

Figure 6. Value function of an agents while mapping a given region. The value function is high around the roads that are close to the agent in a sense that implies that this cluster has been assigned to that agent. The red dot on the left and the blue dot on the right represent the agent's current position.




Figure 7. Value function of an agents while mapping a given region. The value function is high around the roads that are distant and sporadically distributed implying an exploratory strategy. The red dot on the left and the blue dot on the right represent the agent's current position.




Figure 8. Bird's eye view of a partially complete traversal of MARVIN (left) and of the GVIN (right). While slightly more spread out, the traversal of the GVIN leaves many small streets unvisited and is less thorough overall.

MARVIN




Figure 9. Bird's eye view of a partially complete traversal of MARVIN trained with RL (left) and MARVIN trained with IL (right). Both are relatively thorough while expanding to new regions, but the model trained using imitation learning is able to cover the regions in a more efficient manner.



Figure 10. Random graphs sampled from the training set.