Optimization for Locally Optimal Control

Pieter Abbeel
UC Berkeley EECS

Optimal Control (Open Loop)

- Optimal control problem:

$$
\begin{aligned}
&\min_{x,u} \sum_{t=0}^{H} c_t(x_t, u_t) \\
&\text{s.t.} \quad x_0 = \bar{x}_0 \\
&\qquad\quad x_{t+1} = f(x_t, u_t) \qquad t = 0, \ldots, H-1
\end{aligned}
$$

- Solution:
  - = Sequence of controls u and resulting state sequence x
  - If no noise, sufficient to just execute u
- In general non-convex optimization problem, can be solved with sequential convex programming (SCP)

Optimal Control (Closed Loop)

Given: $\bar{x}_0$

For t=0, 1, 2, …, T

- Solve
$$\min_{x,u} \sum_{k=t}^{T} c_k(x_k, u_k)$$
$$\text{s.t.} \quad x_{k+1} = f(x_k, u_k), \qquad \forall k \in \{t, t+1, \dots, T-1\}$$
$$x_t = \bar{x}_t$$
- Execute u_t
- Observe resulting state, $\bar{x}_{t+1}$

= “Model Predictive Control”

Initialize with solution from t - 1 to solve fast at time t

Collocation versus Shooting

- What we considered thus far is a collocation method
  - It considers both x and u simultaneously, optimizes over both of them, and re-linearizes (inside the SCP loop) based on both x and u from the previous round

- Shooting methods
  - Optimize over u directly
  - This can be done as every u results (following the dynamics) in a state sequence x, for which in turn the cost can be computed
  - Upside: Improve sequence of controls over time
    - Versus: collocation might converge to a local optimum that's infeasible
  - Downsides:
    - Derivatives with respect to u as well as the cost for a given u can be numerically unstable to compute (especially in case of unstable dynamical systems)

      [x provides decoupling between time-steps, making computation stable]
    - Not clear how to initialize in a way that nudges towards a goal state