1 Search

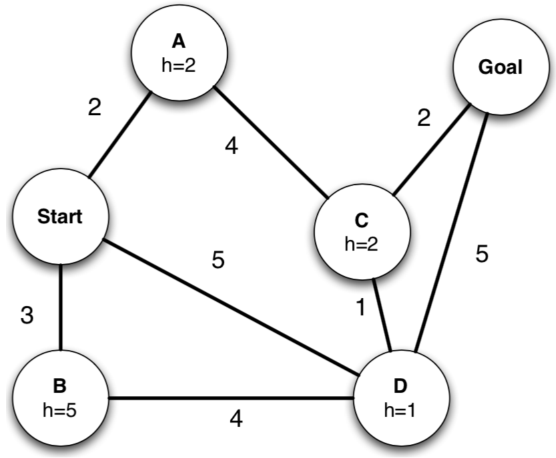

For each of the following graph search strategies, work out the order in which states are expanded, as well as the path returned by graph search. In all cases, break ties in alphabetical order. The start and goal state use letter S and G, respectively. Remember that in graph search, a state is expanded only once.

(a) Depth-first search.

(b) Breadth-first search.

(c) Uniform cost search.

(d) Greedy search with the heuristic values h shown on the graph.

(e) $A^\star$ search with the same heuristic.

2 Adversarial Search

Warm up

1. What is the advantage of adding alpha-beta pruning to a minimax algorithm?

2. Give two advantages of Iterative Deepening minimax algorithms over Depth Limited minimax algorithms.

The following three questions are about the following adversarial "chance" tree.

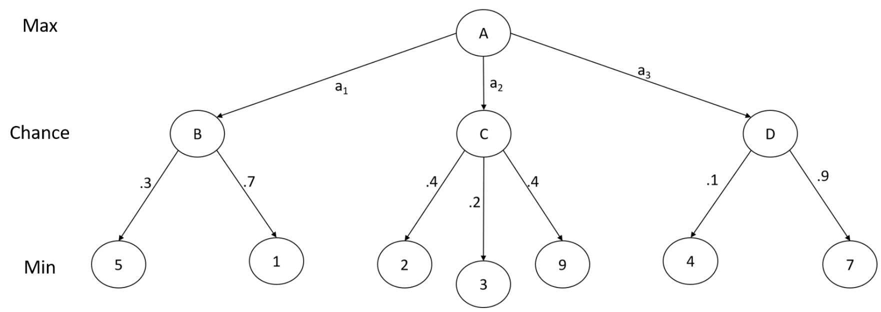

Expectiminimax

1. Calculate the EXPECTIMINIMAX values for nodes B, C and D in the above adversarial "chance" tree.

2. Which action will MAX choose, a_1, a_2, or a_3? Why?

3. If the utility values given for MIN were multiplied with a positive constant c, which action would MAX then choose?

3 CSP Backtracking Search

In this problem, you are given a 3×3 grid with some numbers filled in. The squares can only be filled with the numbers $\{2, 3, \ldots, 10\}$, with each number being used once and only once. The grid must be filled such that adjacent squares (horizontally and vertically adjacent, but not diagonally) are relatively prime.

x_1	x_2	x_3
x_4	x_5	3
4	x_6	2

We will use backtracking search to solve the CSP with the following heuristics:

- Use the Minimal Remaining Values (MRV) heuristic when choosing which variable to assign next.
- Break ties with the Most Constraining Variable (MCV) heuristic.
- If there are still ties, break ties between variables x_i, x_j with $i < j$ by choosing x_i.
- Once a variable is chosen, assign the minimal value from the set of feasible values.
- For any variable x_i, a value v is infeasible if and only if: (i) v already appears elsewhere in the grid, or (ii) a variable in a neighboring square to x_i has been assigned a value u where $\gcd(v, u) > 1$, which is to say, they are not relatively prime.

Fill out the table below with the appropriate values.

- Give initial feasible values in set form; x_1 has already been filled out for you.
- Assignment order refers to the order in which the final value assignments are given. If x_i is the j^{th} variable on the path to the goal state, then the assignment order for x_i is j.
- In the branching column, write "yes" if the algorithm branches (considers more than one value) at that node in the search tree (for example, x_4 considers more than 1 value), and write "B" if the algorithm backtracks at that node, meaning it is the highest node in its subtree that fails for a value, and has to be chosen again. Also write the values it tried then failed.

Variable	Initial Feasible Values	Assignment Order	Final Value	Branch or Backtrack?
x_1	$\{5, 6, 7, 8, 9, 10\}$	______	______	______
x_2	______	______	______	______
x_3	______	______	______	______
x_4	______	______	______	______
x_5	______	______	______	______
x_6	______	______	______	______

4 Local Search

(a) Which of the following local search algorithm are complete and/or optimal? If necessary, specify the conditions that must be true for completeness or optimality.

- First-choice Hill Climbing
  (i) Complete?
  (ii) Optimal?
- Random-restart Hill Climbing
  (i) Complete?
  (ii) Optimal?
- Simulated Annealing
  (i) Complete?
  (ii) Optimal?
- Genetic Algorithm
  (i) Complete?
  (ii) Optimal?
- Local Beam Search
  (i) Complete?
  (ii) Optimal?

(b) Of the local search algorithms above, which one(s) would perform best in a continuous state space and why?

(c) What are the disadvantages and advantages of allowing sideways moves? How can we modify our search algorithm to address the disadvantages?

5 Propositional Logic

1. Warm Up: Are you familiar with these terms?
   - Symbols
   - Operators
   - Sentences
   - Equivalence
   - Literals
   - Knowledge Base
   - Entailment
   - Query
   - Satisfiable
   - Valid
   - Clause - Definite, Horn clauses
   - Model Checking
   - Theorem Proving
   - Modus Ponens

2. Indicate whether the following sentence is *valid*, *satisfiable*, or *unsatisfiable*. If satisfiable, give a model such that the sentence is satisfied. Prove your answer by reducing the sentence to its simplest form. Remember to **show all the steps and write down an explanation of each step.** Let T stand for the atomic sentence *True* and F for the atomic sentence *False*.

   $((T \Leftrightarrow \neg(x \vee \neg x)) \vee z) \wedge \neg(z \wedge ((z \wedge \neg z) \Rightarrow x))$

3. Indicate whether the following sentence is *valid*, *satisfiable*, or *unsatisfiable*. If satisfiable, give a model such that the sentence is satisfied. Prove your answer by reducing the sentence to its simplest form. Remember to **show all the steps and write down an explanation of each step.** Let T stand for the atomic sentence *True* and F for the atomic sentence *False*.

   $(\neg(x \vee \neg x) \wedge y) \vee ((x \vee (z \Rightarrow \neg z)) \wedge ((x \Rightarrow z) \vee \neg(F \Rightarrow T)))$

6 First Order Logic

(a) For each of the logical expressions, state whether it correctly expresses the English sentence and explain.

i. All the Kardashians love Kim: $\forall x$ Kardashian(x) $\wedge$ Love(x, Kim)

ii. Some Kardashian dislikes sugar: $\exists x$ Kardashian(x) $\Rightarrow$ Dislikes(x, sugar).

(b) Forward Chaining: Using the following statements and generalized modus ponens, derive: RidesPrivateJet(North) (In English, this means derive that North West rides a private jet).

1. Kardashian(Kim)
2. Rich(Kim)
3. Parent(North, Kim)
4. $\forall x$ Kardashian(x) $\Rightarrow$ Celebrity(x)
5. $\forall x, y$ Kardashian(x) $\wedge$ Parent(y,x) $\Rightarrow$ Kardashian(y)
6. $\forall x, y$ Rich(x) $\wedge$ Parent(y,x) $\Rightarrow$ Rich(y)
7. $\forall x$ Rich(x) $\wedge$ Celebrity(x) $\Rightarrow$ RidesPrivateJet(x)

(c) For each pair of atomic sentences, give the most general unifier if it exists:

i. ***Q(f(A), f(B)), Q(f(x),f(x))***

ii. ***R(w, w, z, f(z)), R(f(x), f(m(5)), x, y)***

7 Satisfiability and Planning

In the recent pandemic, suppose we are tasked with making a plan to deliver N-95 masks around the U.S. We first try taking a SATplan (logical planning) approach, and formulate the following propositions:

- $at(loc, t)$: our cargo plane is at location loc at time t
- $fuel(x, t)$: the fuel level is at x at time t, $x \in [0, 5]$.
- $hasFuel(loc, t)$: location loc has fuel (to re-fuel the plane with) at time t
- $hasMasks(loc, t)$: location loc has masks at time t

Our starting state is $at(Pittsburgh, 0) \wedge fuel(5, 0)$.

1. Using the above predicates, formulate successor-state axioms for the actions $refuel(t)$, $deliver(t)$, and $fly(origin, destination, t)$. We can only refuel at a location that has fuel, and the fuel level jumps to 5 as a result of refueling. We can fly between any distinct locations, as long as the fuel level is less than 3; after flying, fuel level decreases by 1. We can deliver masks anywhere, and the result is that the location of the plane now has masks.

2. Convert your *deliver* axiom into conjunctive normal form. You may want to abbreviate each proposition. What's the purpose of converting logical sentences into CNF (besides solving recitation problems)?

3. Suppose our goal is to deliver presents to NYC. Describe an algorithm which uses a SAT solver to find a plan for this goal.

4. Run DPLL to determine whether the goal $hasMasks(Pittsburgh, 1)$ is feasible with our knowledge base $at(Pittsburgh, 0) \wedge fuel(0, 0) \wedge D$, where D is your *deliver* axiom instantiated with $loc = Pittsburgh$, $t = 0$ (we can leave the other axioms out because they won't be relevant). What's a possible model found by DPLL?

5. Suppose DPLL returned $False$ on some sentence $A \wedge B$. What entailment conclusions can we draw involving A and B?

6. Now we take a GraphPlan approach. Define each action as an operator in the following table (note that we can drop the t parameter from each predicate and action):

	$refuel$	$fly(o,d)$	$deliver$
Precondition			
Add			
Delete			

7. Now draw the GraphPlan graph up to proposition level S_1. Suppose NYC is the only other location besides Pittsburgh.

8. Which operators are mutually exclusive in A_0? Which propositions are mutually exclusive in S_1?

9. In general, when does GraphPlan stop extending the graph?

10. Is GraphPlan sound? complete? optimal? What about the SATPlan algorithm you described above?

8 Probability

1. *Independence*

   For each of the following four subparts, you are given three joint probability distribution tables. For each distribution, please identify if the given independence / conditional independence assumption is true or false.
   For your convenience, we have also provided some marginal and conditional probability distribution tables that could assist you in solving this problem.

   (a) X is independent from Y.

X	Y	$P(X,Y)$
0	0	0.240
1	0	0.160
0	1	0.360
1	1	0.240

X	$P(X)$
0	0.600
1	0.400

Y	$P(Y)$
0	0.400
1	0.600

   (b) X is independent from Y.

X	Y	$P(X,Y)$
0	0	0.540
1	0	0.360
0	1	0.060
1	1	0.040

X	$P(X)$
0	0.600
1	0.400

X	Y	$P(X \mid Y)$
0	0	0.600
1	0	0.400
0	1	0.600
1	1	0.400

(c) X is independent from Y given Z.

X	Y	Z	$P(X,Y,Z)$
0	0	0	0.280
1	0	0	0.070
0	1	0	0.210
1	1	0	0.140
0	0	1	0.060
1	0	1	0.060
0	1	1	0.030
1	1	1	0.150

X	Z	$P(X,Z)$
0	0	0.700
1	0	0.300
0	1	0.300
1	1	0.700

Y	Z	$P(Y,Z)$
0	0	0.500
1	0	0.500
0	1	0.400
1	1	0.600

X	Y	Z	$P(X,Y \mid Z)$
0	0	0	0.400
1	0	0	0.100
0	1	0	0.300
1	1	0	0.200
0	0	1	0.200
1	0	1	0.200
0	1	1	0.100
1	1	1	0.500

(d) X is independent from Y given Z.

X	Y	Z	$P(X,Y,Z)$
0	0	0	0.140
1	0	0	0.140
0	1	0	0.060
1	1	0	0.060
0	0	1	0.048
1	0	1	0.192
0	1	1	0.072
1	1	1	0.288

X	Z	$P(X,Z)$
0	0	0.500
1	0	0.500
0	1	0.200
1	1	0.800

Y	Z	$P(Y,Z)$
0	0	0.700
1	0	0.300
0	1	0.400
1	1	0.600

X	Y	Z	$P(X,Y \mid Z)$
0	0	0	0.350
1	0	0	0.350
0	1	0	0.150
1	1	0	0.150
0	0	1	0.080
1	0	1	0.320
0	1	1	0.120
1	1	1	0.480

2. *Chain Rule*

(a) When is $P(A, B \mid C)$ equivalent to the following?

(i) $\frac{P(C|A)P(A|B)P(B)}{P(C)}$

(ii) $\frac{P(B,C|A)P(A)}{P(B,C)}$

(iii) $P(A \mid B, C)P(B \mid C)$

(iv) $\frac{P(A|C)P(B,C)}{P(C)}$

(b) When is $P(A \mid B, C)$ equivalent to the following?

(i) $\frac{P(C|A)P(A|B)P(B)}{P(C)}$

(ii) $\frac{P(B,C|A)P(A)}{P(B,C)}$

(iii) $\frac{P(A|C)P(C|B)P(B)}{P(B,C)}$

(iv) $\frac{P(C|A,B)P(B|A)P(A)}{P(B|C)P(C)}$

3. *Probability Tables*

   Let A be a random variable representing the choice of protein in the sandwich with three possible values, $\{mutton, bacon, egg\}$, let B be a random variable representing the choice of bread with two possible values, $\{toast, naan\}$, and let K be a random variable representing the presence of ketchup or not, $\{+k, k\}$.

   How many values are in each of the probability tables and what do the entries sum to?
   Write **'?'** if there is not enough information given.

Table	**num**	**sum**
$P(A, B)$		
$P(A, B, +k)$		
$P(A, B \mid +k)$		
$P(B \mid +k, A)$		

9 Bayes' Nets: Representation, Independence

For this problem, any answers that require division can be left written as a fraction.

PacLabs has just created a new type of mini power pellet that is small enough for Pacman to carry around with him when he's running around mazes. Unfortunately, these mini-pellets don't guarantee that Pacman will win all his fights with ghosts, and they look just like the regular dots Pacman carried around to snack on.

Pacman (P) just ate a snack , which was either a mini-pellet ($+p$), or a regular dot ($-p$), and is about to get into a fight (W), which he can win ($+w$) or lose ($-w$). Both these variables are unknown, but fortunately, Pacman is a master of probability. He knows that his bag of snacks has 5 mini-pellets and 15 regular dots. He also knows that if he ate a mini-pellet, he has a 70% chance of winning, but if he ate a regular dot, he only has a 20% chance.

(a) What is $P(+w)$, the marginal probability that Pacman will win?

(b) Pacman won! Hooray! What is the conditional probability $P(+p \mid +w)$ that the food he ate was a mini-pellet, given that he won?

Pacman can make better probability estimates if he takes more information into account. First, Pacman's breath, B, can be bad ($+b$) or fresh ($-b$). Second, there are two types of ghost (M): mean ($+m$) and nice ($-m$). Pacman has encoded his knowledge about the situation in the following Bayes' Net:

P(M)	
+m	0.50
-m	0.50

P(P)	
+p	0.25
-p	0.75

P(W|M, P)			
+m	+p	+w	0.60
+m	+p	-w	0.40
+m	-p	+w	0.10
+m	-p	-w	0.90
-m	+p	+w	0.80
-m	+p	-w	0.20
-m	-p	+w	0.30
-m	-p	-w	0.70

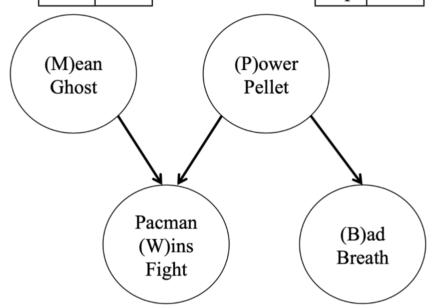

P(B|P)		
+p	+b	0.80
+p	-b	0.20
-p	+b	0.40
-p	-b	0.60

(c) What is the probability of the event $(-m, +p, +w, -b)$, where Pacman eats a mini-pellet and has fresh breath before winning a fight against a nice ghost?

For the remaining of this question, use the half of the joint probability table that has been computed for you below:

$P(M, P, W, B)$				
$+m$	$+p$	$+w$	$+b$	0.0800
$+m$	$+p$	$+w$	$-b$	0.0150
$+m$	$+p$	$-w$	$+b$	0.0400
$+m$	$+p$	$-w$	$-b$	0.0100
$+m$	$-p$	$+w$	$+b$	0.0150
$+m$	$-p$	$+w$	$-b$	0.0225
$+m$	$-p$	$-w$	$+b$	0.1350
$+m$	$-p$	$-w$	$-b$	0.2025

(d) What is the marginal probability, $P(+m, +b)$ that Pacman encounters a mean ghost and has bad breath?

(e) Pacman observes that he has bad breath and that the ghost he's facing is mean. What is the conditional probability, $P(+w \mid +m, +b)$, that he will win the fight, given his observations?

(f) Pacman's utility is +10 for winning a fight, -5 for losing a fight, and -1 for running away from a fight. Pacman wants to maximize his expected utility. Given that he has bad breath and is facing a mean ghost, should he stay and fight, or run away? Justify your answer.

10 Bayes' Nets: Sampling

Consider the following Bayes Net and corresponding probability tables.

R	P(R)
+r	0.4
-r	0.6

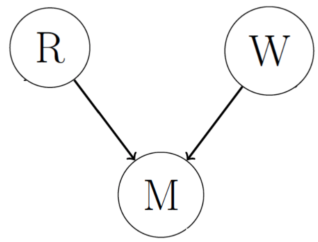

W	P(W)
+w	0.2
-w	0.8

M	R	W	P(M | R,W)
+m	+r	+w	0.1
-m	+r	+w	0.9
+m	+r	-w	0.45
-m	+r	-w	0.55
+m	-r	+w	0.35
-m	-r	+w	0.65
+m	-r	-w	0.9
-m	-r	-w	0.1

Consider the case where we are sampling to approximate the query $P(R \mid +m)$.

Fill in the following table with the probabilities of drawing each respective sample given that we are using each of the following sampling techniques. Let $P(+m) = a$.

Method	$< +r, -w, +m >$	$< +r, +w, -m >$
Prior sampling		
Rejection sampling		
Likelihood weighting		

We are going to use Gibbs sampling to estimate the probability of getting the sample $< +r, -w, +m >$. We will start from the sample $< -r, +w, +m >$ and resample W first then R. What is the of drawing sample $< +r, -w, +m >$?

11 HMMs and Particle Filtering

Consider the following hidden Markov model with a binary hidden state X. The transition probabilities and initial distribution are:

X_0	$P(X_0)$
0	0.5
1	0.5

X_t	X_{t+1}	$P(X_{t+1}\|X_t)$
0	0	0.9
0	1	0.1
1	0	0.5
1	1	0.5

(a) After one timestep (i.e., after a dynamics update), what is the new belief distribution $P(X_1)$?

X_1	$P(X_1)$
0	
1	

Now, we incorporate sensor readings as our observations. The sensor model is parameterized by some value $\beta \in [0, 1]$:

X_t	E_t	$P(E_t\|X_t)$
0	0	β
0	1	$1 - \beta$
1	0	$1 - \beta$
1	1	β

(b) At $t = 1$, we get the first sensor reading, $E_1 = 0$. Find $P(X_1 = 0|E_1 = 0)$ in terms of β.

(c) For what range of values of β will a sensor reading $E_1 = 0$ increase our belief that $X_1 = 0$? In other words, what is the range of β for which $P(X_1 = 0|E_1 = 0) > P(X_1 = 0)$?

(d) Now, we want to use particle filtering to predict what state value our model currently assumes. At time t, there are 2 particles in state value 0, and 3 particles in state value 1. What is the prior belief distribution $\hat{P}(X_t)$?

X_t	$\hat{P}(X_t)$
0	
1	

(e) At some time t, we receive our first sensor reading $E_t = 1$. Given $\beta = 0.6$ and the previous table for $P(E_t|X_t)$, how many particles will be in each state value after updating our belief and resampling? When resampling, use this list of numbers as a source of randomness: [0.182, 0.703, 0.471, 0.859, 0.382] and fix the order of states to be $X_t = 0$, $X_t = 1$.

(f) Suppose we now have the following modified HMM structure, in which the hidden variables now have a parent variable Y_t, starting at $t = 1$:

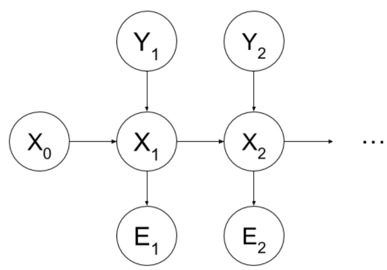

Write expressions for answering the following queries. Make sure your expressions are solely in terms of the probability tables from the HMM, and that they are in the simplest possible form (hint: conditional independence!). You must explicitly write out any normalization constants.

(i) $P(X_1 \mid E_1)$

(ii) $P(Y_1 \mid X_1, X_0)$

(iii) $P(Y_1 \mid E_1)$

(iv) $P(Y_2 \mid E_1, E_2)$

12 LP

The Easter bunny is running late and is struggling to pack all the candy that he wants to hand out. This year, he is giving out massive amounts of chocolate and Peeps.

Let x_1 represent pounds of chocolate and let x_2 represent pounds of Peeps. We assume we can deliver a fraction of a pound of chocolate or Peeps. Unfortunately, the Easter bunny's basket can only fit 10.5 pounds of sweets. Furthermore, the Easter bunny wants to provide enough candy for 10 kids. Each pound of chocolate is enough for 2 kids while a pound of Peeps is only enough for 1 kid. However, the Easter bunny also wants to maximize the children's happiness. Chocolate brings 4 units of happiness while Peeps only bring 1.

1. Represent the following problem as an LP and graph the constraints in the provided graph.

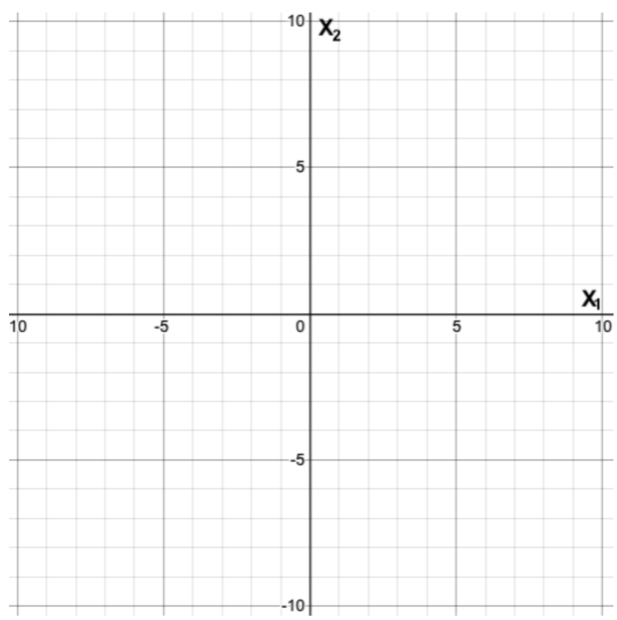

2. What would the optimal solution be?

3. If the Easter bunny didn't know how much happiness chocolate and peppermints bring, what would be a cost vector that makes the optimal solution (0, 10.5)?

4. List three cost vectors that will lead to an infinite number of solutions.

5. If the Easter bunny can only give out chocolate in the form of 1 pound bars that cannot be divided up, which constraints will you have to add in the first iteration of Branch and Bound?

13 MDPs/RL

1. Warm Up

   - What does the Markov Property state?
   - What are the Bellman Equations, and when are they used?
   - What is a policy? What is an optimal policy?
   - How does the discount factor γ affect the agent's policy search? Why is it important?
   - What are the two steps to Policy Iteration?
   - What is the relationship between $V^*(s)$ and $Q(s, a)$?
   - Exploration, exploitation, and the difference between them? Why are they both useful?
   - What is the difference between on-policy and off-policy learning?
   - What is the difference between model-based and model-free learning?
   - We are given a pre-existing table of Q-values (and its corresponding policy), and asked to perform ϵ-greedy Q-learning. Individually, what effect does setting each of the following constants to 0 have on this process?

     (i) α:

     (iii) ϵ:

- For each of the following functions, write which MDP/RL value the function computes, or none if none apply. We are given an MDP (S, A, T, γ, R), where R is only a function of the current state s. We are also given an arbitrary policy π.

  Possible choices: $V^*, Q^*, \pi^*, V^\pi, Q^\pi$.

  (i) $f(s) = R(s) + \sum_{s'} \gamma T(s, \pi(s), s') f(s')$

  (ii) $g(s) = \max_a \sum_{s'} T(s, a, s')[R(s) + \gamma \max_{a'} Q^*(s', a')]$

2. MDPs - Micro-Blackjack: In micro-blackjack, you repeatedly draw a card (with replacement) that is equally likely to be a 2, 3, or 4. You can either Draw or Stop if the total score of the cards you have drawn is less than 6. Otherwise, you must Stop. When you Stop, your utility is equal to your total score (up to 5), or zero if you get a total of 6 or higher. When you Draw, you receive no utility. There is no discount ($\gamma = 1$).

   (a) What are the states and the actions for this MDP?

   (b) What is the transition function and the reward function for this MDP?

   (c) Fill out the value iteration table below. We have filled out the first row for you. (Recall that we always initialize $V_0(s)$ to 0 for all states s.) Then, perform policy extraction and give the optimal policy for this MDP.

V	0	2	3	4	5	Done
V_0	0	0	0	0	0	0
V_1						
V_2						
V_3						
Policy Extraction						

3. Ms.Pacman: While Pacman is busting ghosts, Ms. Pacman goes treasure hunting on GridWorld Island. She has a map showing where the hazards are, and where the treasure is. From any unmarked square, Ms. Pacman can take any of the deterministic actions (N, S, E, W) that doesn't lead off the island. If she lands in a hazard square or a treasure square, her only action is to call for an airlift (X), which takes her to the terminal *Done* state; this results in a reward of -64 if she's escaping a hazard, or +128 if she reached the treasure. There is no living reward.

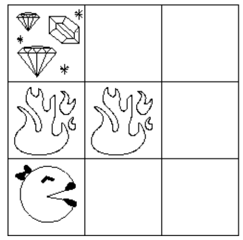

(a) Let $\gamma = 0.5$. What are the optimal values V^* of each state in the grid above?

(b) How would we compute the Q-values for each state-action pair?

(c) What's the optimal policy? You may use the grid below to fill in the optimal action for each state.

Call this policy π_0.

Ms. Pacman realizes that her map might be out of date, so she uses Q-learning to see what the island is really like. She believes π_0 is close to correct, so she follows an ϵ-random policy, ie., with probability ϵ she picks a legal action uniformly at random (otherwise, she does what π_0 recommends). Call this policy π_ϵ.

π_ϵ is known as a *stochastic* policy, which assigns probabilities to actions rather than recommending a single one. A stochastic policy can be defined with $\pi(s, a)$, the probability of taking action a when the agent is in state s.

(d) Write a modified Bellman update equation for policy evaluation when using a stochastic policy $\pi(s, a)$.

14 Game Theory

1. Warm Up

   (a) To formulate a game, what needs to be defined?

   (b) How can we define a strategy?

   (c) What is the type of a solution concept of a game?

   (d) What is a Nash Equilibrium?

   (e) Does a Nash Equilibrium always exist? Does a pure Nash always exist?

   (f) Give an example of a game with infinite actions such that no Nash Equilibrium exists.

2. Equilibria

   With the Grinch now reformed, he has started helping Santa deliver presents. They both can either take a northern path or a southern path to deliver presents. Santa prefers to deliver on the northern path and the Grinch prefers the southern path. However, both are happier when they deliver together. If they deliver together, the one who prefers that location gets a payoff of 9, and the other gets a payoff of 5. If they both deliver to their preferred place, Santa gets a payoff of 3 and the Grinch gets a payoff of 4. If they both deliver to their unpreferred place, they both get a payoff of 1.

   (a) Finish formulating the game by filling in the values and actions for the two players.

		The Grinch	
SANTA			

   (b) Identify the pure strategy Nash Equilibria in this game.

(c) Determine the mixed strategy Nash Equilibrium in this game.

(d) With the mixed strategy Nash Equilibrium, what is the probability each action outcome actually gets played.

(e) Use the probabilities found in the last part to find the expected utility for each Santa and the Grinch in the mixed Nash. Who gets a higher utility?

(f) Now let's consider the important fact that Santa is still the boss. And so the game here does not reflect a one-shot game, as before, but a Stackelberg game with commitment and a leader, Santa. With Santa as a leader, find the Stackelberg Equilibrium. What are the utilities for each player with this dynamic, and how does it compare to the mixed nash utility for each player?

(g) The Grinch is a little annoyed about the fact that Santa is seeming dictatorial in the previous part. To combat this, he tells Santa that he will always choose south no matter what Santa does, so Santa should choose south. Still, Santa has the first official choice. Why does the Grinch's statement not matter?