ESE 650 - Learning In Robotics
Project 5
Learning Planning Costs

- Varsha Shankar
vasha@seas.upenn.edu

This report describes project #5 for ESE 650; Learning Planning Costs and is divided into the following sections:
Section I - Introduction
Section II - Data
Section III - Feature Extraction
Section IV - Planning
Section V - Conclusions

Section I - Introduction

This project aims to use an aerial map to learn optimal weights to convert local features in the map into traversal costs and thus match hand labelled routes on the map. The problem of finding allowable and safe paths in a mapped environment is a fundamental issue in robotics and this project serves as a quick introduction to the same.

Section II - Data

Input data for this project was in the form of an aerial view map of the Penn campus.

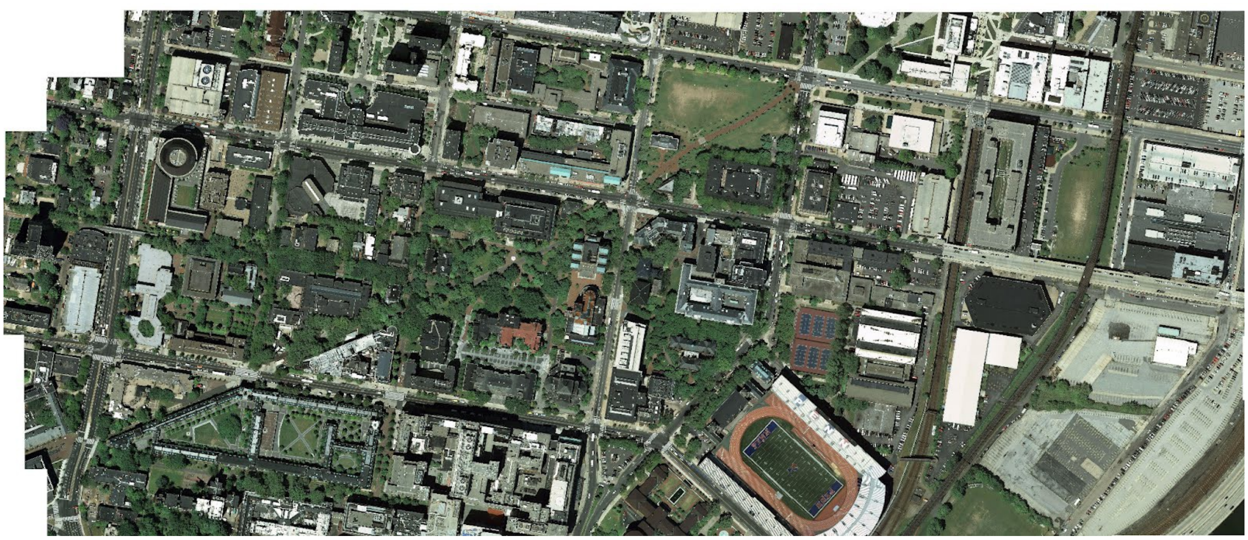

Figure 1 : Aerial view of the Penn campus

This image was downsampled for processing in the remainder of this project. From this map, we derived features for two modes of travel viz. driving and walking. The weights associated with each feature given the mode of travel varied significantly. For example, in the case of the person, green regions of the map are walkable, the roads are not and the sidewalks are permitted. In the case of a car, the roads are allowable, the green regions are not and the buildings are also not permitted to be traversed through.

Section III - Feature Extraction

There are two scenarios for which we plan paths in this project viz. car driving and pedestrian walking. The features used are as follows:

Car driving : Road features, building features, vegetation features

Pedestrian walking : Road features, building features, vegetation features, sidewalk features

The details of extraction on each feature are detailed below.

Road detection

A supervised learning method using a multivariate gaussian model was trained to recognise roads. The hue, saturation and intensity channels of a representative section of the road was selected from the map to train the model on. These three channels were found to be capture the significant features of a road the best. Once mean and the covariance of these three channels is calculated, we find the value of the probability density function at every point in the image. This gives us a probability of each pixel in the map being a part of road.

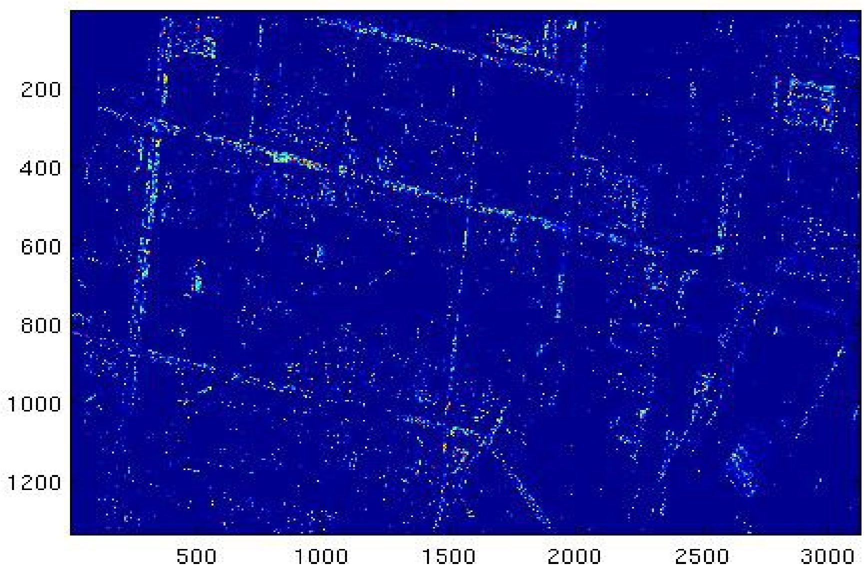

Figure 2 : Output of the road detector

Vegetation detection

The procedure for vegetation detection is very similar to the road detection. The color channels used for vegetation detection were hue, saturation and green. This detector similar to the road detector gives us a probability map of each pixel being part of vegetation.

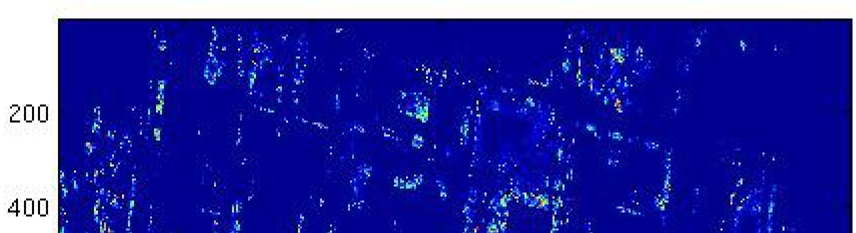

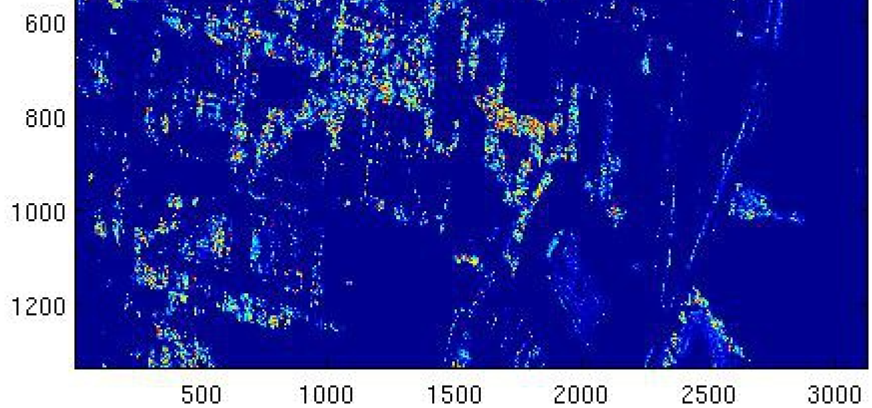

Figure 3 : Output of the vegetation detector

Building detection

HSV colorspace was found to bring out the color of buildings and equivalently roads prominently. The image was thus converted to HSV space. The image was thresholded to convert it to black and white. A connected components analysis was run on the thresholded image. Based on a minimum area and a aspect ratio range criteria, potential buildings are selected and given a value of 1 in the building detector map. All other areas are given a value of 0. This detector serves as a 0/1 classifier for buildings. Its output is shown below. The pixels that are deemed to be part of a building are coloured black.

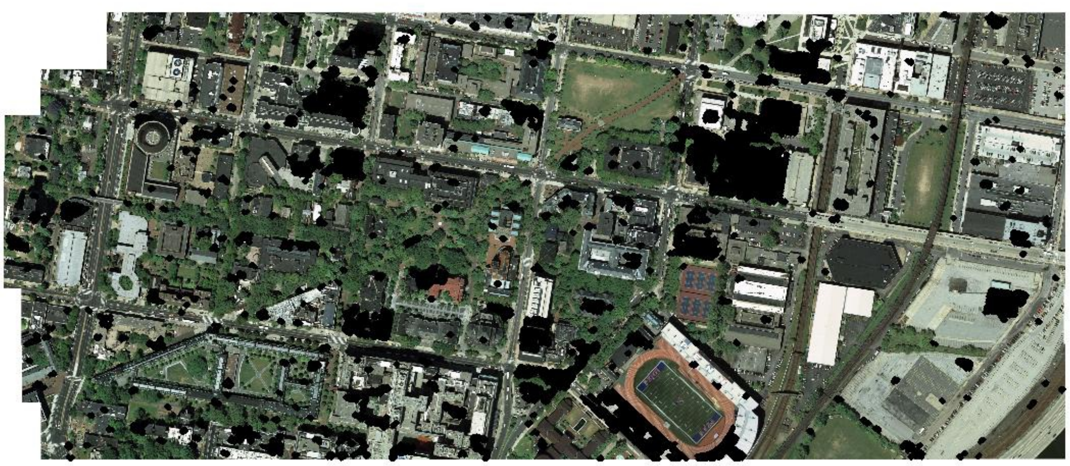

Figure 4 : Output of the building detector

Sidewalk detection

Sidewalk detection followed a similar process to road and vegetation detection. The color channels used for this were hue, saturation, intensity and green.

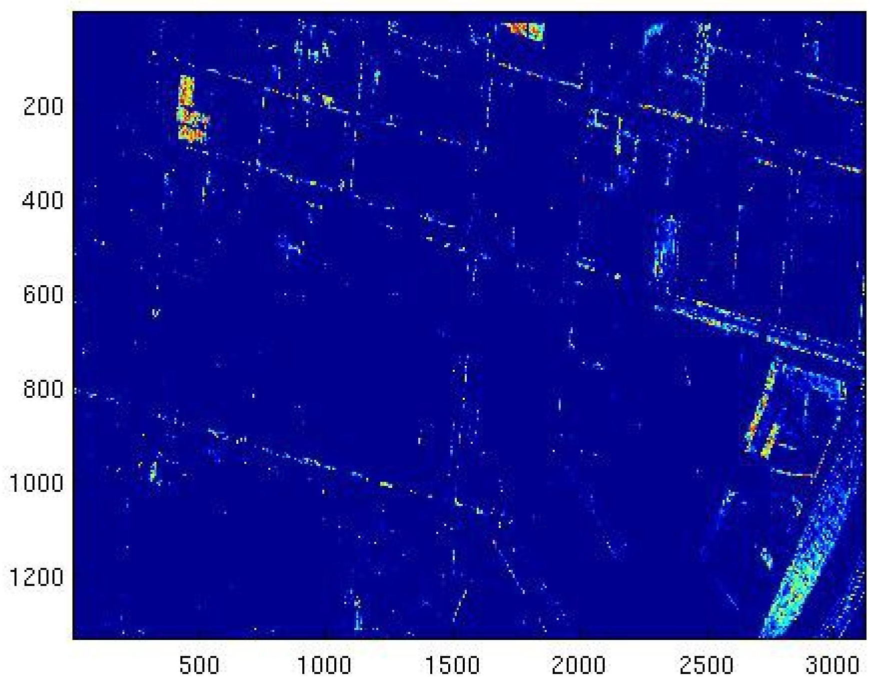

Figure 5 - Output of the sidewalk detector

Section IV - Planning

A cost map is built using the feature detectors described above for each scenario (car and pedestrian modes of travel). For the car, a cost map is built by weighting the road feature positively and the vegetation and building features negatively. The feature maps are summed are passed through a sigmoid function

$$\frac{1}{1+e^{w_1*vegetation+w_2*road+w_3*building}}$$

The output of this sigmoid is used to build the cost map.

Similarly, for the pedestrian, the cost map is built using positive weights for vegetation and sidewalks and negative weights for roads and buildings.

This cost map is used as input to a Dijkstra's planning algorithm to get a planned path.

Section V - Results

Path planning for a car

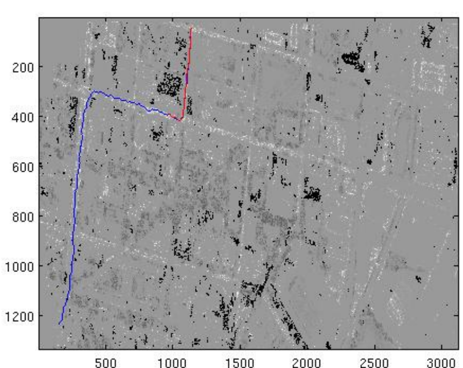

Figure 5 : Example

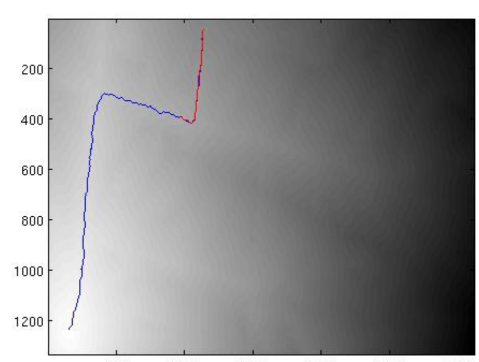

of a car
path

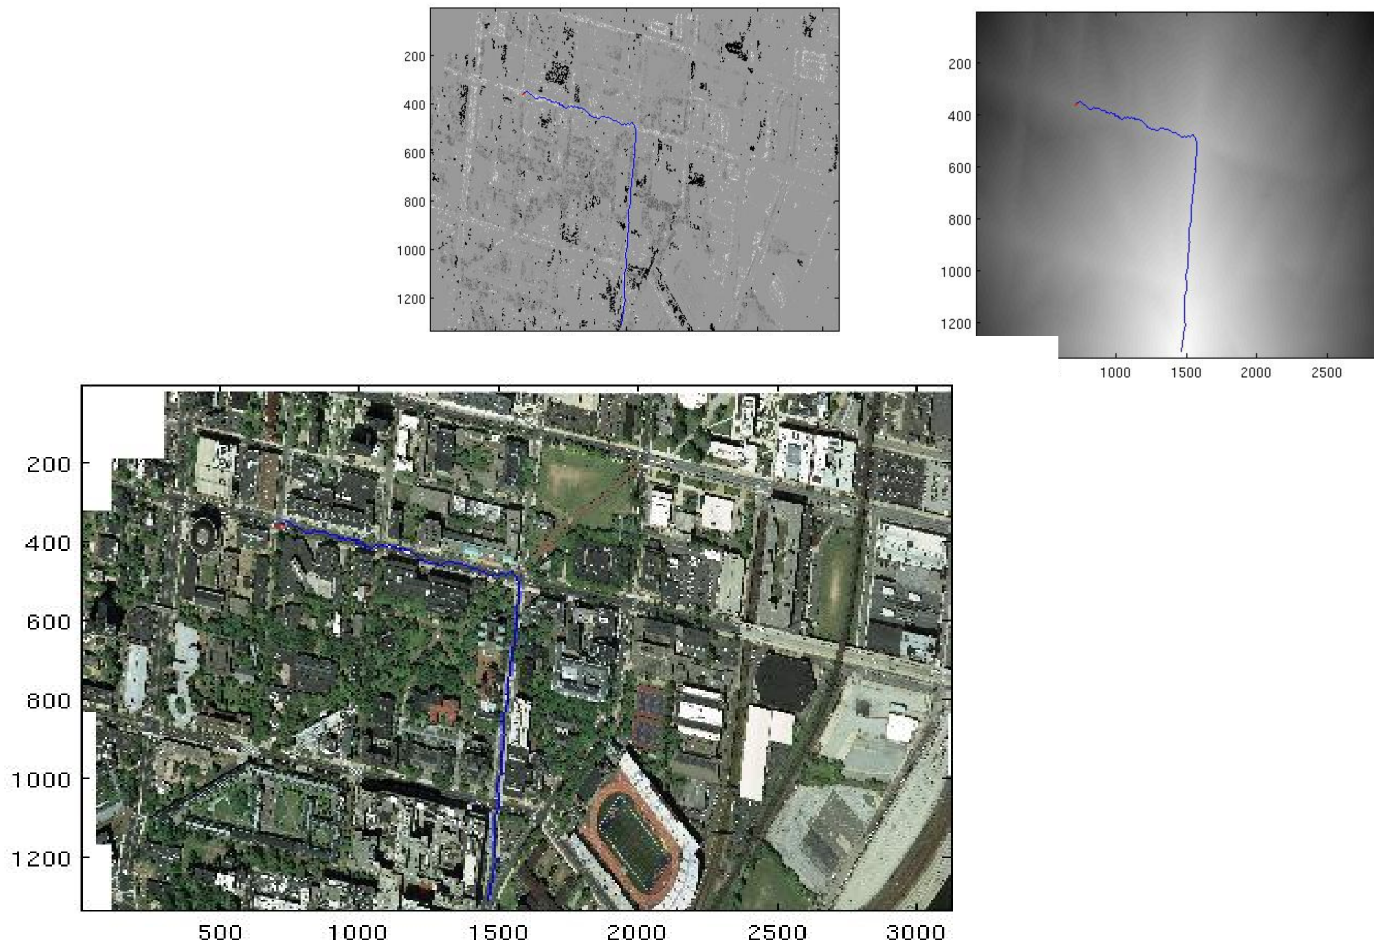

Figure 6 : Example of another car path

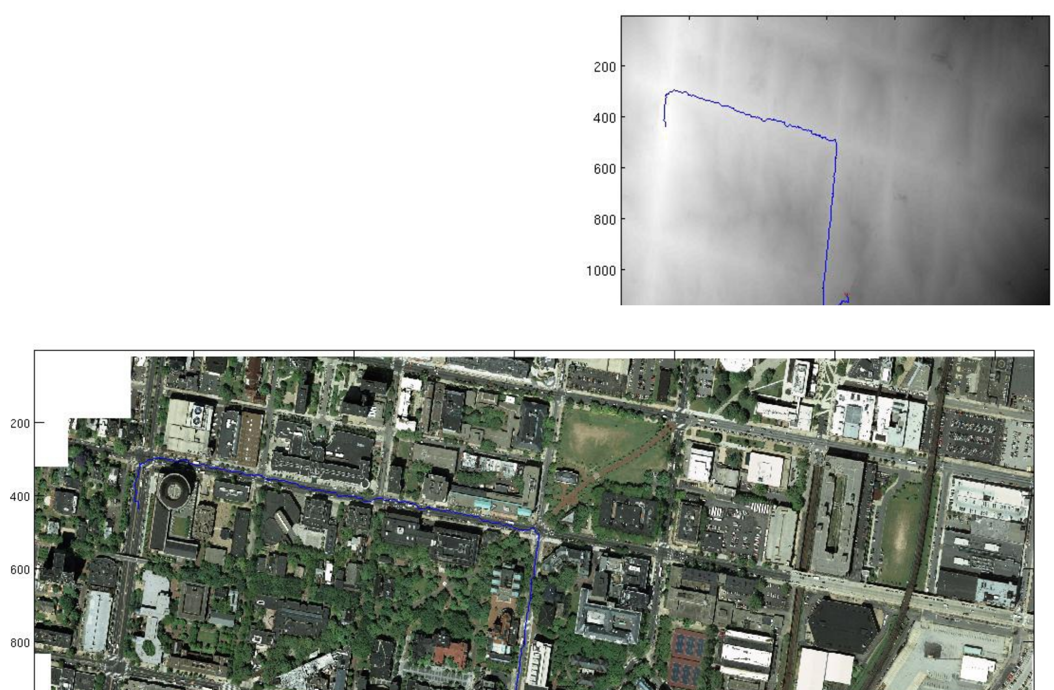

Figure 7 : Yet another example car path

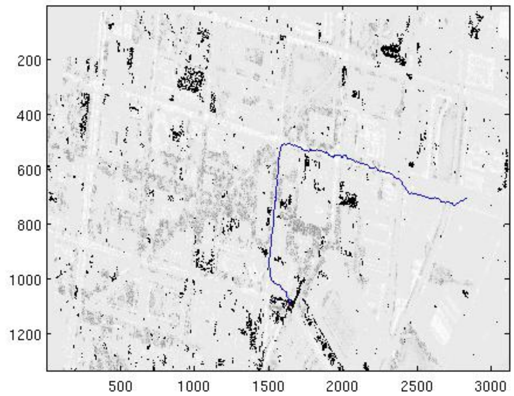

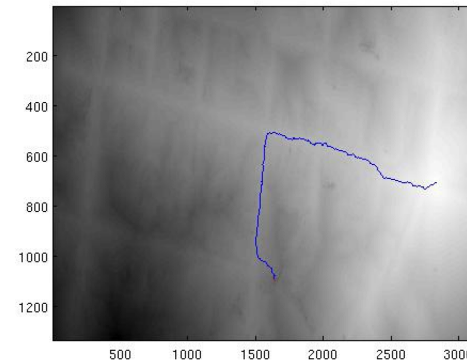

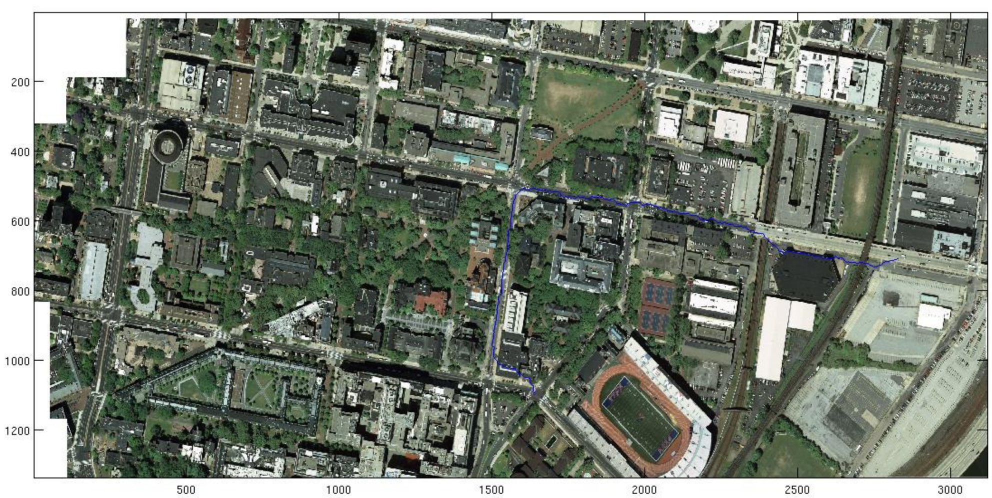

Path planning for a pedestrian

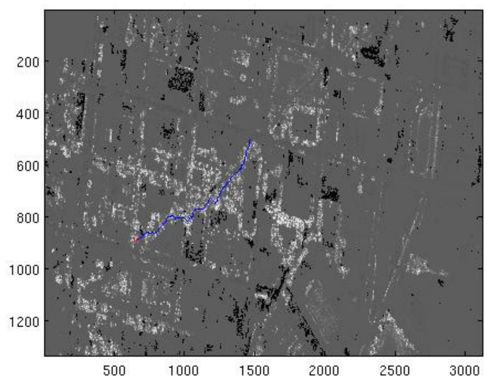

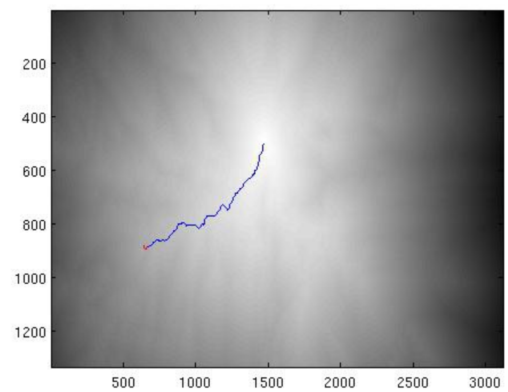

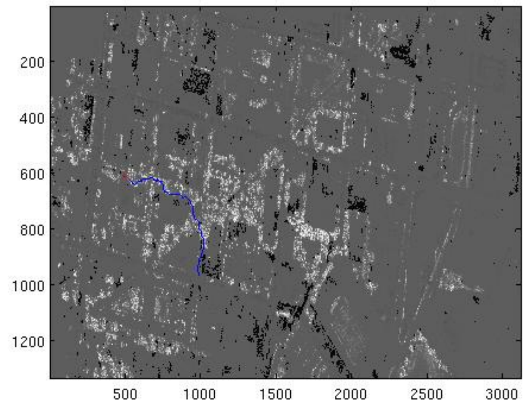

Figure 9 : Example of a pedestrian path

Figure 10 : Example of pedestrian path

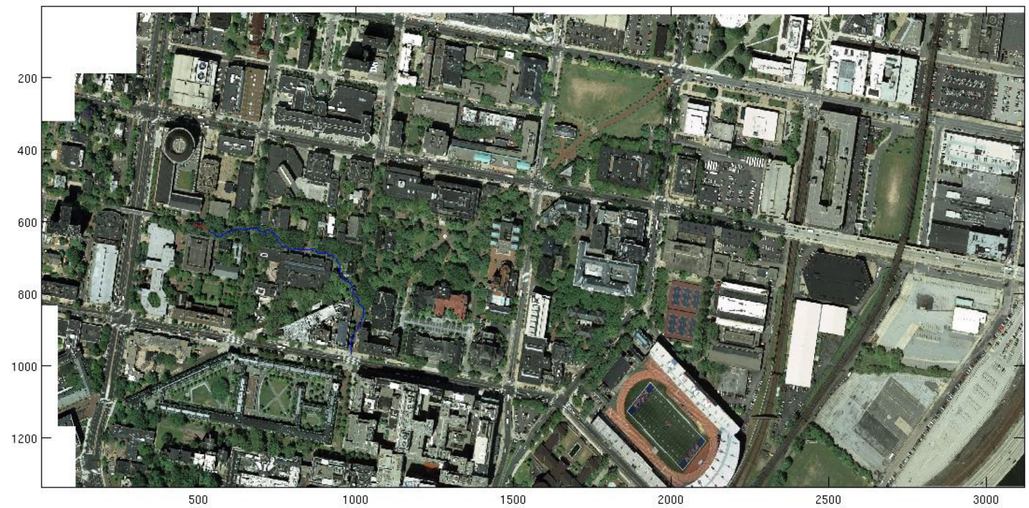

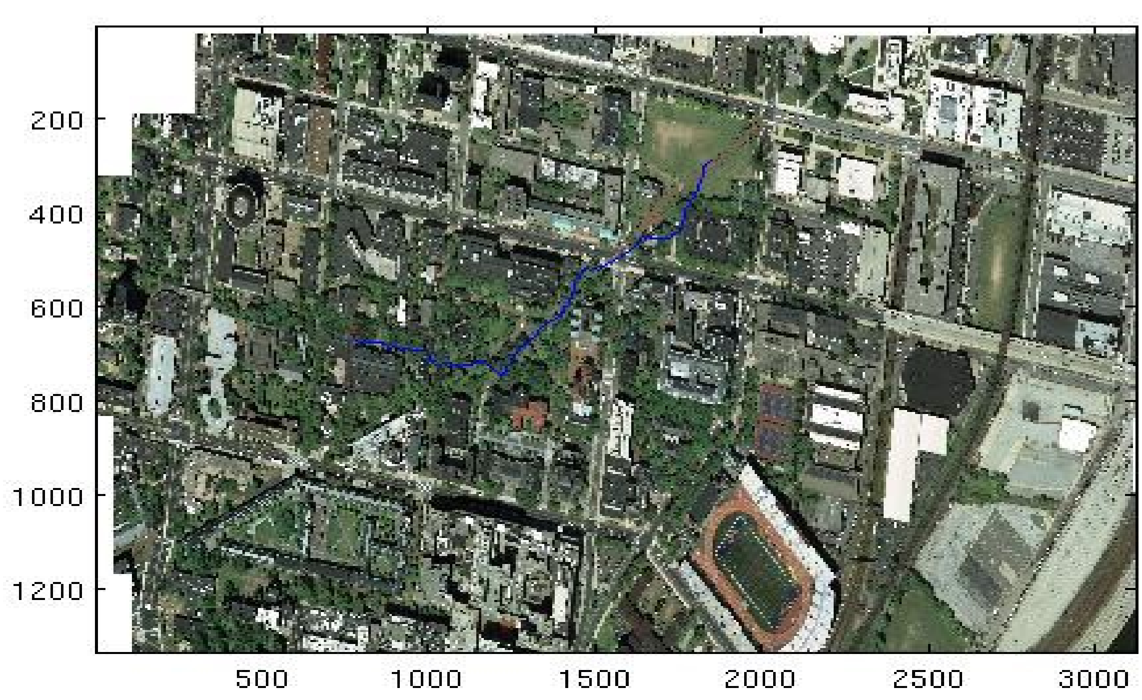

Figure 11: Another example of a pedestrian path

Section VI - Conclusion

For this project, I manually adjusted weights for the feature values in order to get a cost map for car and pedestrian modes of travel. Learning the weights on the features using a regressor and training paths to indicate positive and negative features would help to get more accurate estimates and better paths. Future work on this project would be directed towards learning weights through such a supervised learning technique. Further, improvements might also be achievable in the form of max margin planning as described in RatCliff et all Learning to Search.