

Learning to navigate without a map

Goal Learn a reinforcement policy to navigate from the start to the goal.

Description:

The goal of the project is to determine how well deep learning is suited for planning under incomplete information. You will use the policy gradients algorithm [1] to learn to navigate a grid world presented in [2] (best paper at NIPS' 16). While [2] assumes a fully-observable grid, you will work with the case in which you don't know where the obstacles are in advance, and need to discover that. You will compare this partially-observed scenario with the fully-observed one from [2] to discover how much more difficult it is. You will also compare the deep net performance with the A-star algorithm [3].

[1] <http://karpathy.github.io/2016/05/31/r1/>

[2] <https://papers.nips.cc/paper/6046-value-iteration-networks.pdf>

[3] https://en.wikipedia.org/wiki/A*_search_algorithm

Requirements / Tools:

Required: python, machine learning knowledge

Supervisor:

Nikolay Savinov

nikolay.savinov@inf.ethz.ch




Figure 1: Two instances of a grid-world domain. Task is to move to the goal between the obstacles.