

This is a cheat sheet for the RL-MPC tutorial.

Keep this handy and also feel free to chime in during the tutorial for clarification :-)

Reinforcement learning

s, s'	State	s_t, s_{t+1}
a	Action	a_t
p	State transition probability	$s' \sim p(s' s, a)$
r	Reward	$r_t = r(s_t, a_t)$
π	Policy	$a \sim \pi(a s), a = \pi(s)$
γ	Discount factor	$\gamma \in [0, 1]$
G_t	Discounted return	$G_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k}$
Q^π	State-action value function	$Q^\pi(s, a) = \mathbb{E}[G_0 s_0=s, a_0=a]$
Q^*	Optimal value function	The above, but better
V^π	Value function	$V^\pi(s) = \mathbb{E}[G_0 s_0 = s]$
V^*	Optimal value function	$V^*(s) = \max_a Q^*(s, a)$
π^*	Optimal policy	$\pi^*(s) = \arg \max_a Q^*(s, a)$

Control

x	State	x_t
u	(Control) input	u_t
f	State transition function	$x_{t+1} = f(x_t, u_t)$
ℓ	(Stage) cost	$\ell(x, u) = x^T M x + u^T R u$
K	Gain matrix	$u = -Kx$

Acronyms

RL	Reinforcement learning
MPC	Model predictive control
LQR	Linear quadratic regulator
PID	Proportional-integral-derivative