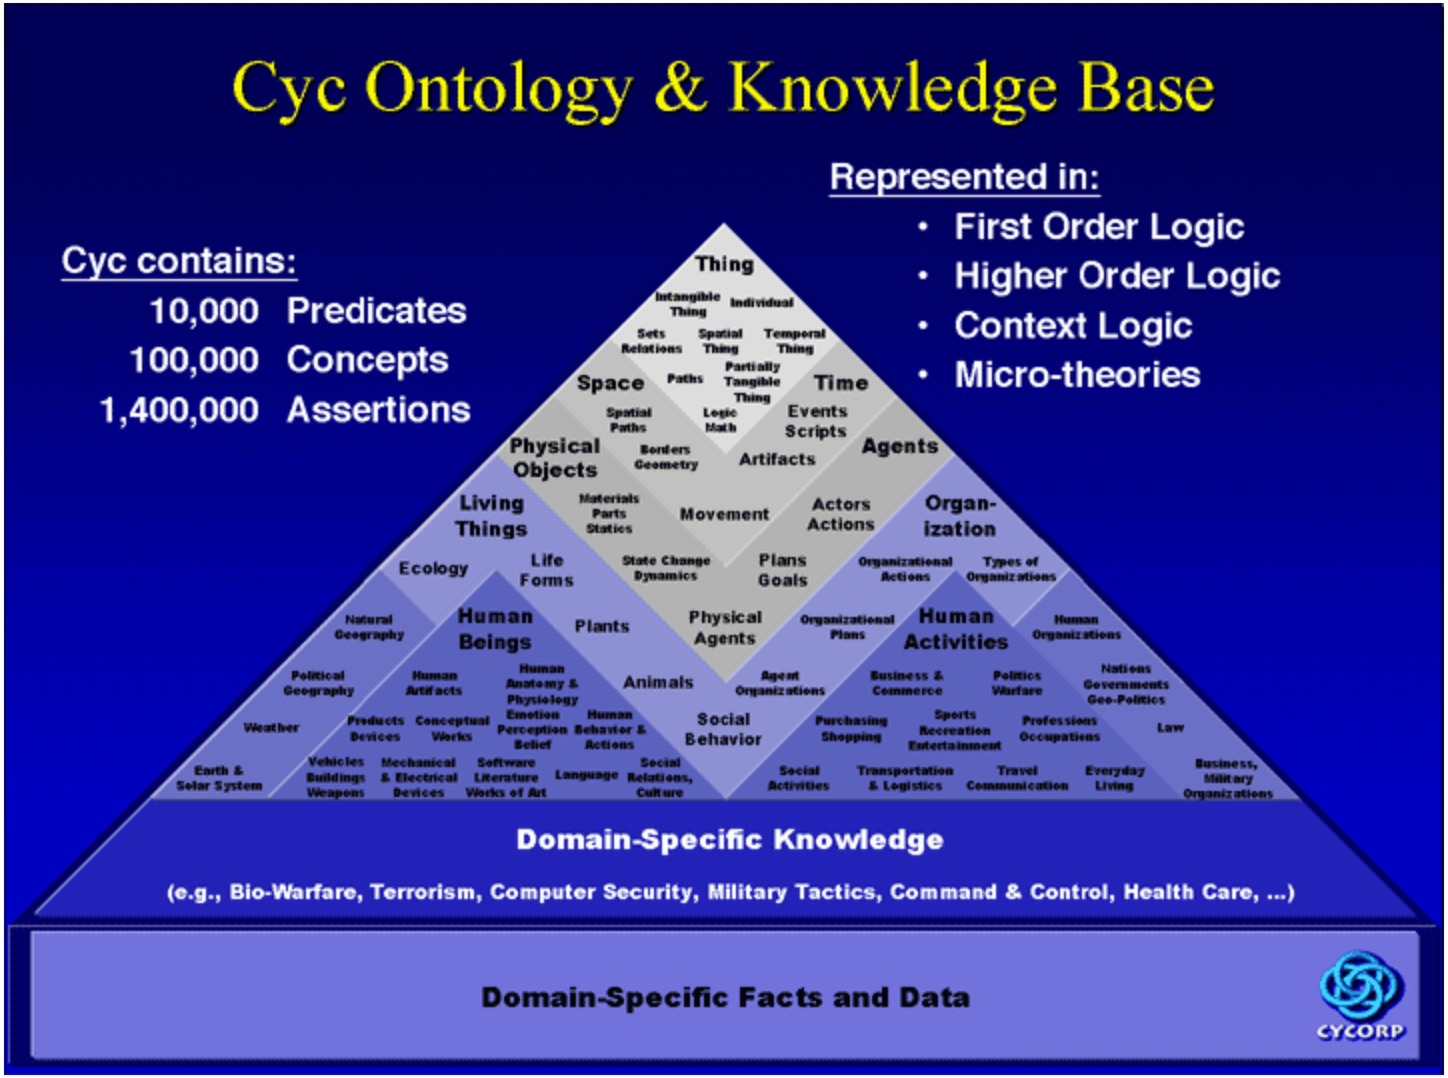

Cyc Ontology & Knowledge Base
Cyc contains:
10,000 Predicates
100,000 Concepts
1,400,000 Assertions
Represented in:
• First Order Logic
• Higher Order Logic
• Context Logic
• Micro-theories
Thing
Intangible Thing
Individual
Sets Relations
Spatial Thing
Temporal Thing
Space
Paths
Partially Tangible Thing
Time
Spatial Paths
Logic Math
Events Scripts
Physical Objects
Borders Geometry
Artifacts
Agents
Living Things
Materials Parts Statics
Movement
Actors Actions
Organ-ization
Ecology
Life Forms
State Change Dynamics
Plans Goals
Organizational Actions
Types of Organizations
Natural Geography
Human Beings
Plants
Physical Agents
Organizational Plans
Human Activities
Human Organizations
Political Geography
Human Artifacts
Human Anatomy & Physiology
Animals
Agent Organizations
Business & Commerce
Politics Warfare
Nations Governments Geo-Politics
Weather
Products Devices
Conceptual Works
Emotion Perception Belief
Human Behavior & Actions
Social Behavior
Purchasing Shopping
Sports Recreation Entertainment
Professions Occupations
Law
Earth & Solar System
Vehicles Buildings Weapons
Mechanical & Electrical Devices
Software Literature Works of Art
Language
Social Relations, Culture
Social Activities
Transportation & Logistics
Travel Communication
Everyday Living
Business, Military Organizations
Domain-Specific Knowledge
(e.g., Bio-Warfare, Terrorism, Computer Security, Military Tactics, Command & Control, Health Care, ...)
Domain-Specific Facts and Data
CYCORP

CSP/logic: formal verification

Hardware verification
(e.g. IBM, Intel)

Software verification
(small to medium programs)

Most progress in the last 10 years based on encodings into propositional satisfiability (SAT)

Course Map

Dimen-sions / Course Modules	**Deterministic vs. Stochastic**	**Static vs. Sequential**	**States vs. Features vs. Relations**
1. Search	Deterministic	Static	States
2. CSPs	Deterministic	Static	Features
3. Planning	Deterministic	Sequential	States or Features
4. Logic	Deterministic	Static	Relations
5. Uncertainty	Stochastic	Static	Features
6. Decision Theory	Stochastic	Sequential	Features

Reasoning under Uncertainty

Sample application: Microsoft Kinect

- Sensor: IR camera for depth perception sensing projected pattern
- Noise: no fixed reference points; movements in the background

Source: Microsoft & YouTube

Uncertainty/Decision Theory: Diagnosis

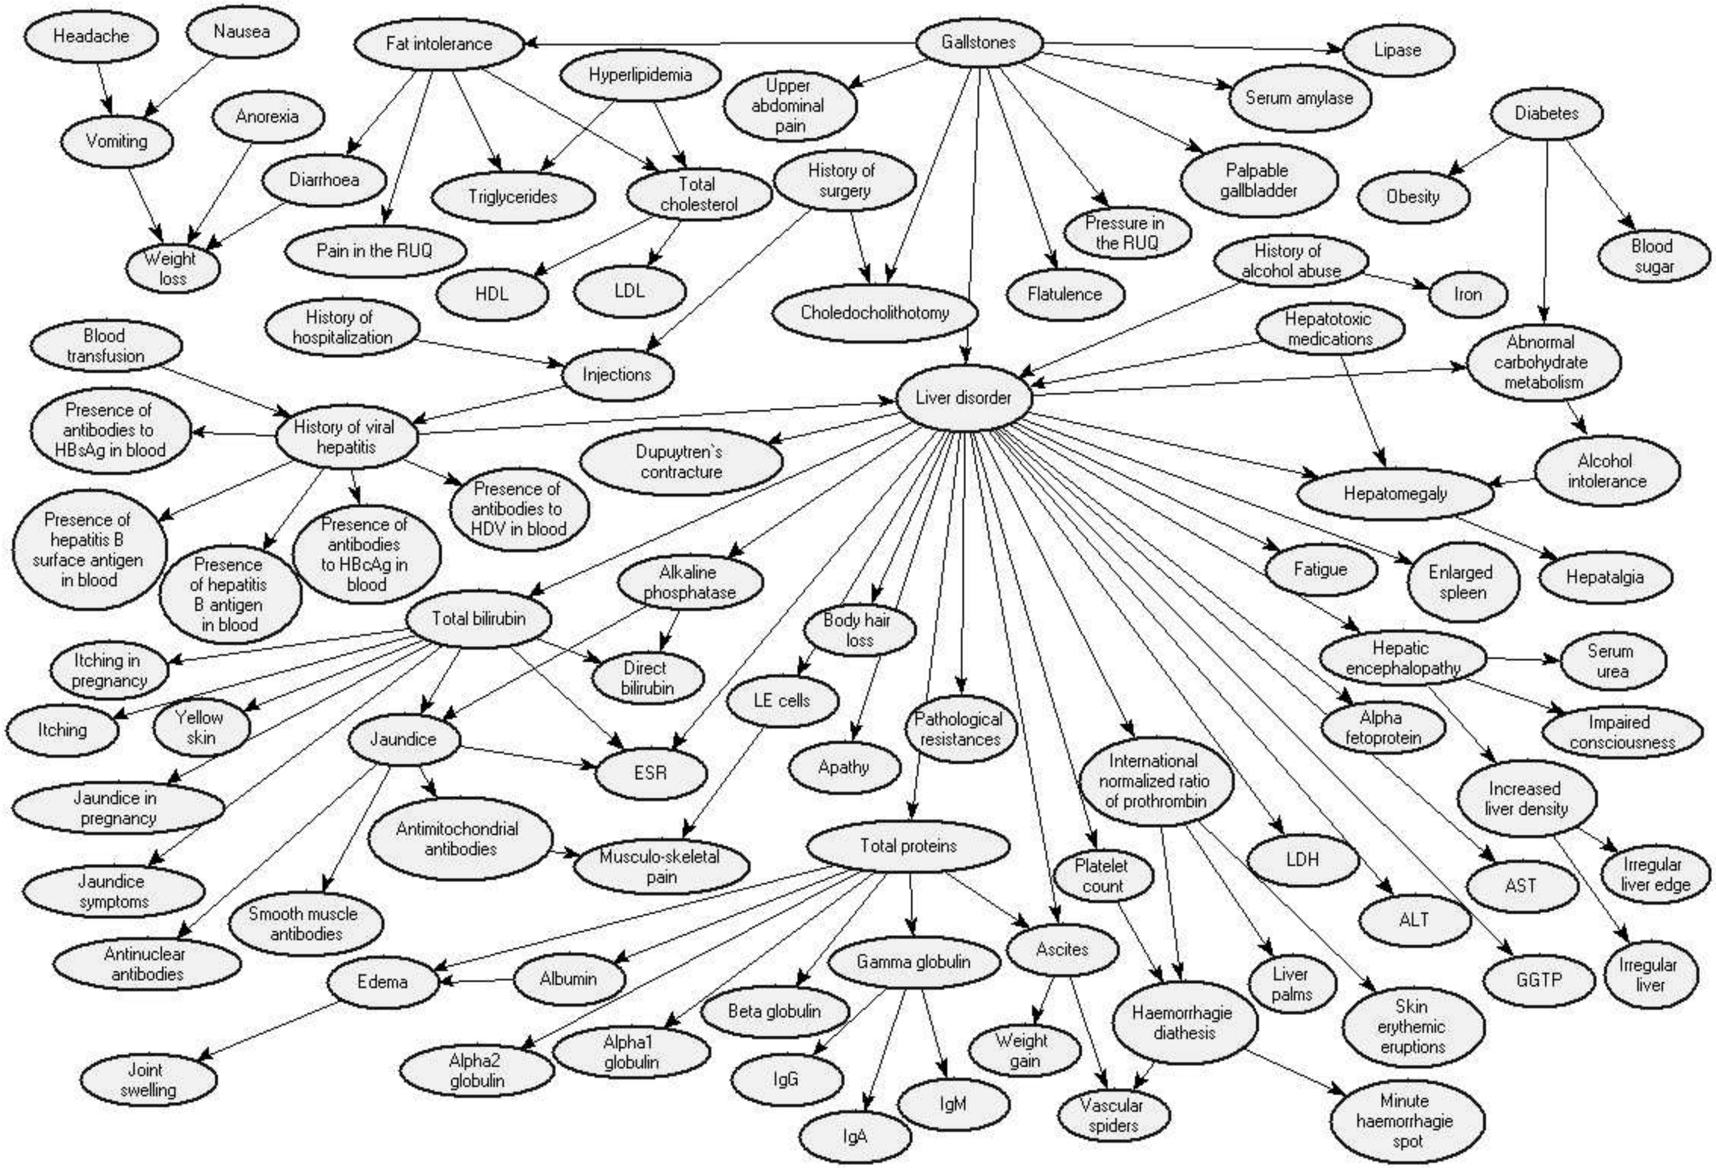

Course Map

Dimen-sions / Course Modules	Deterministic vs. Stochastic	Static vs. Sequential	States vs. Features vs. Relations
1. Search	Deterministic	Static	States
2. CSPs	Deterministic	Static	Features
3. Planning	Deterministic	Sequential	States or Features
4. Logic	Deterministic	Static	Relations
5. Uncertainty	Stochastic	Static	Features
6. Decision Theory	Stochastic	Sequential	Features

Decision Theory: Decision Support Systems

E.g., Computational Sustainability

- New interdisciplinary field, AI is a key component
  - Models and methods for decision making concerning the management and allocation of resources
  - to solve most challenging problems related to sustainability
- Often constraint optimization problems. E.g.
  - Energy: when and where to produce green energy most economically?
  - Which parcels of land to purchase to protect endangered species?
  - Urban planning: how to use budget for best development in 30 years?

Source: http://www.computational-sustainability.org/

Planning Under Uncertainty

Helicopter control: MDP, reinforcement learning

Source: *Andrew Ng*

Planning Under Uncertainty

Autonomous driving:

DARPA Urban Challenge – Stanford's Junior

Source: *Sebastian Thrun*

Planning Under Uncertainty

- Autonomous driving: Dickmanns (1986), Google, Audi, Toyota, Mercedes-Benz, …
- Self-driving cars are now street legal in Florida, California and Nevada.

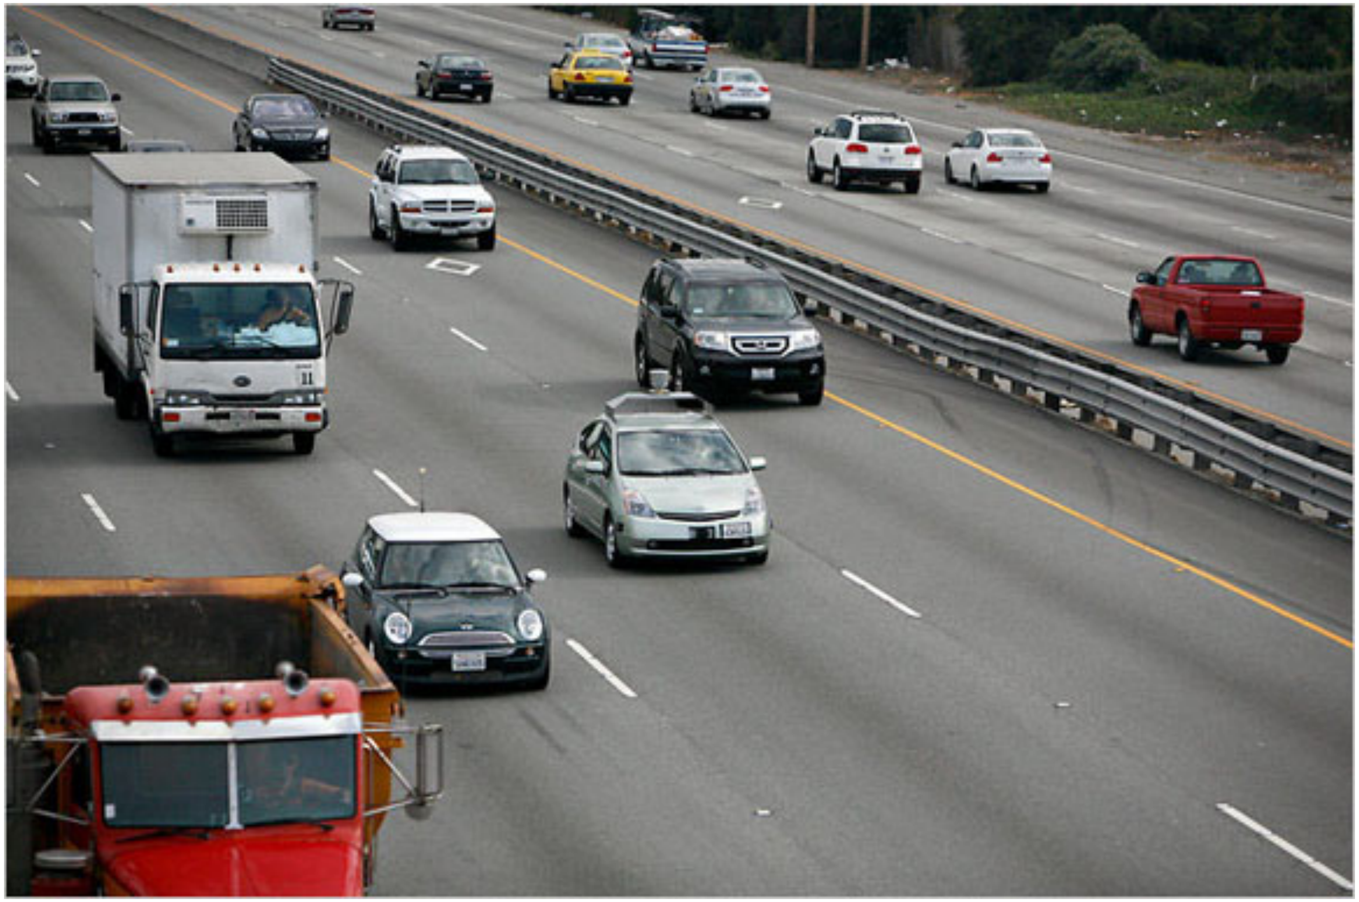

Image source: *geek.com*

Military applications: ethical issues

- Robot soldiers
  - Existing: robot dog carrying heavy materials for soldiers in the field
  - The technology is there
- Unmanned airplanes
- Missile tracking
- Surveillance
- …

Image Source: Boston Dynamics

Multiagent Systems: Robot Soccer

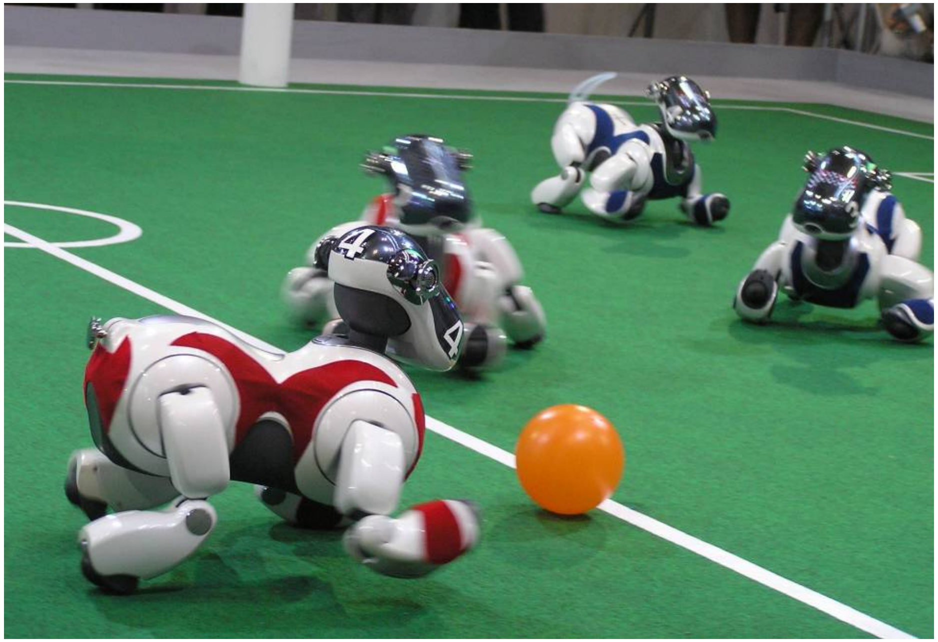

Source: RoboCup web site

RoboCup

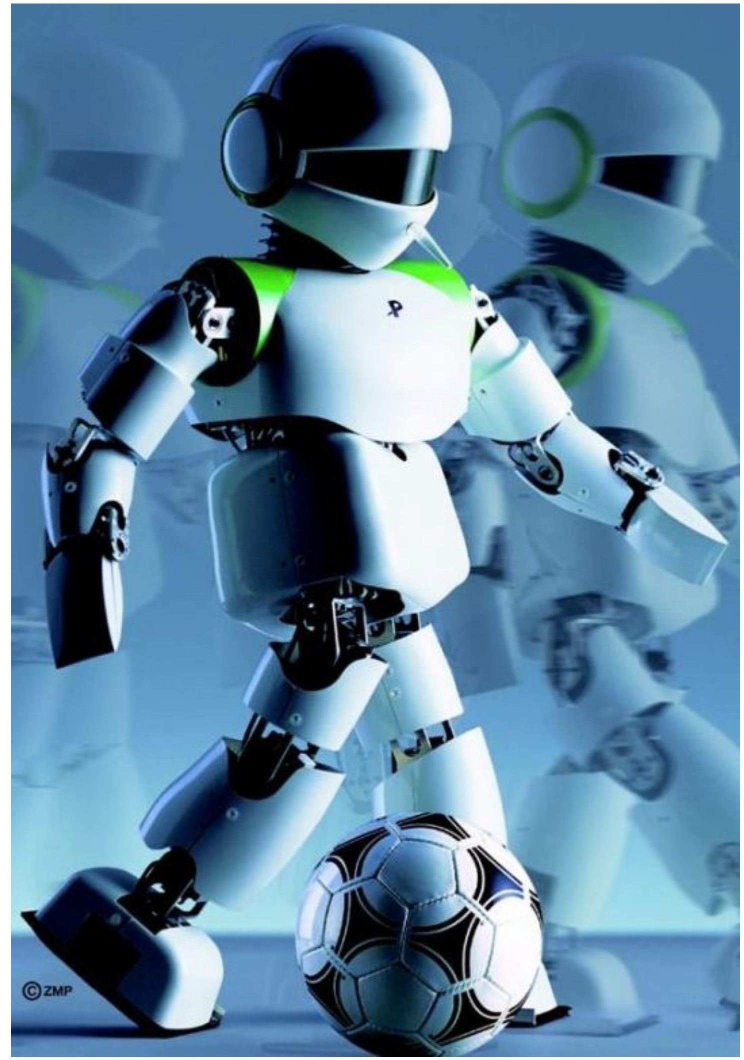

The Dynamites: Two on Two

World's First Soccer Playing Robots (UBC, 1993)

Robot Soccer: Humanoids

Source:
Darmstadt Dribbling Dackels

Summary(1)

We would like most general agents possible, but to start with we need to restrict scope:

4. Flat representations (vs. hierarchical)
5. Knowledge given (vs. knowledge learned)
6. Goals and simple preferences (vs. complex preferences)
7. Single-agent scenarios (vs. multi-agent scenarios)
8. Perfect rationality (vs. bounded rationality)

Extensions we will cover:

1. Deterministic versus stochastic domains
2. Static vs. Sequential domains
3. Representation: Explicit state or features or relations

Summary(2)

- Huge diversity of applications
- More than I could possibly show here
- We shall focus on their common foundations

Coming up …

- For Friday, 1pm: Assignment 0
  - Available on Connect
  - Section 1.5 & 1.6 in the textbook will be particularly helpful
- We'll start the search module: ***read Sections 3.0-3.4***
- Please continue to bring coloured cards (we shall use them next class)