ECE209AS (Fall 2025)

Computational Robotics

Prof. Ankur Mehta ⟨mehtank@ucla.edu⟩

Lecture 3 | Discretization and function approximation

Addendum to lecture videos

Clarification

When coming up with features for function approximation, you have the option of what basis functions to use. You could choose to have a generic, universal, unbiased set of functions (e.g. polynomial regression, Fourier series, neural net), or you could instead hand-craft a set of features targetted towards the problem you are trying to solve (distance to the ice cream store, points scored, etc.). Similar to the dangers of reward-shaping, you will likely influence the policy you discover by your choice of basis functions.

However, you will be using this basis to approximate the value, which explicitly characterizes the policy (= how to achieve the goal), and so this influence is more overt and even sometimes desired (so that you can reduce computation and settle for a not-quite-optimal-but-good-enough policy). In contrast, the reward is meant to characterize the goal (and not how to achieve it); if you changed that based on an assumed policy, you would be incorrect in then claiming that your MDP solution was actually an optimal policy.

Additional References

- This material is covered in Pieter Abbeel's CS287 at UC Berkeley with slightly different notation in lectures 3 and 4.
  - https://people.eecs.berkeley.edu/~pabbeel/cs287-fa19/