Perspectives and Final Review

CPSC 322 – last class

April 6, 2011

Announcements

- TA evaluations and feedback forms
  - Please fill them out
  - Feedback is very important to improve the course
- We are gathering feedback for the practice exercises
  - You should get (/have gotten) an email about a survey
  - You can earn 10 bonus marks for Assignment 4 by filling in survey
- Final exam is Monday, April 11. DMP 310, 3:30-6pm. Be on time!
  - The list of short questions is online … please use it!
  - Also use the practice exercises (online on course website)
- Remaining office hours this week
  - Thursday: Frank (11-12, X530), Vasanth (3-5pm, learning center)
  - Friday: Mike (10-12am, learning center), Frank (3-5, DMP 110)
- Optional Rainbow Robot tournament: this Friday, 3-4pm
  - Vasanth will run the tournament in DMP110,
    and I will answer questions on the side (office hours)

Course Overview

Course Module

Representation

Reasoning Technique

Environment

Problem Type

		Deterministic	Stochastic	
Static	Constraint Satisfaction	***Variables + Constraints*** Arc Consistency Search		
Static	Logic	***Logics*** Search	***Bayesian Networks*** Variable Elimination	Uncertainty
Sequential	Planning	***STRIPS*** Search As CSP (using arc consistency)	***Decision Networks*** Variable Elimination ***Markov Processes*** Value Iteration	Decision Theory

Big Picture: Heuristic Search

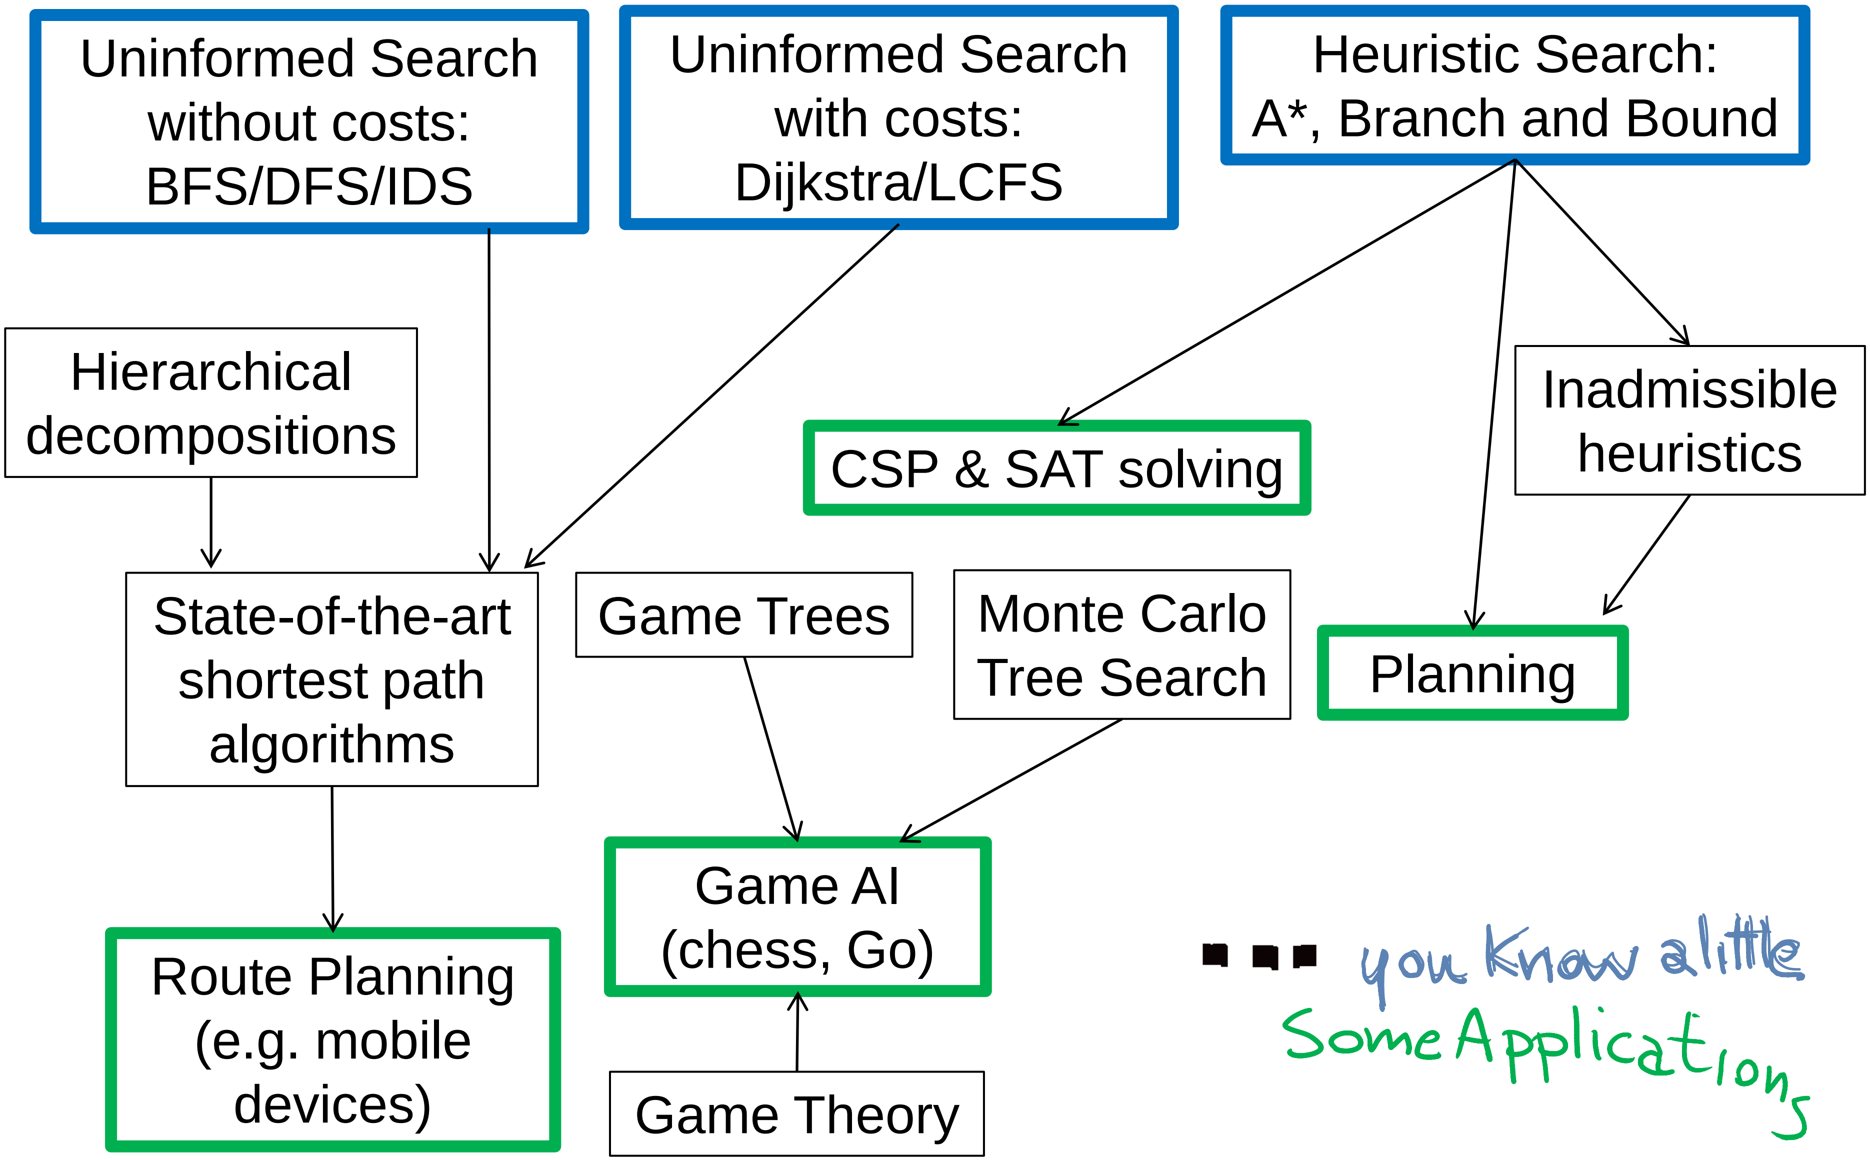

Game Tree Search

- Directed graph
  - Two types of nodes: 1) I choose. 2) opponent chooses

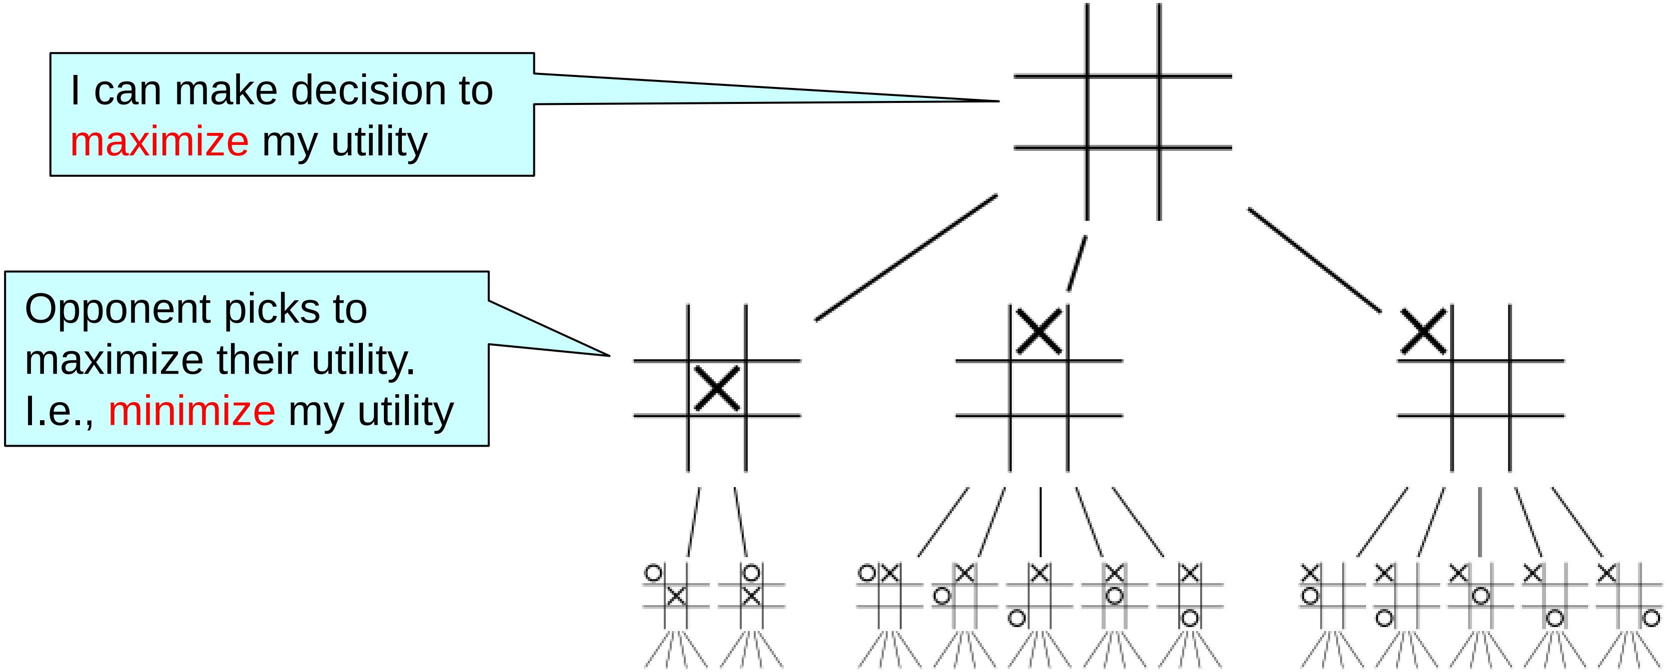

  - Minimax tree search
    - Best strategy: make decision such that utility is maximized across all opponent moves
    - Look forward to depth d, estimate utility of reached state
      - Human grandmasters in chess: up to 9 levels
      - Chess computers: up to 13 levels

Big Picture: Constraint Satisfaction Problems

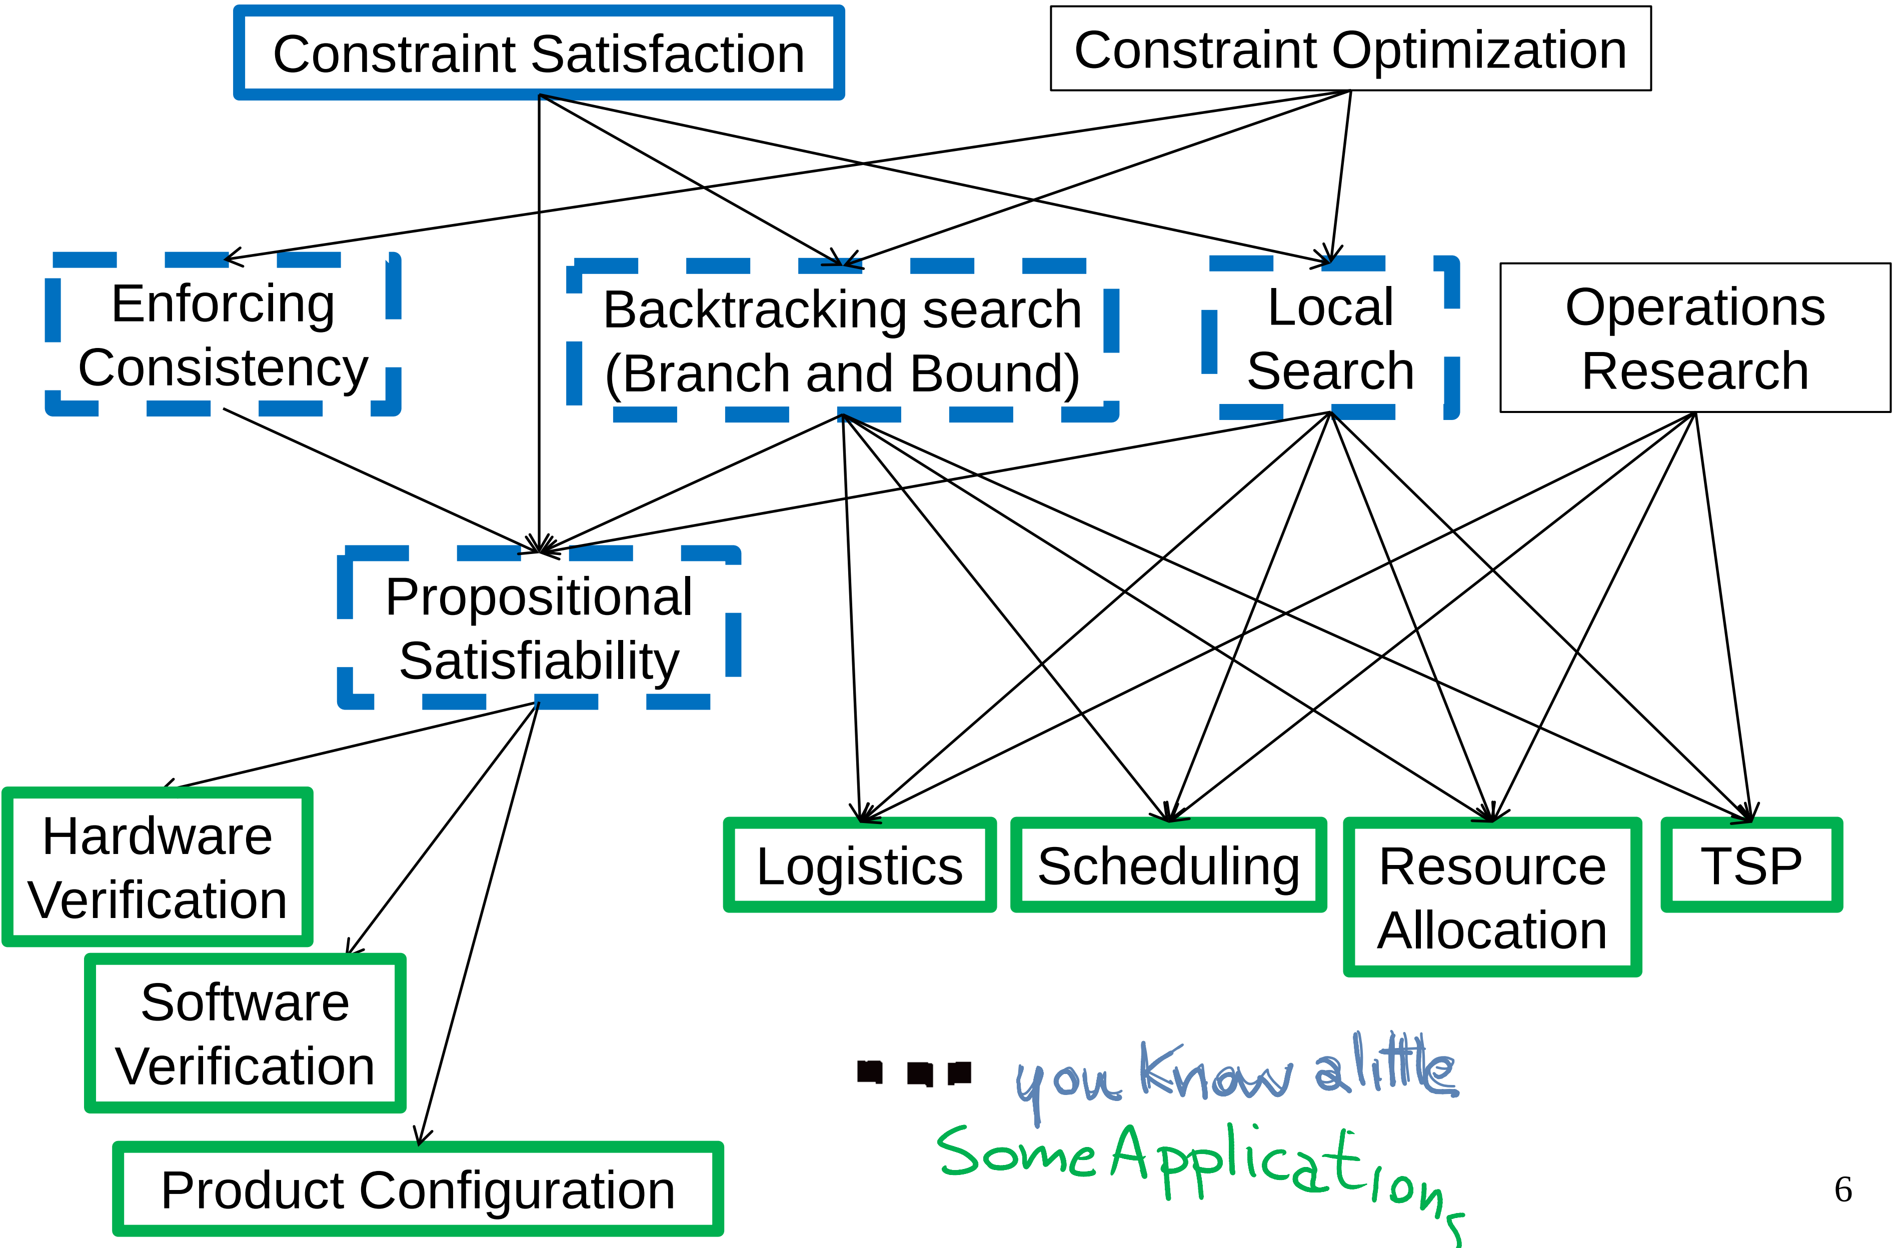

Constraint optimization problems

- 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, $x_i \in \mathbb{Z}$ (integers)
  - NP-hard, widely used in operations research and in industry

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

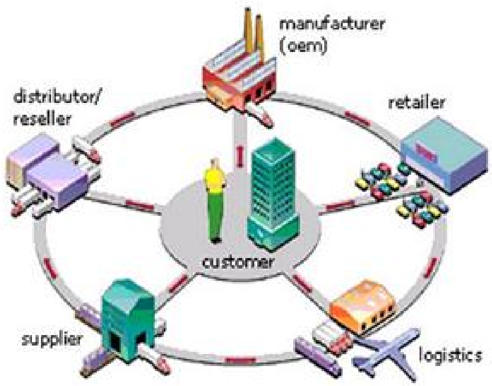

Supply chain
management
software
Oracle,
SAP,…

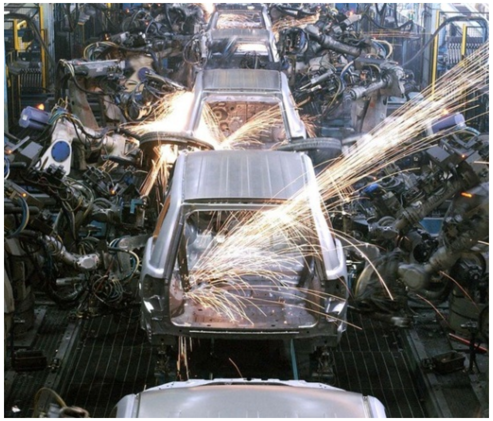

Production planning
and optimization
Airbus, Dell, Porsche,
Thyssen Krupp,
Toyota, Nissan, ...

Big picture: Deterministic Planning

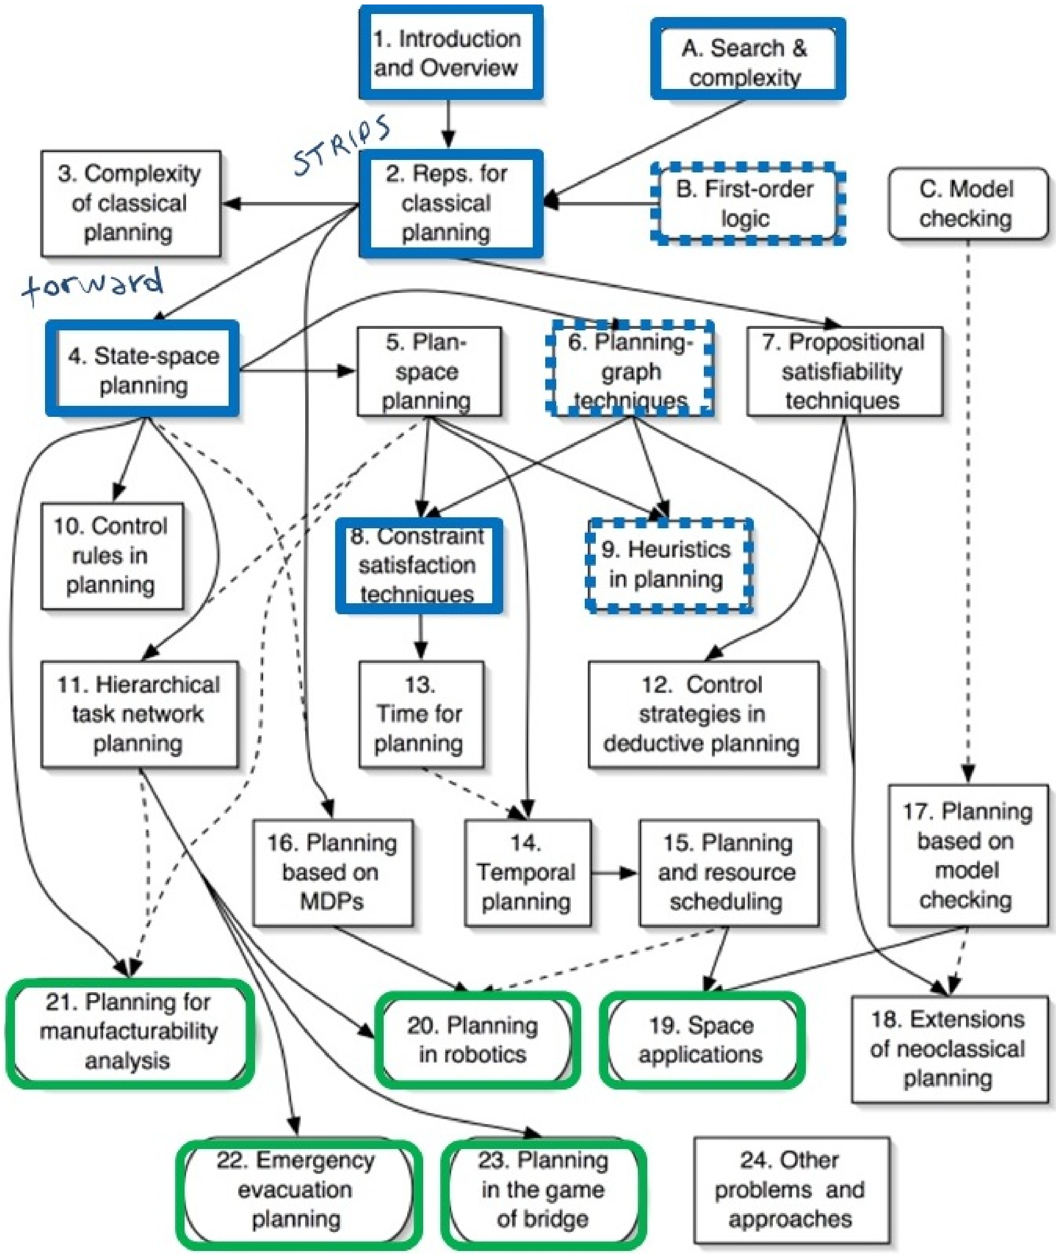

- Chapter overview of book: ***Automated Planning: Theory and Practice***
- By Ghallab, Nau, and Traverso. Web site:
  - http://www.laas.fr/planning
  - Also has lecture notes

Big Picture: Logics in AI

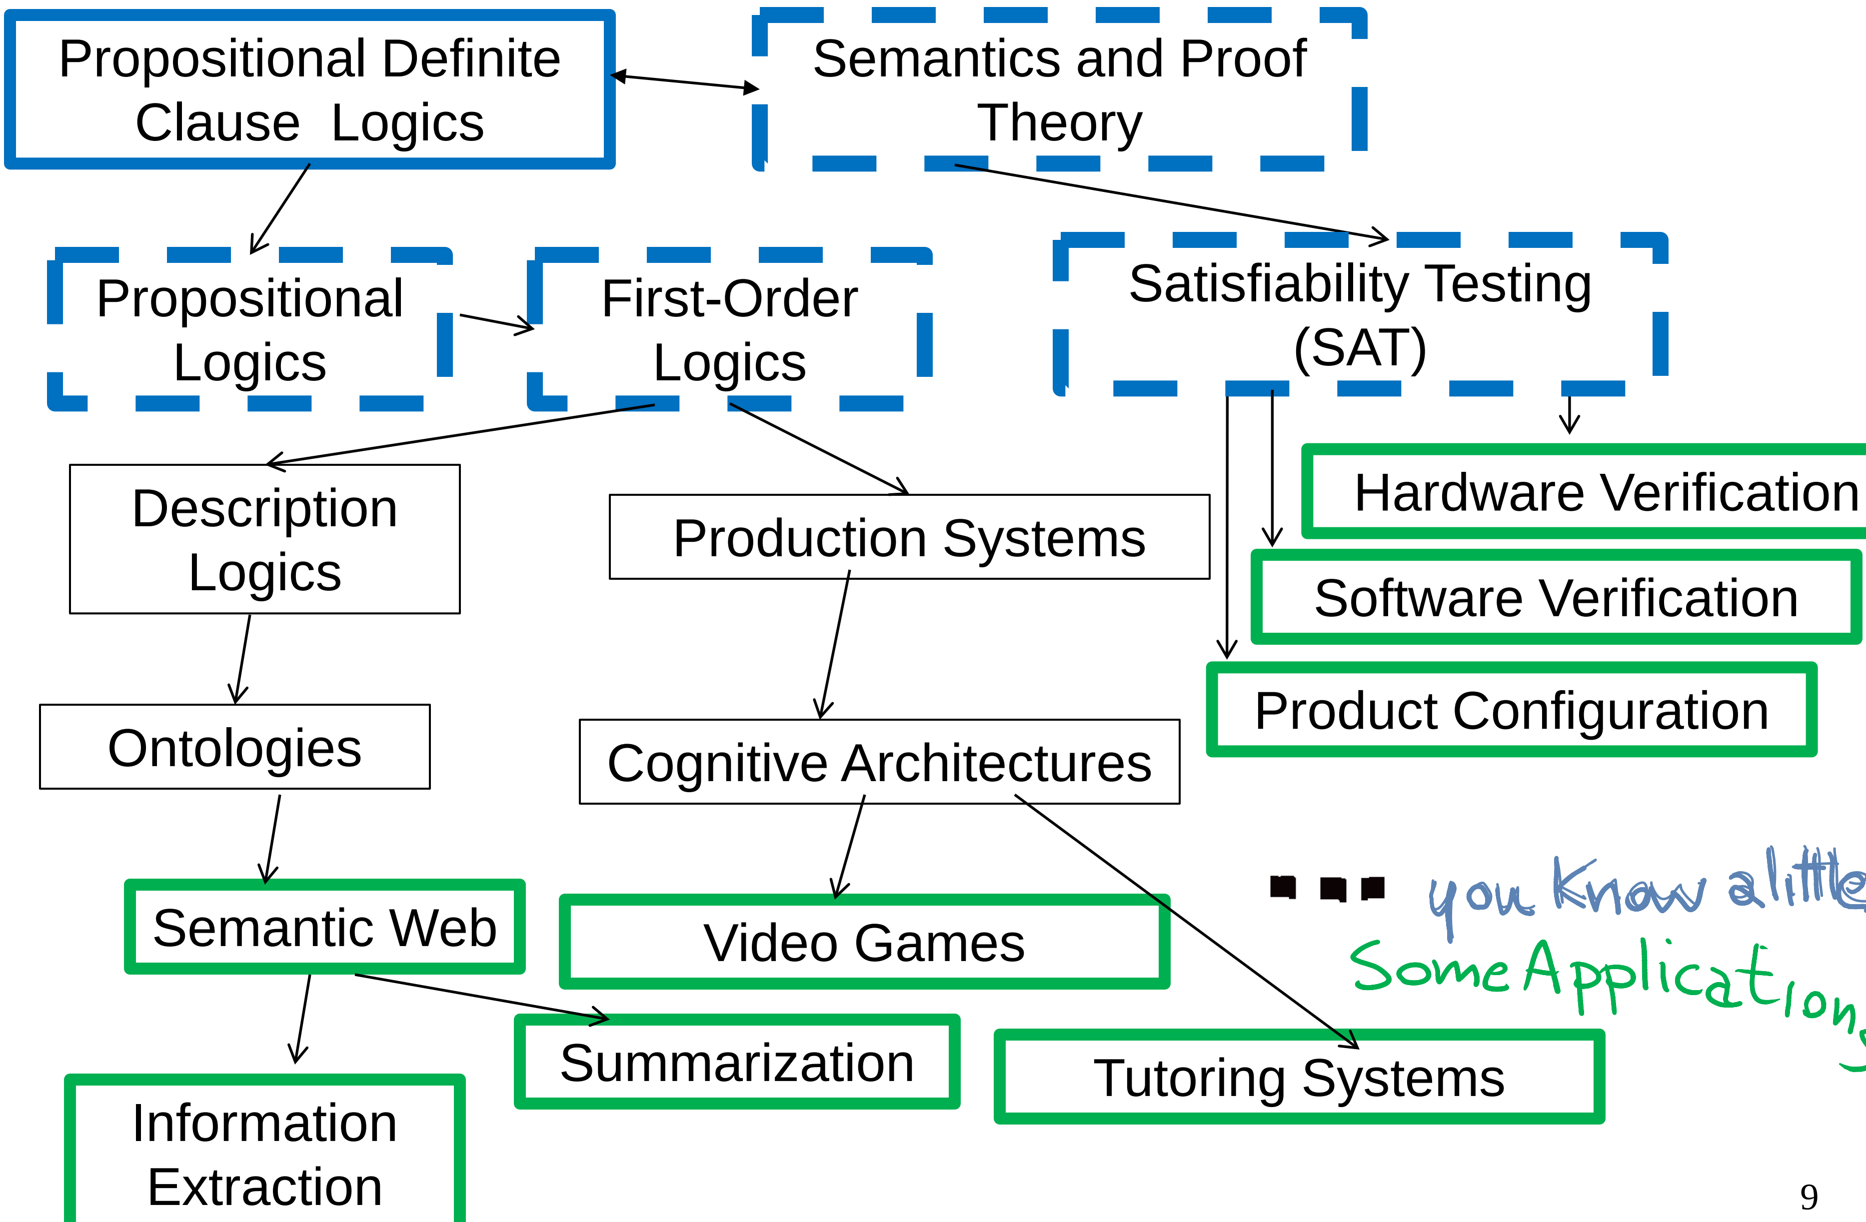

CSP/logic: formal verification

Hardware verification
(e.g., IBM)

Software verification
(small to medium programs)

Most progress in the last 10 years based on:

- Encodings into propositional satisfiability (SAT)
- Best methods: extensions of arc consistency + domain splitting

Big picture: Reasoning Under Uncertainty

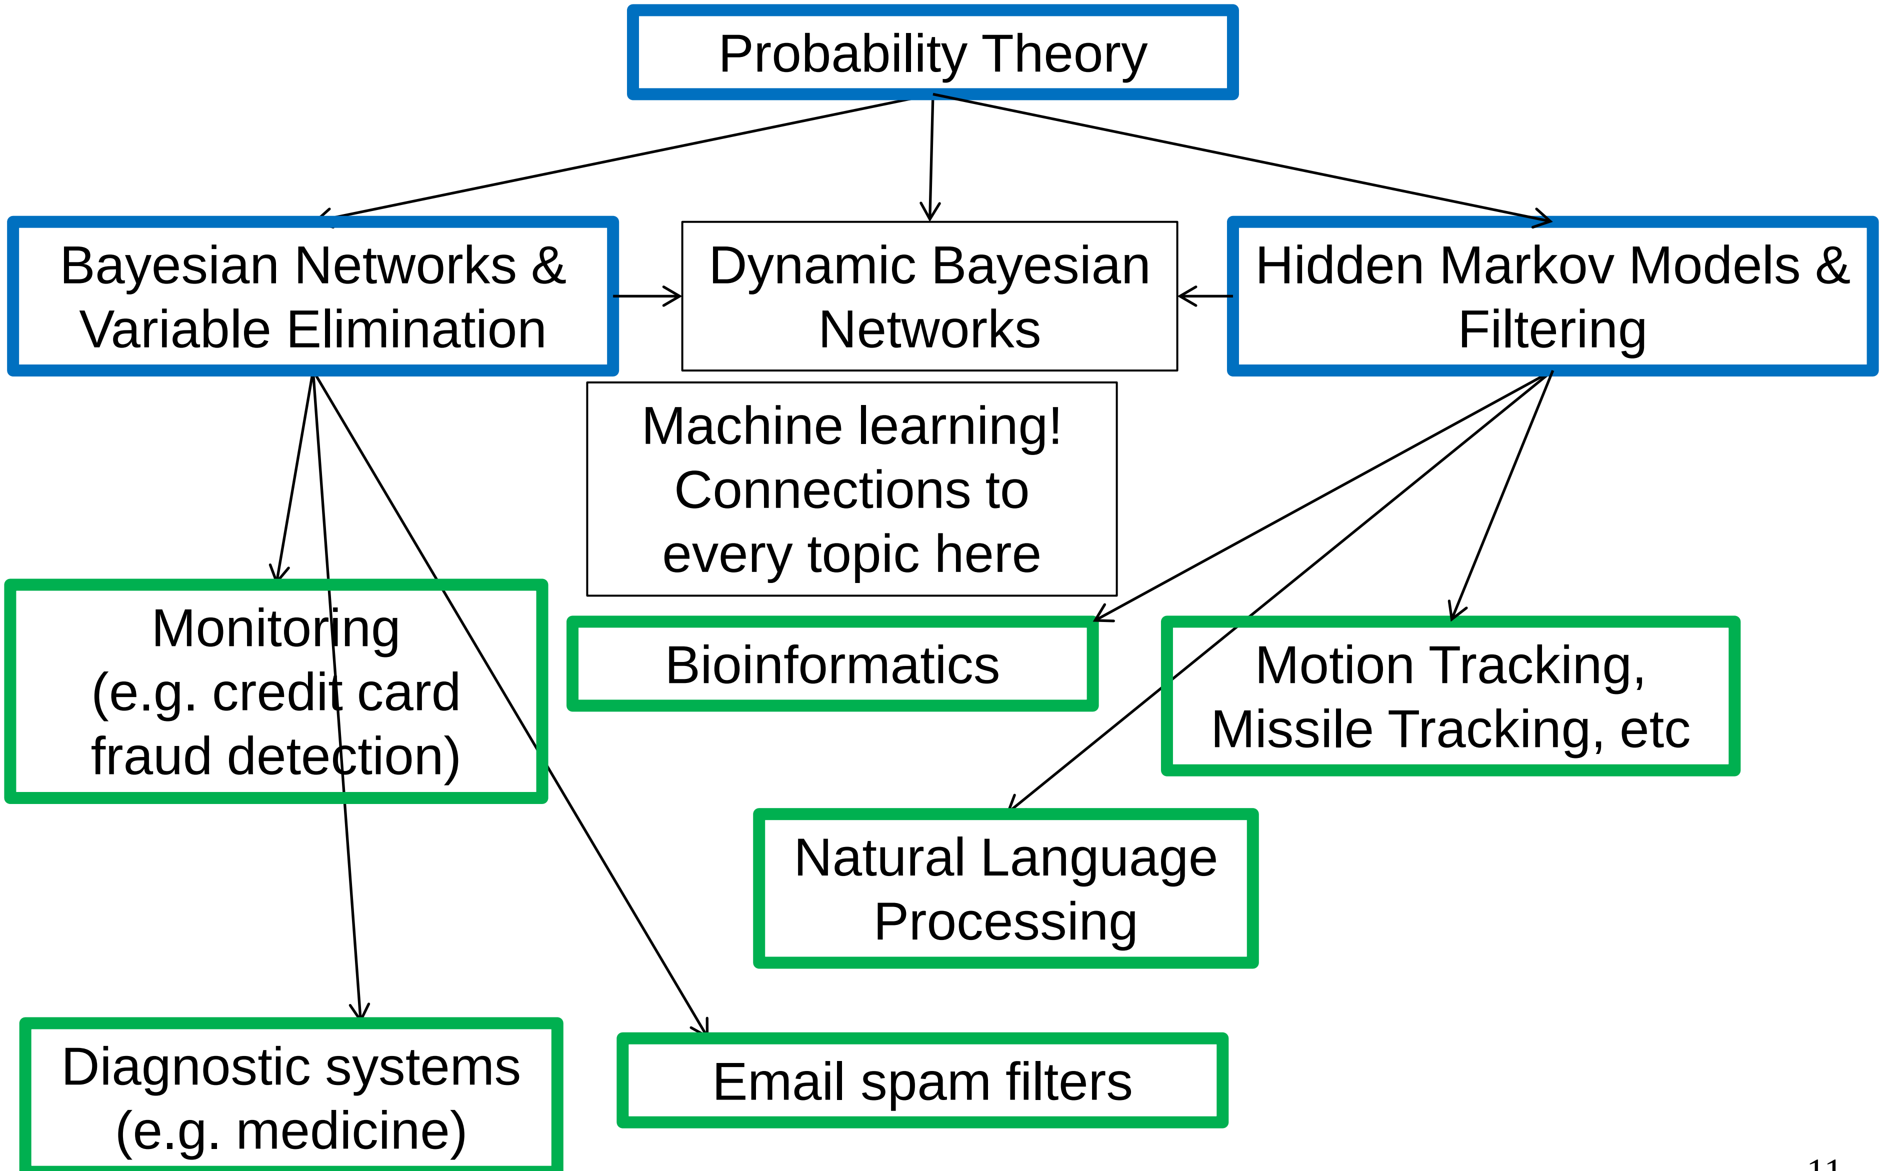

Reasoning Under Uncertainty

E.g. motion tracking: track a hand and estimate activity:

– drawing, erasing/shading, other

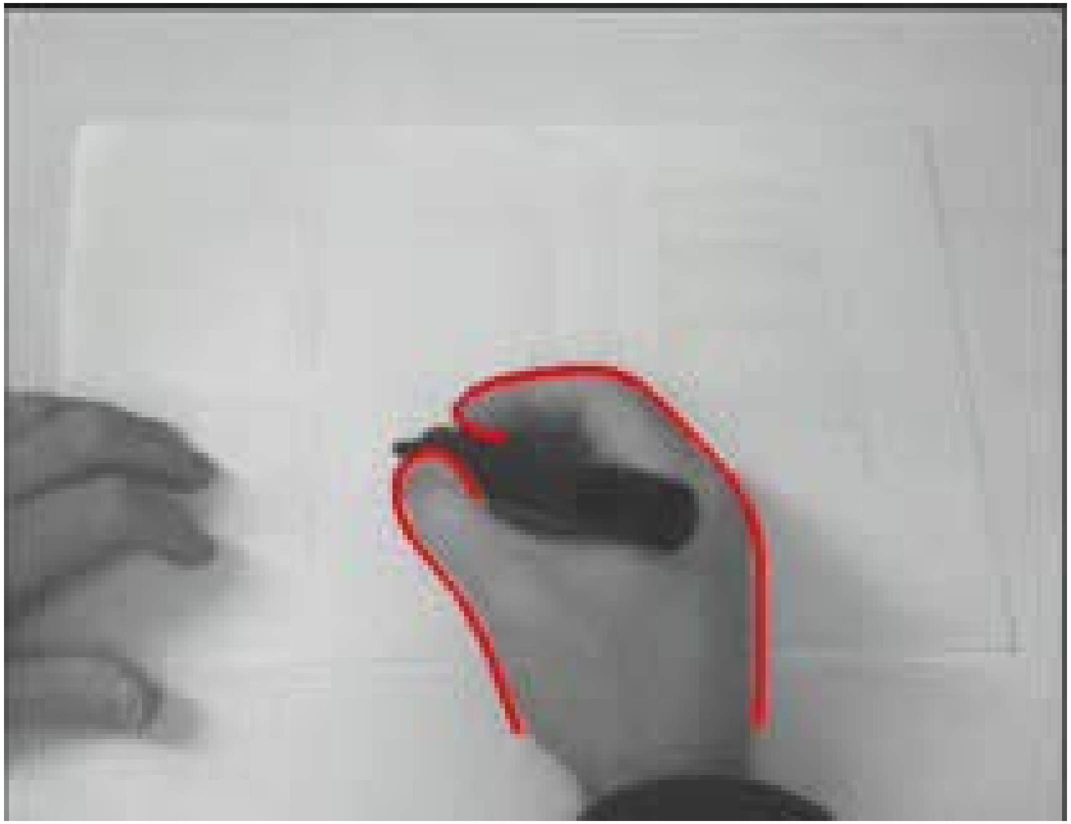

Source: *Kevin Murphy, UBC*

– Machine Learning is all about reasoning under uncertainty: e.g., CPSC 340

Big Picture: Planning under Uncertainty

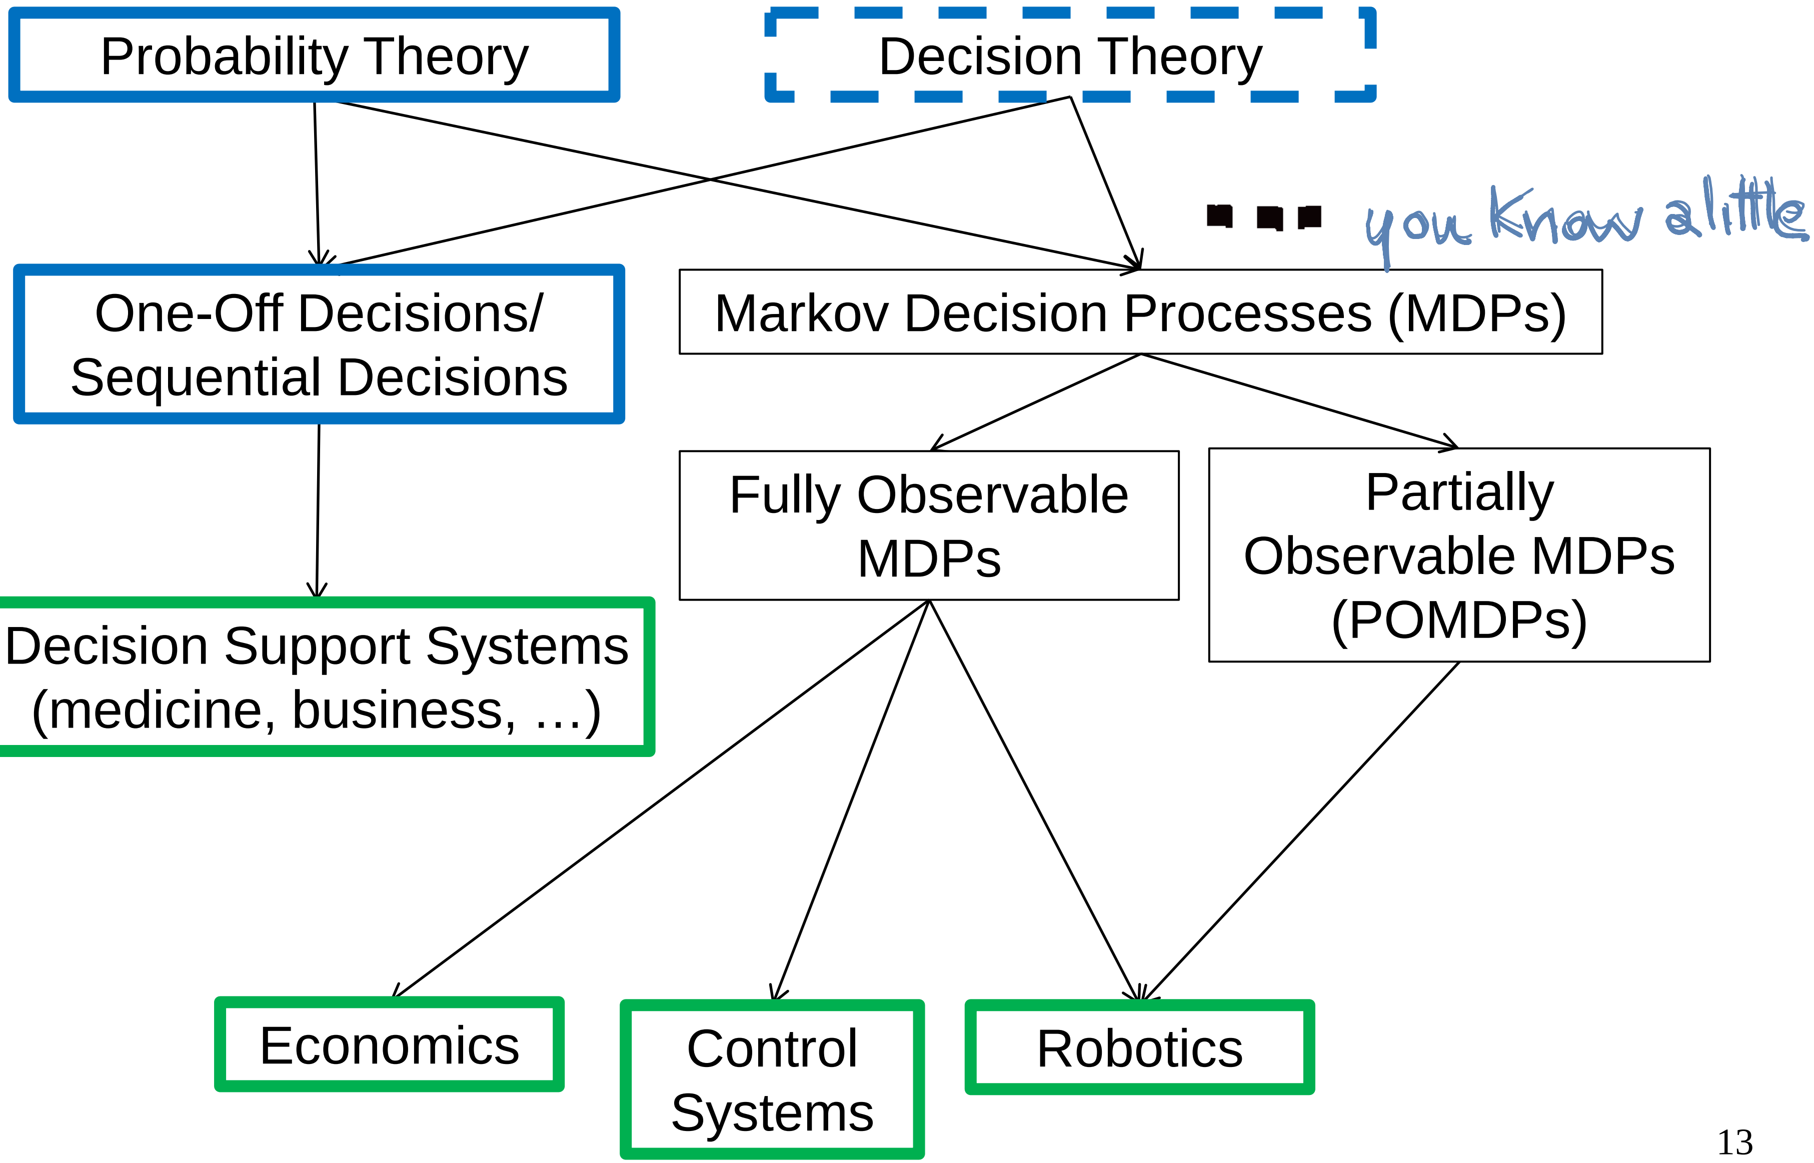

Decision Theory: Decision Support Systems

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 are 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/

- CPSC422 is all about uncertainty and decision theory

Dimensions of Representational Complexity

We 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
   - Solve subproblems of smaller size
5. Knowledge given vs.
   knowledge learned from experience
   - Machine learning. Learn e.g. Bayesian network structure & probabilities
6. Goals vs. complex preferences
   - Multi-objective optimization is an active research area
7. Single-agent vs. multi-agent
   - Game theory
8. Perfect rationality vs. bounded rationality
   - Typically only limited time available, so we need approximation algorithms

Decisions under uncertainty in a multi-agent setting: robot soccer

Source:
Darmstadt Dribbling Dackels

If you are considering grad school …

- Consider applying for NSERC funding this coming September!
  - You apply in September for the school year starting the September after (or even the January 1.5 years after)
  - Writing a good application takes a while (so do it early)
  - Funding helps
    - Increases your chances of getting accepted
    - Gives you more flexibility
    - You typically earn more if you are an a scholarship

Review for Final

- Primarily worked examples on the whiteboard
- Review will focus on uncertainty and decision theory
  - This is by popular demand
  - But keep in mind:
    - Focus of final exam will be on
      - Planning
      - Logic
      - Uncertainty and decision theory
- See practice exercises on the course website
  - These can be very useful for studying

Review: conditional independence

- Two variables X and Y are conditionally independent given a set of observed variables E, if and only if
  - There is no path along which information can flow from X to Y
  - Information can flow along a path if it can flow through all the nodes in the path.
    - Note: observation status of A and C does not matter

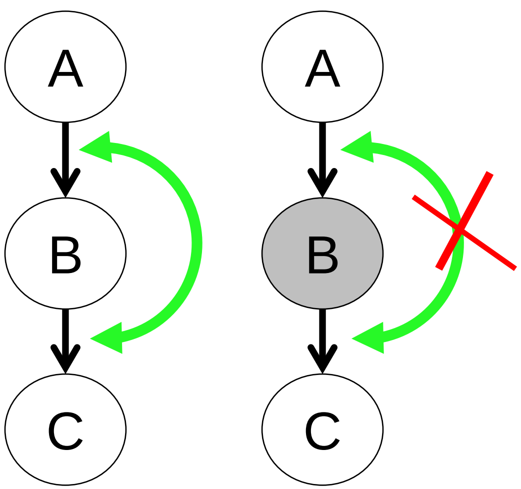

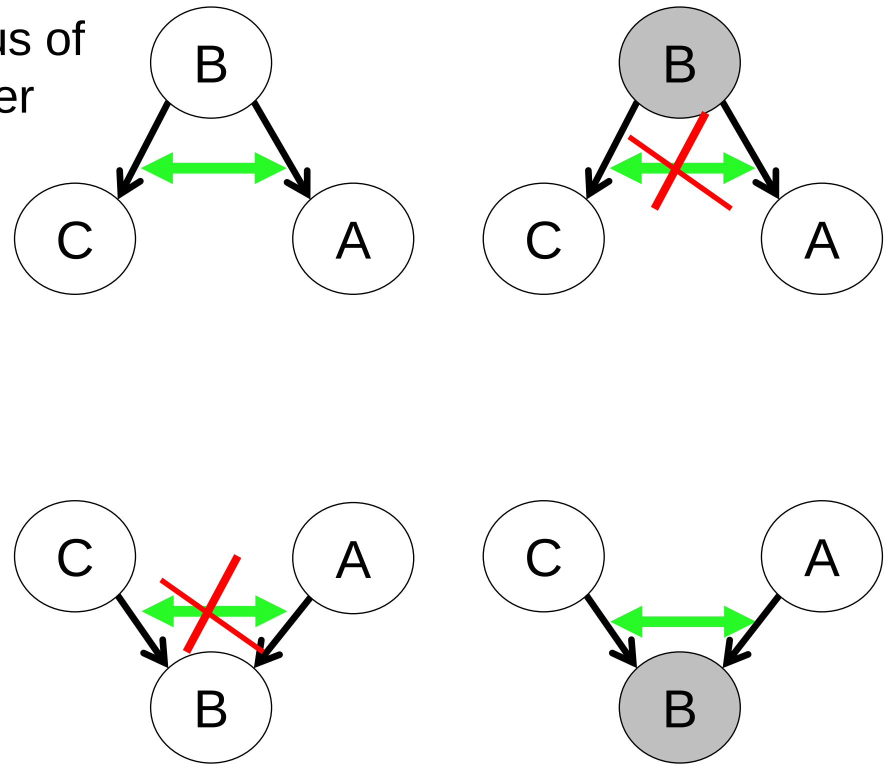

Review: Variable Elimination (VE) in BNs

- The joint probability distribution of a Bayesian network is $P(X_1, \ldots, X_n) = \prod_{i=1}^{n} P(X_i|pa(X_i))$
  - We make a factor f_i for each conditional probability table $P(X_i|pa(X_i))$
  - So we have $P(X_1, \ldots, X_n) = \prod_{i=1}^{n} f_i$

- The variable elimination algorithm computes $P(Y|\,E=e)$ as $\sum_{Z_1,\ldots,Z_k} \prod_{i=1}^{n} P(X_i|pa(X_i), E=e)$:
  - where $Z_1, \ldots, Z_k$ are all non-query and non-evidence variables

  1. Create a factor for each conditional probability
  2. Assign *E*=*e*, one evidence variable at a time
  3. Sum out all non-query variables $Z_1, \ldots, Z_k$, one at a time
     - To sum out Z_i:
       - Multiply factors containing Z_i
       - Then marginalize out Z_i from the product
  4. Normalize the final factor f(Y).
     - The resulting factor is exactly $P(Y|\,E=e)$

Review: Variable Elimination for single decisions

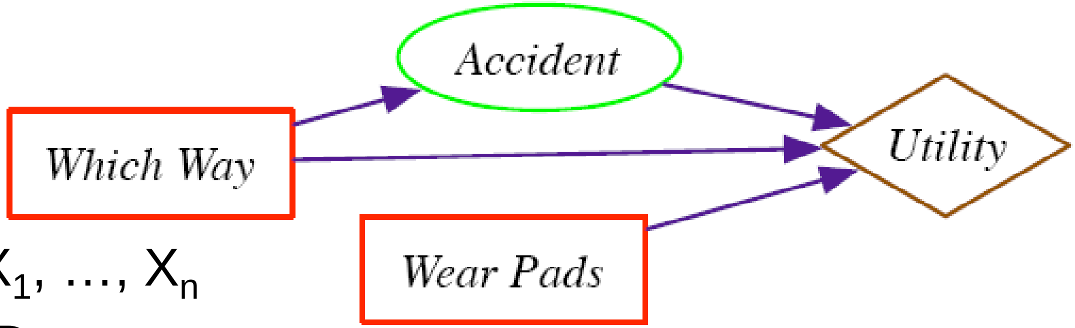

- Denote
  - the random variables as $X_1, \ldots, X_n$
  - the decision variables as D
  - the parents of node N as pa(N)

$$E(U) = \sum_{X_1,\ldots,X_n} P(X_1,\ldots,X_n \mid D)\, U(pa(U))$$
$$= \sum_{X_1,\ldots,X_n} \prod_{i=1}^{n} P(X_i \mid pa(X_i))\, U(pa(U))$$

- The variable elimination algorithm computes optimal decision:
  1. Create a factor for each conditional probability and for the utility
  2. Sum out all random variables, one at a time
     - This creates a factor on D that gives the expected utility for each d_i
  3. Choose the d_i with the maximum value in the factor

Review: Variable Elimination for sequential decisions

1. Create a factor for each CPT and a factor for the utility
2. While there are still decision variables
   - 2a: Sum out random variables that are not parents of a decision node.
     - E.g Tampering, Fire, Alarm, Smoke, Leaving
   - 2b: Max out last decision variable D in the total ordering
     - Keep track of decision function
3. Sum out any remaining variables:
   result is the expected utility of the optimal policy.

$$E[\pi] = E[U|\pi] = \sum_{w \models \pi} P(w)\, U(w)$$

$$= \sum_{X_1,\ldots,X_n,D_1\ldots,D_m} \prod_{i=1}^{n} P(X_i|pa(X_i)) \prod_{j=1}^{m} (\delta_j(pa(D_j)) = D_j)\ U(pa(U))$$

This term is zero if D_j's value does not agree with what the policy dictates given D_j's parents.

Announcements

- TA evaluations and feedback forms
  - Please fill them out
  - Feedback is very important to improve the course
- We are gathering feedback for the practice exercises
  - You should get (/have gotten) an email about a survey
  - You can earn 10 bonus marks for Assignment 4 by filling in survey
- Final exam is Monday, April 11. DMP 310, 3:30-6pm. Be on time!
  - The list of short questions is online … please use it!
  - Also use the practice exercises (online on course website)
- Remaining office hours this week
  - Thursday: Frank (11-12, X530), Vasanth (3-5pm, learning center)
  - Friday: Mike (10-12am, learning center), Frank (3-5, DMP 110)
- Optional Rainbow Robot tournament: this Friday, 3-4pm
  - Vasanth will run the tournament in DMP110,
    and I will answer questions on the side (office hours)