

Midterm Exam Solutions

16-311: Introduction to Robotics

2019

Name: _____

Andrew ID: _____

Team Number: _____

- You will have 1 hour and 15 minutes to complete this exam
- There are 7 sections on 20 pages. Make sure you have all of the pages. Write your Andrew ID on all the sections and keep your work in that section (they will be graded separately). There are blank pages throughout the sections.
- When making drawings - be precise. Rounded edges should look rounded, sharp edges should look sharp, sizes should be close to scale. Neatness counts.
- Show your work. Partial credit may apply. Likewise, justify algebraically your work to ensure full credit, where applicable.
- It should be very clear what your final answer is, circle it if necessary.
- You may need to make certain assumptions to answer a problem. State them (e.g. what is optimal).
- You are allowed one handwritten two-sided reference sheet for the exam. No cell phones, laptops, neighbors, etc. allowed.
- Good luck and you can do it.

PAGE INTENTIONALLY LEFT BLANK

1 Vision

Andrew ID: _____

For this example, you have a downward-facing camera capturing a conveyor belt with recyclable items. Every second, your conveyor belt stops and the camera takes an image. You have saved a variable *image*, which is a $m \times n \times 3$ array with values from 0 to 255 taken from an image taken at one instance.

The following code has been written for you. First figure out what it does and then determine what needs to be changed.

```
1 % Create a copy of image to store output
2 outputImage = image;
3
4 % Get size info about the image
5 [imLength, imWidth, imDepth] = size(image);
6
7 % ???
8 for rows = 1:imLength
9     for cols = 1:imWidth
10         if image(rows, cols, 3) > 200
11             outputImage(rows, cols, 1:imDepth) = 0;
12         end
13     end
14 end
15
16 % Plot the original and output images next to each other
17 figure()
18 subplot(1,2,1), imshow(image)
19 subplot(1,2,2), imshow(outputImage)
```

1. What line(s) need to be added to this code so that it could produce the output below?

Between lines 11 and 12, we need to set all non-blue pixels to white.

In MATLAB:

else

image(rows, cols, 1:imDepth) = 255;

(4 points) Accept either verbal explanation or pseudocode. No points off for syntax error.




Figure 1: Left: Input image. Right: Output image

2. Two different coworkers wrote their own code to group together pixels of the same object. Both chose to implement the floodfill algorithm, which recursively renames neighboring pixels if they have the same value. In one to two sentences, explain which implementation is correct.

Both coworkers are correct, they are just using a different definition of neighboring pixels. Coworker 1 used the 4-point connectivity and coworker 2 used the 8-point connectivity

(2 points)

Input:




Figure 2: Input image.

Outputs:




Figure 3: Left: Output from Coworker 1. Right: Output from Coworker 2.

3. At certain times of the day, light from a nearby window casts shadows on the conveyor. What is one thing you could do algorithmically so that your segmentation would group a single object separated by a streak of black as one object? Assume that the shadows are small compared to the objects. Explain in one to two sentences.

Before thresholding or before segmentation, you could blur your image (apply a Gaussian filter, etc.). You would form the filter in such a way that it would "spread" some of the color from the object to the black pixels separating an object.

(4 points) Accept alternate methods that are still correct.




Figure 4: Image showing shadows cast across conveyor belt.

4. If your camera is 20 cm away from the conveyor belt, has a resolution of 10 pixels per cm, and a distance from image plane to focal point of 0.5 cm, how wide is an object directly below the camera that is 20 pixels wide in your image? Assume that you can ignore the lens geometry, i.e. the camera is a pinhole camera as we examined in class. Show your work in at least one equation

$$\frac{d}{x} = \frac{d_i}{x_i} = \frac{20\text{cm}}{x} = \frac{0.5\text{cm}}{(20\text{pixels} * 1\text{cm} / 10 \text{ pixels})} \text{ or } \frac{x}{d} = \frac{x_i}{d_i} = \frac{x}{20} = \frac{(20\text{pixels} * 1\text{cm} / 10 \text{ pixels})}{0.5\text{cm}}$$

$$x = 80 \text{ cm}$$

(3 points) 2 points for correct equation. 1 point for correct solution. -0.5 for no units. -0.5 for order of magnitude error. -0.5 for arithmetic mistake.

2 Machine Learning

Andrew ID: _____

1. In the diagram below, what do the blue neurons make up?

The hidden Layer




Figure 5: Neural Net Diagram.

2. Name 3 applications where neural networks are or could be used.

Image classification, reinforcement learning of cartpole balancing, scene understanding, economic prediction, sports predictions, etc.

(3 points) Any reasonable applications with a lot of data.

3. Name 1 scenario or conditions where neural networks would not be the best choice.

A scenario where the relationship between inputs and outputs could be represented through a simple relationship (a neural network would complicate and obscure the connection). An application with limited data (neural nets need a lot of data). A scenario without much processing time or power (neural nets can take a long time to train).

(2 points) Any one of these cases. No explanation needed for full credit.

4. In one to two sentences, what is the difference between supervised and unsupervised learning?

With supervised learning, you have labeled training data. In unsupervised learning, you do not.

(2 points)

5. To classify Waldo images from boxes, did you use supervised or unsupervised learning?

Supervised learning (we made images and labeled them as Waldos or boxes.)
(1 point) No alternate solution, everyone had to do this homework this way.

6. In one to two sentences, what is the difference between reinforcement and imitation learning?

Reinforcement learning trains an agent to map observations to actions for a particular goal (input is state, output is probability distribution over actions).
Imitation learning is mimicking an expert.

(2 points)

7. To train your cartpole, did you use reinforcement learning or imitation learning?

Imitation learning for the PID portion and RL for the policy improvement portion.

(2 points) One point for each part.

3 Control

Andrew ID: _____


In this scenario, you are controlling a toy submarine. This toy has fans that allow it to go down into the water. The submarine will float if these fans are off. You have added a depth sensor to the submarine that can give very accurate and fast readings. We are only worried about controlling the submarine in the vertical direction. Assume that you are operating at depths where all external forces (gravity, pressure, etc.) are constant.

1. What type of controller would be most helpful for this system? An open loop controller? A controller with just feedback? A controller with feedback and feedforward? Explain your answer in one sentence.

Controller with feedback and feedforward. Feedback to get information from the depth sensor, feedforward to account for the buoyancy of the toy (that it would float up if the fans were off).

(2 points) -1 for just feedback.

2. Draw a block diagram of your chosen controller. Include at least the following labels in boxes or on arrows: Plant, Desired State, Controller, Sensor, and Actual State.



3. What are two different ways you could tune a PID controller to reduce the oscillations in the following graph:

Increase D to dampen the oscillations or decrease P to decrease the frequency of the oscillations.

(2 points) 1 point for each solution.





Figure 6: Graph with oscillations.

4 Path Planning


Andrew ID: _____

1. Draw the visibility graph on the following environment. Assume a point robot:



(2 points)

2. Draw the shortest path from start to goal. Assume a point robot:



(1 point)

3. Is this path optimal? Respond in one sentence.

Depends on what you consider optimal. For example, this path is optimal wrt path length but not to maximizing distance from obstacles.

(2 points)

4. Depict the configuration space obstacle (just the obstacle, there are no walls here) for a robot that can translate and rotate. For simplicity, imagine that the robot can only be at orientations in multiples of 90° as shown to the right of the obstacle.





Figure 7: Three-dimensional depiction of configuration space obstacle.




Figure 8: Two-dimensional depiction of configuration space obstacle.

(8 points, 2 for each orientation) -50% if did a configuration space using and/or to combine orientations.

5. You are working on a vacuuming robot that has a map of its environment. Divide the environment below in such a way that you could be sure that you entirely covered it by vacuuming in a simple up and down (boustrophedon) pattern.




Figure 9: Sample environment with typical trapezoidal decomposition.

Problem didn't specify trapezoidal decomp, so give credit for less lines as long as the robot does not have to backtrack.

(2 points)

6. Would this be a good division to abstract into a graph where each node of the graph is one of the sections? Why or why not? Explain in one sentence.

Yes. If the world is divided using trapezoidal decomposition and each cell is a node, you know that adjacent cells are traversible and can be represented by a vertex between the nodes

(2 points)

PAGE INTENTIONALLY LEFT BLANK

5 Graph Search

Andrew ID: _____

This section will evaluate understanding of concepts related to graph search.

1. In the graph below, you want to see if node H is reachable from node A. Would Breath First Search (BFS) or Depth First Search (DFS) tell you that H is reachable from A faster (in a shorter number of node visits)? Break ties alphabetically. So if two nodes are on the same "level" for your algorithm, visit the one earlier in the alphabet first.

BFS





Figure 10: Sample tree for BFS and DFS.

2. Is this always the case? Will the answer to the previous question always allow you to find a node faster? Answer in one sentence.

No. It depends on your tree structure and where the nodes happen to be.
Counter example: E would be faster to find through DFS in this scenario

(2 points) 1 for answer, 1 for explanation.

3. The following nodes are waypoints around the REL. Note that each edge between two nodes has an associated cost with it that takes into account distance, terrain, etc. The numbers inside the circles represent the heuristic guess of the distance to the goal which, for this example, is euclidean distance. You start at node A and there is a key rock formation at node H that you are trying to get to.



Using A*, find the path from start to goal. List the nodes that are EXPANDED in the A* algorithm above in the order in which they are expanded:

ACEFBG

(6 points) Can only receive partial credit if it is apparent that they are implementing A* by adding real costs plus heuristic cost at that node and just made an arithmetic error.

4. List the nodes in the final path that you found using A*.

AGH

(1 point)

5. What is the total real cost (not heuristic) of this path?

15

(1 point)

6 Localization

Andrew ID: _____


The following questions cover concepts of localization covered in lecture and lab.

1. Why is the term "encoder error" misleading? Answer in one sentence.


The encoders are not malfunctioning, the wheels may be slipping or gears backlash, etc. which is causing you to be off path, but the encoders are likely recording the correct number of revolutions.

(2 points)

2. You run a basic differentially driven robot with two motors and no control system 20 feet and record the ending location. You repeat this process 100 times. Draw what the distribution of endpoints might look like.



3. In class we looked at an example of a robot moving in one dimension by a series of flowers that can sense flowers if they are right next to it. If you sense a flower, what might your discrete probability distribution look like? Draw below the image.



(2 points)

4. In Lab 6, did you sample continuously (or as fast as you could get sensor readings) or discretely? Explain how and why you implemented this in two sentences.

Accept either answer as long as they explain how and why

(2 points)

7 Guest Speaker Topics

Andrew ID: _____

Select ONE of the following questions to answer in one sentence. If you answer both, we will just grade the first one.

(2 points total)

1. Earlier this semester, Prof Oliver Kroemer gave a brief introduction to some of his work in machine learning. What does self-supervised learning mean?

The program labels the data on its own.

2. One of Prof Choset's students, Shuo Yang, came in to show us the Biorobotics Lab's six-legged robot. What type of motors does this robot have and what is their benefit?

Series elastic actuators, they give the robot compliance.

PAGE INTENTIONALLY LEFT BLANK
This is the end of the test.