Midterm Study Guide

What Will it Be Like?

- Broadly:
  - Turn a problem description into a solution
  - Work through a problem to reach a solution
  - Demonstrate a conceptual grasp of the material
- Concepts ←→ Algorithms and Implementations
- Basic idea: you need to **understand the ideas** behind the material we have covered, and be **able to apply it** to solving problems.

What Kind of Questions?

- T/F, multiple choice, fill in the blank
- Write definitions of terms
- Work through an {algorithm|solution type|problem}
- Draw something – search trees, states, Bayes nets, paths through a map, …
- Write a **short** answer to English questions
  - E.g.: What approach would you use to solve this problem?
  - E.g.: "We know these are independent. Why?"

What Do I Need To Do?

- Homeworks **and lectures** should be good practice
- Coding questions (not minor syntax mistakes, etc.)
  - We're looking for "I understand this well enough to implement it," not "I know Python really well"
- Look at **homeworks**, **sample problems** in lectures, and **class exercises**
- Look at lectures' "Why?" questions.

Scoring

- Follow directions.
- Start with a perfect score, mark down for mistakes
  - If I ask for 2 examples, and you give 3, one of which is wrong, it's $-1/2$, not $-1/3$
- Read carefully.
  - You have time.
  - "I didn't see the part that said…"
- Ask for clarification on, e.g., unfamiliar words

Topics: AI

- What is intelligence?
- What is AI?
- What is it used for? Good for?
- Historical events and figures

Topics: Agents

- Agents
  - What kinds are there?
  - What do they do?
  - How do we characterize them (what traits do they have)?
    - Autonomy, rationality, ..?
  - How do they interact with an environment?
- Environments
  - What's an environment?
  - How is it characterized?

Topics: Search

- What is it for?
- Elements of a search problem
  - State spaces, actions, costs, …
  - How do state spaces pertain to search?
  - To problem-solving?
- Exploring search space (selecting, expanding, generating, evaluating)
- Specific algorithms: How do they work? What are they good for? What are their weaknesses?

Topics: Formalizing Search

- What are the elements of a search problem?
  - "Express [*X*] as a search problem." What does that mean?
- **States:** every state a puzzle can be in
- **Actions/Operations:** how you get between states
- **Solutions:** you need a goal test (and sometimes a *heuristic*, or estimate of distance from goal)
  - Sometimes we care about path (planning), sometimes just goal (identification). Can you say which, for a given problem?
- **Costs:** not all solutions or actions are equal

Topics: Uninformed Search

- Why do uninformed search?
- Come up with some examples of uninformed search problems
- Important algorithms: BFS, DFS, iterative deepening, uniform cost
- A (very) likely question: "What would be the best choice of search method for [*problem*], **and why**?"
- Characteristics of algorithms
  - Completeness, optimality, time and space complexity, …

Topics: Informed Search

- Some external or pre-existing information says what part of state space is **more likely** to have a solution
- **Heuristics** encode this information: $h(n)$
  - Pop quiz: What does $h(n) = 0$ mean?
- Admissibility & Optimality
  - Some algorithms can be optimal when using an admissible heuristic

A heuristic applies to a **node**, can give optimal solution with the right **algorithm**

- Algorithms: best-first, greedy search, A*, IDA*, SMA*
- What's a good heuristic for a problem? Why?

Sample Problem

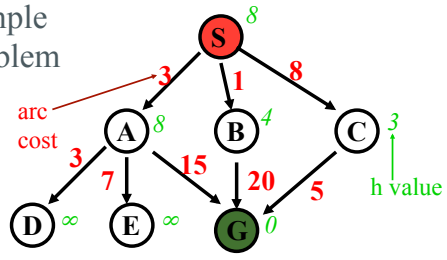

Apply the following to search this space. At each search step, show: current node being expanded, $g(n)$ (path cost so far), $h(n)$ (heuristic estimate), $f(n)$ (evaluation function), and $h^*(n)$ (true goal distance).

Depth-first search
Breadth-first search
A* search
Uniform-cost search
Greedy search

Topics: Local Search

- Idea: Keep a single "current" state, try to improve it
  - Don't keep path to goal
  - Don't keep entire search in memory
  - Go to "successor states"
- Concepts: hill climbing, local maxima/minima, random restarts
- Important algorithms: hill climbing, local beam search, simulated annealing

HW2

- How many states are there?
- What operations fully encode this search problem?
  - That is: how can you reach every state?
- Are there loops?
- How many states does pure DFS visit?
  - If there are loops?
- What's a good algorithm? A bad one?

Topics: CSPs

- Constraint Satisfaction: subset of search problems
  1. **State** is defined by **random variables** X_i
  2. With values from a **domain** D
  3. Knowledge about problem can be expressed as **constraints** on what values X_i can take
- Special algorithms, esp. on constraint networks
- **How would you express [X] as a CSP?** As a search? How would you represent the constraints?
  - E.g.: "Must be alphabetical"

Topics: Constraint Networks

- **Constraint propagation**: constraints can be propagated through a constraint network
  - Goal: maintain **consistency** (constraints aren't violated)
- Concepts: Variable ordering, value ordering, fail-first
- Important algorithms:

1. **Backtracking:** DFS, but at each point:
   - Only consider a single variable
   - Only allow legal assignments
2. **Forward checking**

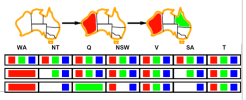

Topics: Games

- Why play games? What games, and how?
- Characteristics: zero-sum, deterministic, perfect information (or not)
- What's the search tree for a game? How big is it?
- How would you express game [X] as a search? What are the states, actions, etc.? How would you solve it?
- Algorithms: (expecti)minimax, alpha-beta pruning
  - Many examples on slides

Topics: Basic Probability

- What is uncertainty?
- What are sources of uncertainty in a problem?
  - Non-deterministic, partially observable, noisy observations, noisy reasoning, uncertain cause/effect world model, continuous problem spaces…
- World of all possible states: a complete assignment of values to random variables
- Joint probability, conditional probability

Topics: Basic Probability

- Independence: A and B are independent
  - $P(A) \perp\!\!\!\perp P(B)$ iff $P(A \wedge B) = P(A)\,P(B)$
  - A and B do not affect each other's probability
- Conditional independence: A and B are independent *given C*
  - $P(A \wedge B \mid C) = P(A \mid C)\,P(B \mid C)$
  - A and B don't affect each other **if C is known**

Topics: Basic Probability

- $P(a \mid b) = \frac{P(a \wedge b)}{P(b)}$
- $P(a \wedge b) = P(a \mid b)\,P(b)$

P(*smart* ∧ *study* ∧ *prep*)	*smart*		¬*smart*	
	study	¬*study*	*study*	¬*study*
prepared	.432	.16	.084	.008
¬*prepared*	.048	.16	.036	.072

- What is the prior probability of smart?
- What is the conditional probability of prepared, given study and smart?
- Is prepared independent of study?

Topics: Probabilistic Reasoning

- Concepts:
  - Posteriors and Priors; Bayesian Reasoning; Induction and Deduction; Probabilities of Events
  - [In]dependence, conditionality, marginalization
- What is Bayes' Rule and what is it useful for?

$$P(H_i \mid E_j) = \frac{P(E_j \mid H_i)P(H_i)}{P(E_j)}$$

$$P(cause \mid effect) = \frac{P(effect \mid cause)P(cause)}{P(effect)}$$

Topics: Joint Probability

What is the **joint probability** of A and B?

- $P(A,B)$
- The probability of any set of legal assignments.
- Booleans: expressed as a matrix/table

	alarm	¬alarm
burglary	0.09	0.01
¬burglary	0.1	0.8

≍

A	B	
T	T	0.09
T	F	0.1
F	T	0.01
F	F	0.8

- Continuous domains → probability functions

22

Conditional Probability Tables

- For X_i, CPD $P(X_i \mid Parents(X_i))$ quantifies effect of parents on X_i
- **Parameters** are probabilities in conditional probability tables (CPTs):

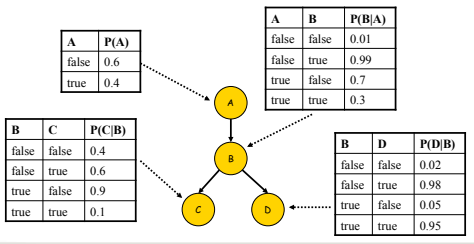

A	P(A)
false	0.6
true	0.4

A	B	P(B|A)
false	false	0.01
false	true	0.99
true	false	0.7
true	true	0.3

B	C	P(C|B)
false	false	0.4
false	true	0.6
true	false	0.9
true	true	0.1

B	D	P(D|B)
false	false	0.02
false	true	0.98
true	false	0.05
true	true	0.95

Example from web.engr.oregonstate.edu/~wong/slides/BayesianNetworksTutorial.ppt

BBN Definition

- AKA Bayesian Network, Bayes Net
- A graphical model (as a DAG) of probabilistic relationships among a set of random variables
- Links represent direct influence of one variable on another

Slides from Dr. Oates, UMBC

Simple Bayesian Network

$S \in \{no, light, heavy\}$ → Smoking → Cancer

$C \in \{none, benign, malignant\}$

P(S=no)	0.80
P(S=light)	0.15
P(S=heavy)	0.05

Smoking=	no	light	heavy
P(C=none)	0.96	0.88	0.60
P(C=benign)	0.03	0.08	0.25
P(C=malig)	0.01	0.04	0.15

Slides from Dr. Oates, UMBC

More Complex Bayesian Network

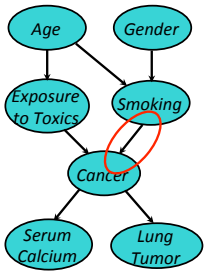

Slides from Dr. Oates, UMBC

More Complex Bayesian Network

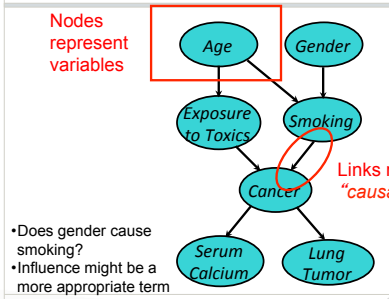

- Does gender cause smoking?
- Influence might be a more appropriate term

Slides from Dr. Oates, UMBC

Independence

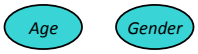

Age and *Gender* are independent.

$P(A,G) = P(G) * P(A)$

$P(A \mid G) = P(A)$
$P(G \mid A) = P(G)$

$P(A,G) = P(G|A)\, P(A) = P(G)P(A)$
$P(A,G) = P(A|G)\, P(G) = P(A)P(G)$

Slides from Dr. Oates, UMBC

Examples of Worked BBNs

- http://tiny.cc/bn-ex
  - *https://www.ics.uci.edu/~rickl/courses/cs-171/[etc]/cs-171-17-BayesianNetworks.pdf*
- http://tiny.cc/bn-ex2
  - *http://chem-eng.utoronto.ca/~datamining/Presentations/Bayesian_Belief_Network.pdf*
- *https://cw.fel.cvut.cz/wiki/_media/courses/ae4m33rzn/bn_solved.pdf*

Topics: Reasoning Under Uncertainty

- How is the world represented over time?
  - Concepts: timesteps, world, observations
  - Transition model captures how the world changes
  - Sensor model capture what we see, given some world
  - Markov assumption (first-order) makes it all tractable
- What can we do with it?
  - Concepts: Filtering, predicting, smoothing, explaining

Topics: Reasoning Under Uncertainty

- How would you represent this problem as a network and set of conditional probability tables?
  - The weather has a 30% chance of changing and a 70% chance of staying the same.
  - If it's raining, the probability of seeing someone carrying an umbrella is 90%; if it's not raining, it's 20%.
- I saw umbrellas Monday and Tuesday, but not today. What is the most likely weather pattern for those days?

Example

R_{t-1}	$P(R_t \mid R_{t-1})$
t	0.7
f	0.3

Weather has a 30% chance of changing and a 70% chance of staying the same.

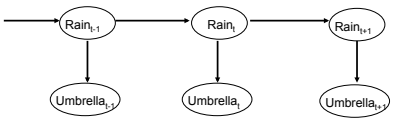

R_t	$P(U_t \mid R_t)$
t	0.9
f	0.2

Fully worked out HMM for rain: www2.isye.gatech.edu/~yxie77/isye6416_17/Lecture6.pdf

Topics: Utility

- How should rational agents make decisions?
- Concepts: rationality, utility functions, value functions, expected value, satisficing, preferences
- Utility is a function of world states
- Must have some preferences that pertain to perceived needs or wants

Topics: Decision Theory

- What is the expected utility of an action?
  - Broadly: its probability times its value
  - The sum of that for all possible outcomes
- Maximum Expected Utility principle

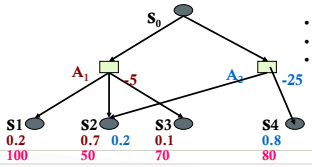

- $U(A_1, S_0) = 62 - 5 = 57$
- $U(A_2, S_0) = 74 - 25 = 49$
- $U(S_0) = \max_a \{U(a, S_0)\} = 57$

Topics: Pareto Optimality

- An outcome is **Pareto optimal** if there is **no other outcome** that **all** players would prefer.
- *S* is a Pareto-optimal solution iff
  - $\forall s' \, (\exists x \, U_x(s') > U_x(s) \rightarrow \exists y \, U_y(s') < U_y(s))$
  - I.e., if X is better off in *s'*, then some Y must be worse off

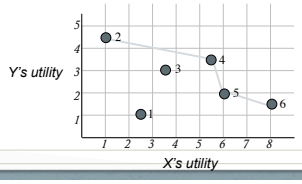

Example questions:

Which solutions are Pareto-optimal?

Which solution(s) maximize global utility (social welfare)?

Topics: Nash Equilibrium

- Occurs when each player's strategy is optimal, **given** strategies of the other players
- No player benefits by **unilaterally** changing strategy while others stay fixed
- Example questions:
  - What strategy should you choose? Why?
  - What strateg(ies) are in a Nash equilibrium?

B \ *C*	Confesses	Denies
Confesses	(3, 3)	(0, 5)
Denies	(5, 0)	(1, 1)

Various Reminders

- **Everything in the readings** is fair game.
- Look at homeworks, sample problems in lectures.
- Look at lectures' "Why?" questions.
- Slides are a good source of **conceptual** understanding.
- Book goes into **detail** and explains more deeply.