



Learning how to autonomously race a car: a predictive control approach


Ugo Rosolia

AMBER Lab
California Institute of Technology


June, 2021

Success Stories from AI


Alpha GO



Waymo's Perception Module



OpenAI




Google



Success Stories from Control Theory

Boston Dynamics




Stanford Dynamic Design Lab




Standard Control Pipeline

Optimal Trajectory




Optimization




Dynamic Programming



Trajectory Tracking



Frequency Domain




Nonlinear Control



System Identification





Tire Dynamics

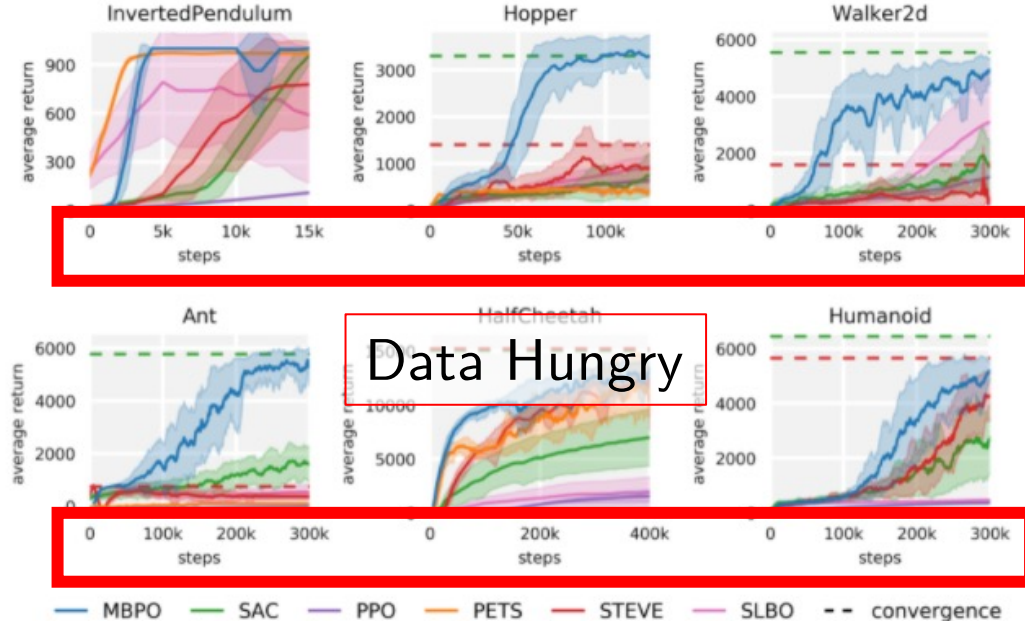


Vehicle Dynamics

Can we simplify the control design?



DeepMind



M. Janner, J. Fu, M. Zhang, and S. Levine. "When to trust your model: Model-based policy optimization." arXiv preprint arXiv:1906.08253 (2019).

Can we simplify the control design?

DeepMind




Force perturbations applied to the torso.



Today's goals:

- ▶ Review standard model-based and model-free RL strategies
- ▶ Design efficient model-based RL framework
- ▶ Summary of the challenges ahead of us

Reinforcement Learning




Today's Example



Learning Model Predictive Controller full-size
vehicle experiments


Credits: Siddharth Nair, Nitin Kapania and Ugo Rosolia

Lessons from Model Predictive Control (MPC)




- ▶ Predicted trajectory given by **Prediction Model**
- ▶ Safe region estimated by the **Safe Set**
- ▶ Predicted cost estimated by **Value Function**

Lessons from Model Predictive Control (MPC)



- ▶ Predicted trajectory given by **Prediction Model** } Identified from historical data
- ▶ Safe region estimated by the **Safe Set**
- ▶ Predicted cost estimated by **Value Function**

Lessons from Model Predictive Control (MPC)



- ▶ Predicted trajectory given by **Prediction Model**
- ▶ Safe region estimated by the **Safe Set**
- ▶ Predicted cost estimated by **Value Function**

Identified from historical data

Estimate these components
to simplify the design

Three key components to learn


Prediction Model

Value Function

Safe Set

Three key components to learn

Prediction Model



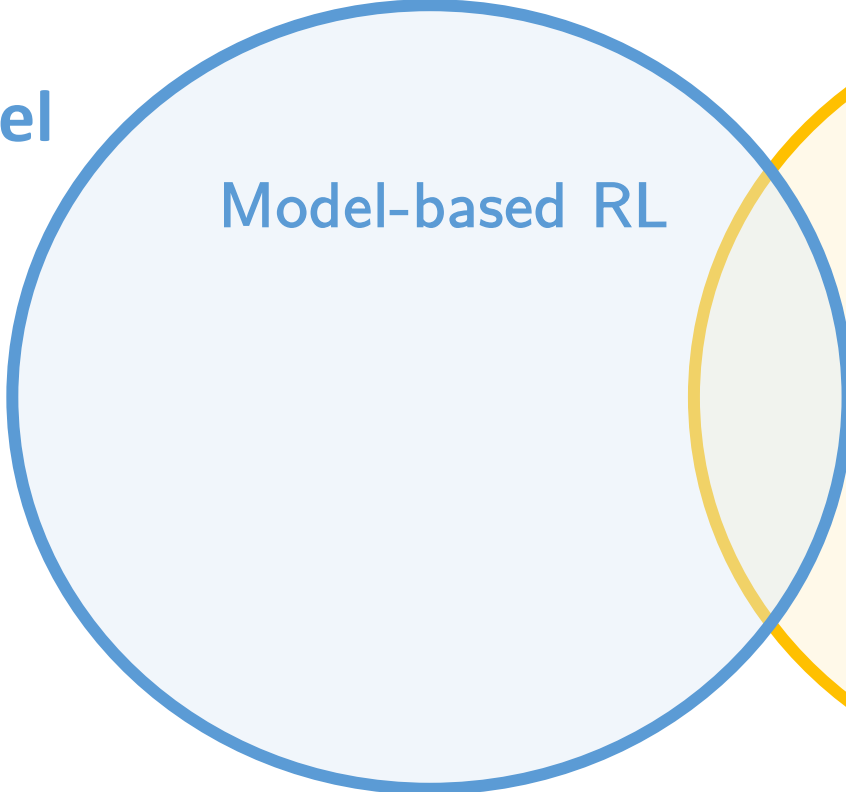
Model-based RL

Value Function

Safe Set

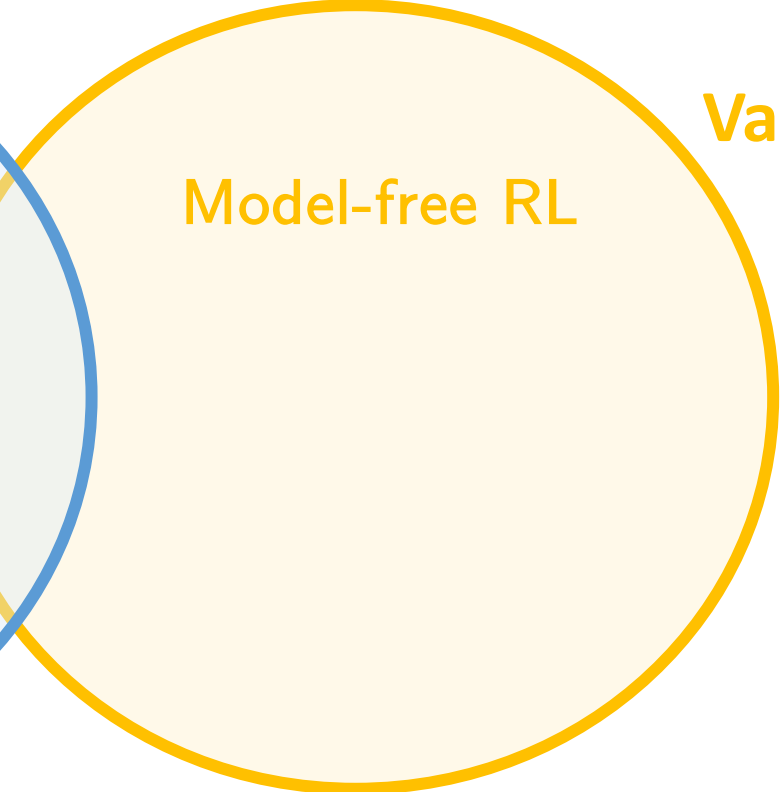
Three key components to learn

Prediction Model



Model-based RL

Value Function




Model-free RL

Safe Set

Three key components to learn

Prediction Model

Value Function



Model-based RL

Model-free RL

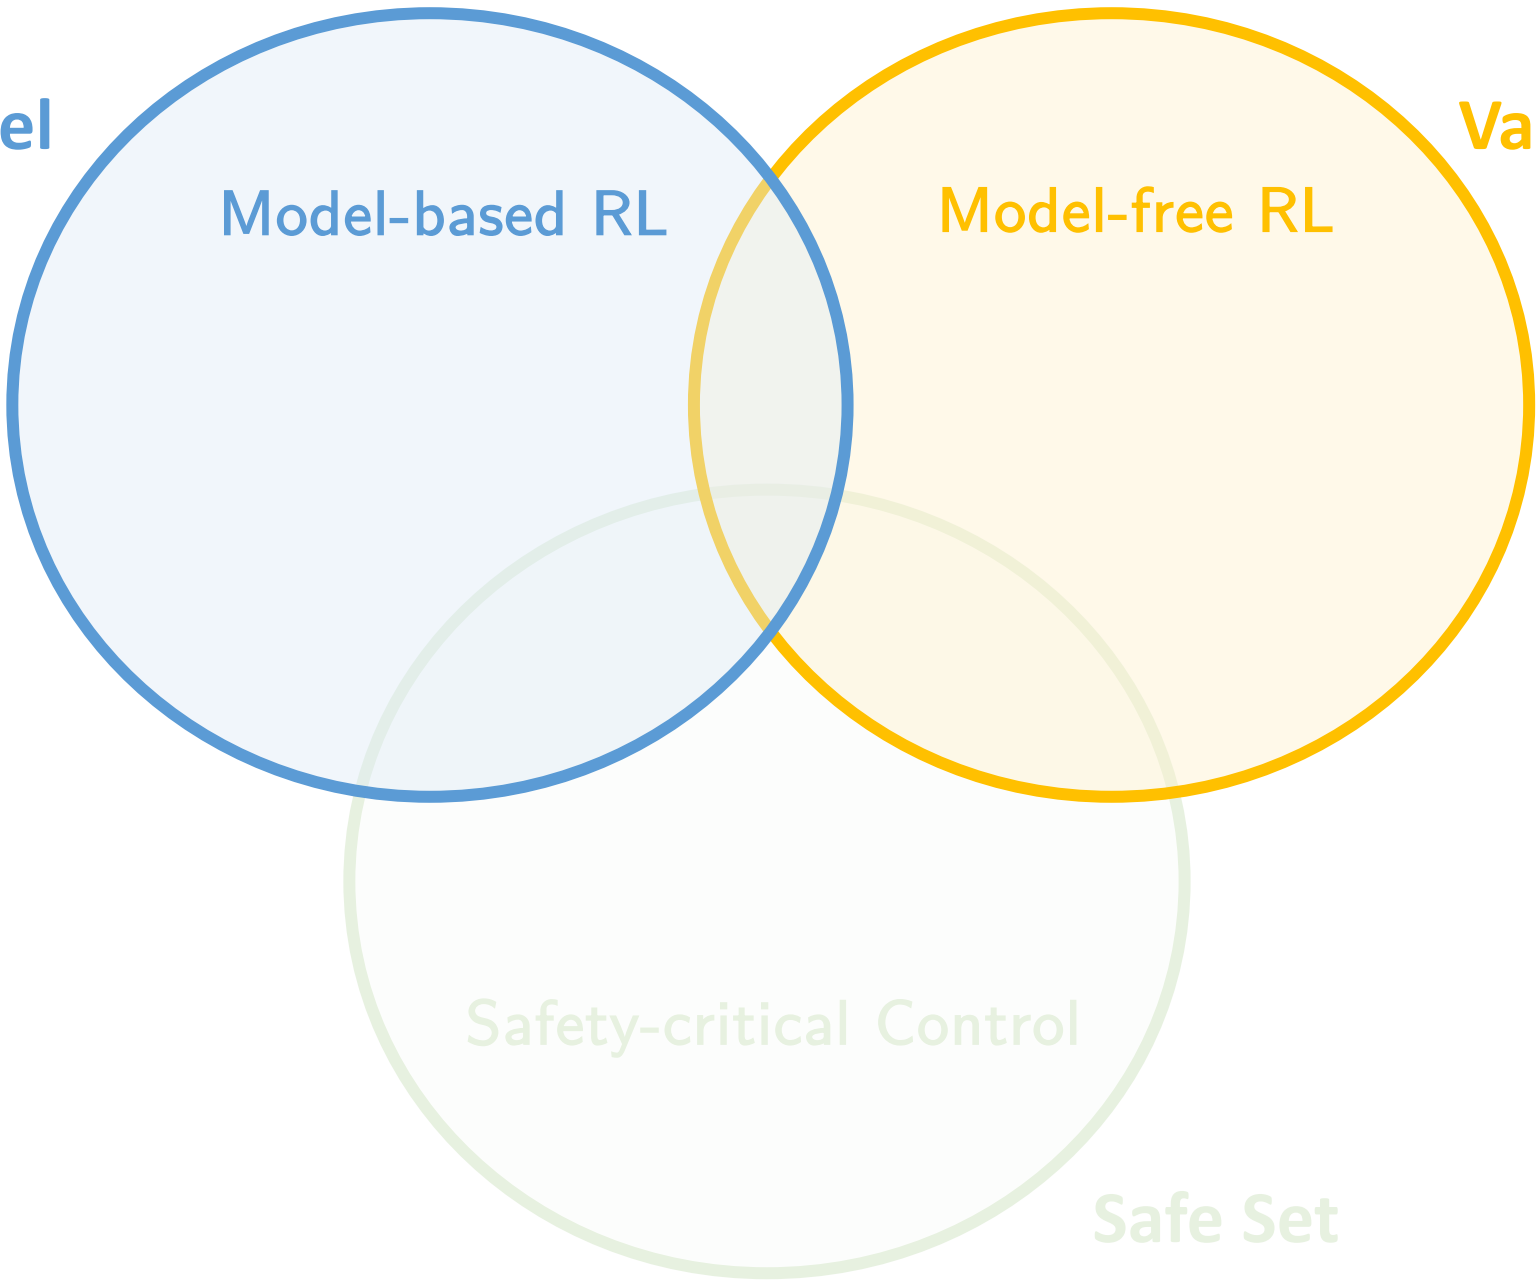
Safety-critical Control

Safe Set

Three key components to learn

Prediction Model

Value Function



Model-based RL

Model-free RL

Safety-critical Control

Safe Set

Theoretical Foundations of Reinforcement Learning



NEURO-DYNAMIC PROGRAMMING

DIMITRI P. BERTSEKAS
JOHN N. TSITSIKLIS



1996

Reinforcement Learning

An Introduction
second edition

Richard S. Sutton and Andrew G. Barto

1998

Reinforcement Learning and Optimal Control

Dimitri P. Bertsekas



2020


Theoretical Foundations of Reinforcement Learning

Principle of Optimality:

$$u^* = \arg \min_{u \in U} [h(x, u) + \mathbb{E}[V^*(x^+) | x, u]]$$

Optimal Control Action Instantaneous cost Future cost Future state

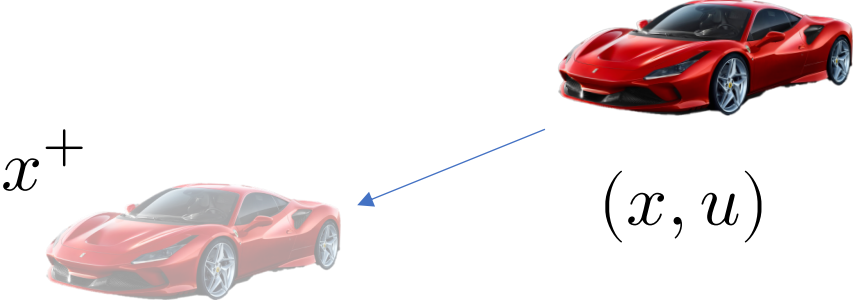
Value function to compute future cost



Map from all possible board configurations to the cost!



Prediction Model to compute x^+




Deep Reinforcement Learning

Principle of Optimality:


$$u^* = \arg \min_{u \in U} [h(x, u) + \mathbb{E}[V^*(x^+) | x, u]]$$

Optimal Control Action Instantaneous cost Future cost Future state

Value function to compute future cost



Prediction Model to compute x^+

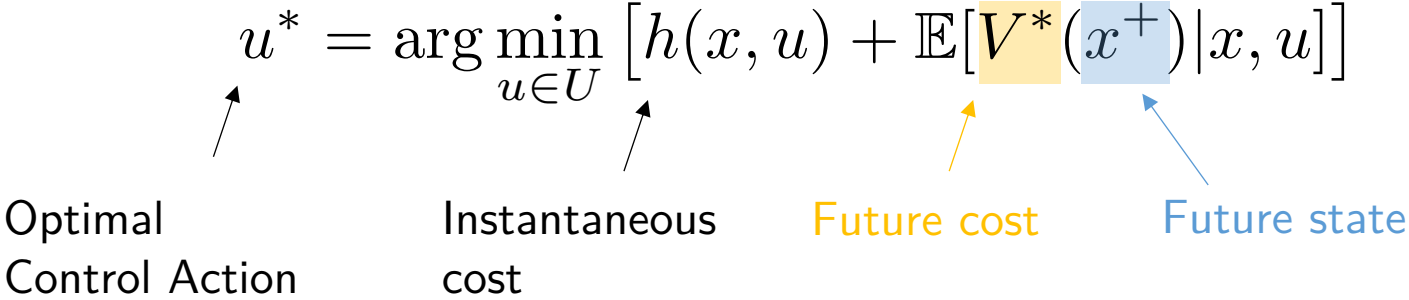


Model-Based vs Model-Free

Principle of Optimality:

$$u^* = \arg \min_{u \in U} [h(x, u) + \mathbb{E}[V^*(x^+) | x, u]]$$

Optimal Control Action Instantaneous cost Future cost Future state



Model-Based vs Model-Free

Principle of Optimality:

$$u^* = \arg \min_{u \in U} [h(x, u) + \mathbb{E}[V^*(x^+) | x, u]]$$

Optimal Control Action Instantaneous cost Future cost Future state

Model-Based RL

$$[u^*, u_1, \dots, u_{N-1}] = \arg \min_{u_k \in U} \mathbb{E} \left[\sum_{k=0}^{N-1} h(x_k, u_k) + V^*(x_N) \right]$$

Model-Based vs Model-Free

Principle of Optimality:

$$u^* = \arg \min_{u \in U} [h(x, u) + \mathbb{E}[V^*(x^+) | x, u]]$$

Optimal Control Action Instantaneous cost Future cost Future state

Model-Based RL

$$[u^*, u_1, \dots, u_{N-1}] = \arg \min_{u_k \in U} \mathbb{E} \left[\sum_{k=0}^{N-1} h(x_k, u_k) + V^*(x_N) \right]$$

The above, for long horizon N is approximated as

$$[u^*, u_1, \dots, u_{N-1}] = \arg \min_{u_k \in U} \mathbb{E} \left[\sum_{k=0}^{N-1} h(x_k, u_k) \right]$$

Model-Based vs Model-Free

Principle of Optimality:

$$u^* = \arg \min_{u \in U} [h(x, u) + \mathbb{E}[V^*(x^+) | x, u]]$$

Optimal Control Action Instantaneous cost Future cost Future state

Model-Based RL

$$[u^*, u_1, \dots, u_{N-1}] = \arg \min_{u_k \in U} \mathbb{E} \left[\sum_{k=0}^{N-1} h(x_k, u_k) + V^*(x_N) \right]$$

The above, for long horizon N is approximated as

$$[u^*, u_1, \dots, u_{N-1}] = \arg \min_{u_k \in U} \mathbb{E} \left[\sum_{k=0}^{N-1} h(x_k, u_k) \right]$$

Model-Free RL

Define the Q-factor:

$$Q^*(x, u) = h(x, u) + \mathbb{E}[V^*(x^+) | x, u]$$


Then the optimal action is

$$u^* = \arg \min_{u \in U} Q^*(x, u)$$

Three key components to learn

Prediction Model

Value Function



Model-based RL

Model-free RL



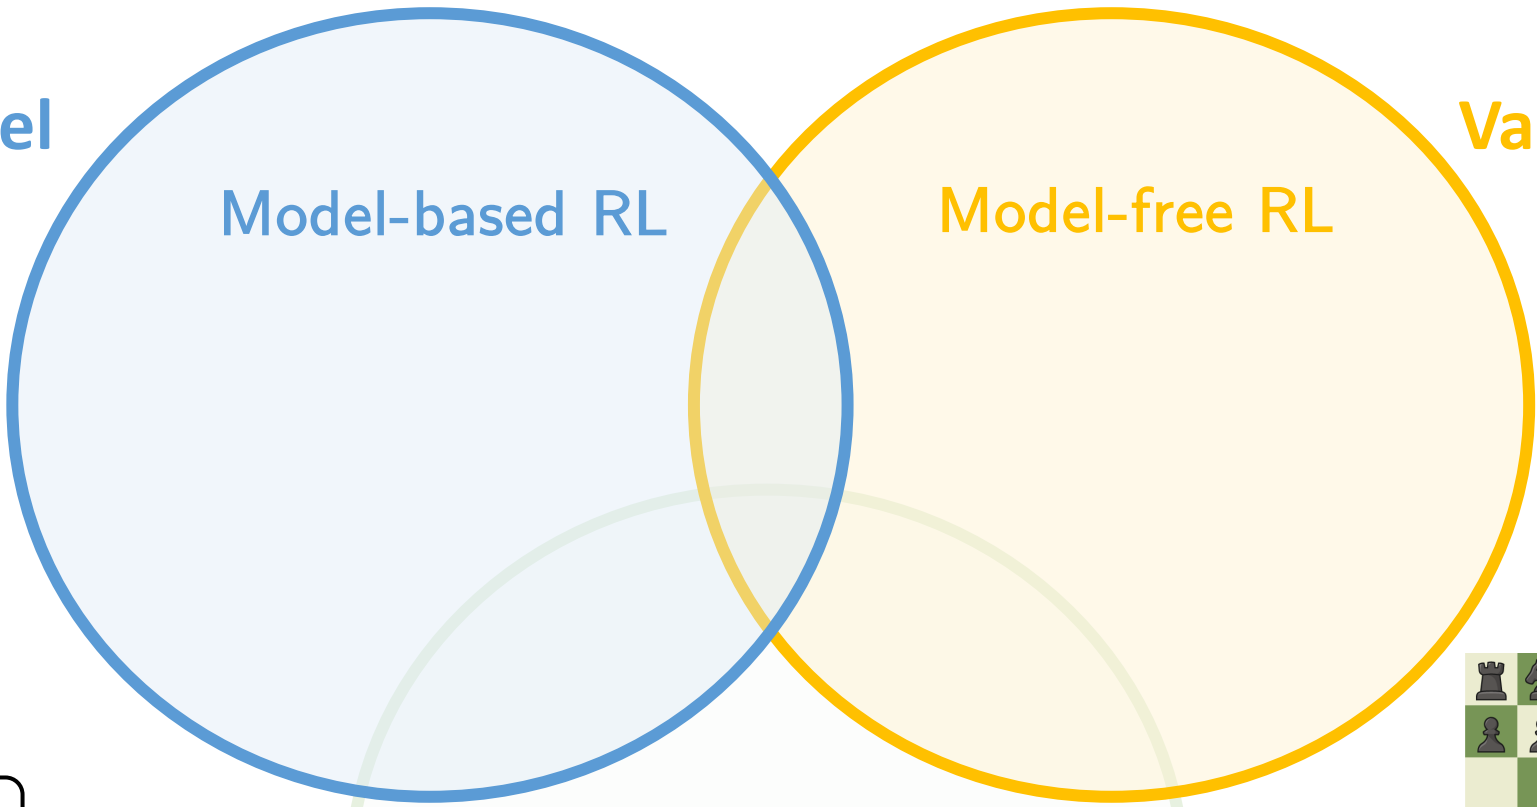
Safety-critical Control

Safe Set

Three key components to learn

Prediction Model

Value Function



Three key components to learn

Prediction Model

Value Function

Model-based RL

Model-free RL

$$\min_{\{u_k\}_{k=0}^N} \mathbb{E} \left[\sum_{t=0}^N h(x_t, u_t) \right] + \min_{u \in \mathcal{U}} Q^*(x, u)$$

Safety-critical Control


Safe Set




What is different in safety-critical systems?

What is different in safety-critical systems? Constraints

What is different in safety-critical systems? Constraints




What is different in safety-critical systems? Constraints



► State

$$x = \begin{bmatrix} p \\ v \end{bmatrix} = \begin{bmatrix} \text{position} \\ \text{velocity} \end{bmatrix}$$

What is different in safety-critical systems? Constraints




► State

$$x = \begin{bmatrix} p \\ v \end{bmatrix} = \begin{bmatrix} \text{position} \\ \text{velocity} \end{bmatrix}$$

► Input $u = a = \text{acceleration}$

What is different in safety-critical systems? Constraints



► State


$$x = \begin{bmatrix} p \\ v \end{bmatrix} = \begin{bmatrix} \text{position} \\ \text{velocity} \end{bmatrix}$$

► Input $u = a =$ acceleration

► Dynamics

$$\begin{bmatrix} p_{k+1} \\ v_{k+1} \end{bmatrix} = \begin{bmatrix} 1 & dt \\ 0 & 1 \end{bmatrix} \begin{bmatrix} p_k \\ v_k \end{bmatrix} + \begin{bmatrix} 0 \\ a_k \end{bmatrix}$$

What is different in safety-critical systems? Constraints



► State

$$x = \begin{bmatrix} p \\ v \end{bmatrix} = \begin{bmatrix} \text{position} \\ \text{velocity} \end{bmatrix}$$


► Input $u = a =$ acceleration

► Dynamics

$$\begin{bmatrix} p_{k+1} \\ v_{k+1} \end{bmatrix} = \begin{bmatrix} 1 & dt \\ 0 & 1 \end{bmatrix} \begin{bmatrix} p_k \\ v_k \end{bmatrix} + \begin{bmatrix} 0 \\ a_k \end{bmatrix}$$

► Cost $x_k^\top Q x_k + u_k^\top R u_k$

What is different in safety-critical systems? Constraints



► State

$$x = \begin{bmatrix} p \\ v \end{bmatrix} = \begin{bmatrix} \text{position} \\ \text{velocity} \end{bmatrix}$$

► Input $u = a =$ acceleration

► Dynamics


$$\begin{bmatrix} p_{k+1} \\ v_{k+1} \end{bmatrix} = \begin{bmatrix} 1 & dt \\ 0 & 1 \end{bmatrix} \begin{bmatrix} p_k \\ v_k \end{bmatrix} + \begin{bmatrix} 0 \\ a_k \end{bmatrix}$$

► Cost $x_k^\top Q x_k + u_k^\top R u_k$

► Constraints

$$\begin{bmatrix} -5 \\ -5 \\ -0.5 \end{bmatrix} \leq \begin{bmatrix} p_k \\ v_k \\ a_k \end{bmatrix} \leq \begin{bmatrix} 5 \\ 5 \\ 0.5 \end{bmatrix}$$

What is different in safety-critical systems? Constraints



► State

$$x = \begin{bmatrix} p \\ v \end{bmatrix} = \begin{bmatrix} \text{position} \\ \text{velocity} \end{bmatrix}$$

► Input $u = a =$ acceleration

► Dynamics

$$\begin{bmatrix} p_{k+1} \\ v_{k+1} \end{bmatrix} = \begin{bmatrix} 1 & dt \\ 0 & 1 \end{bmatrix} \begin{bmatrix} p_k \\ v_k \end{bmatrix} + \begin{bmatrix} 0 \\ a_k \end{bmatrix}$$


► Cost $x_k^T Q x_k + u_k^T R u_k$

► Constraints

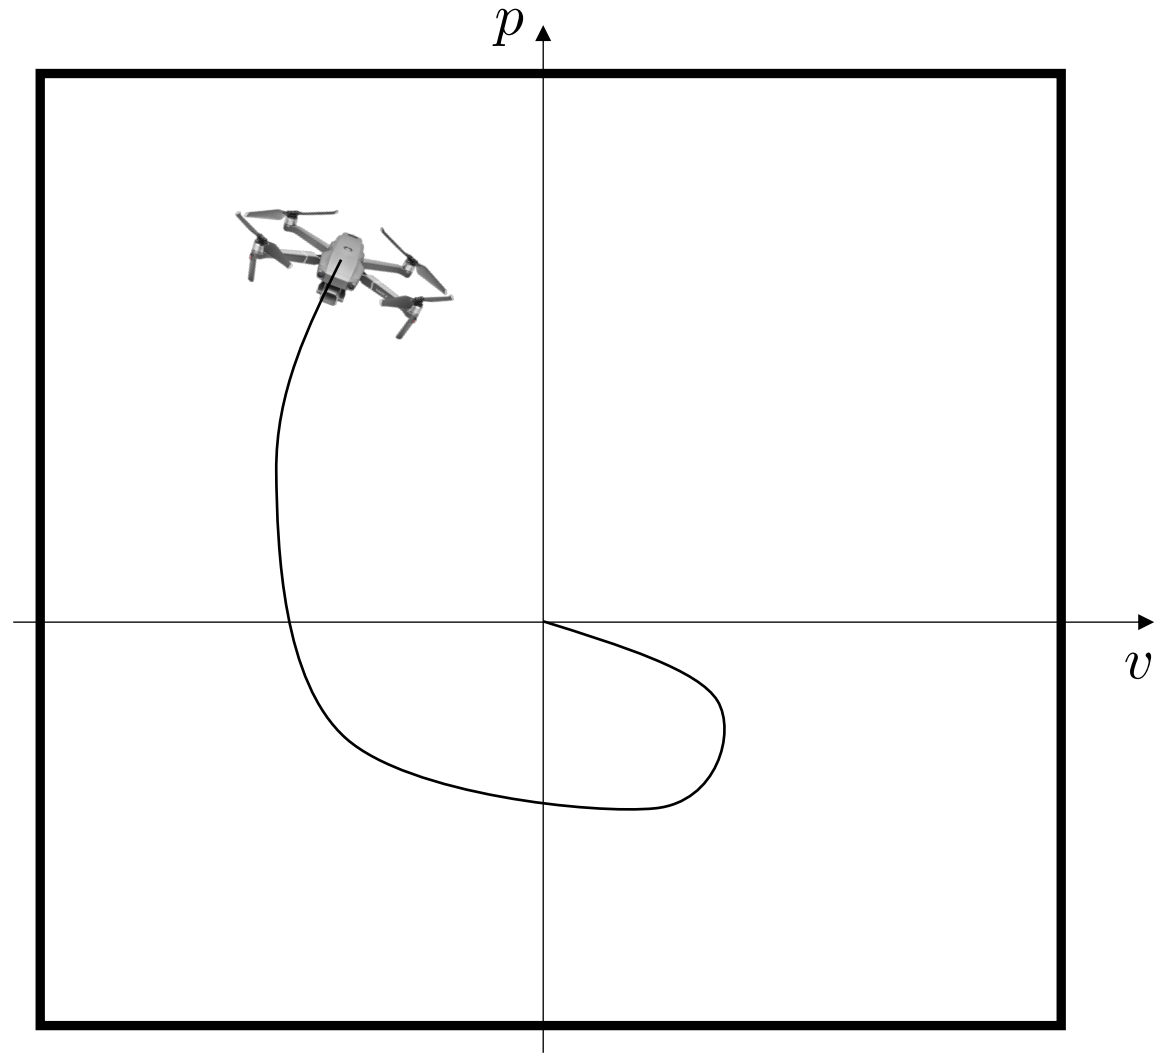
$$\begin{bmatrix} -5 \\ -5 \\ -0.5 \end{bmatrix} \leq \begin{bmatrix} p_k \\ v_k \\ a_k \end{bmatrix} \leq \begin{bmatrix} 5 \\ 5 \\ 0.5 \end{bmatrix}$$

Limited actuation!

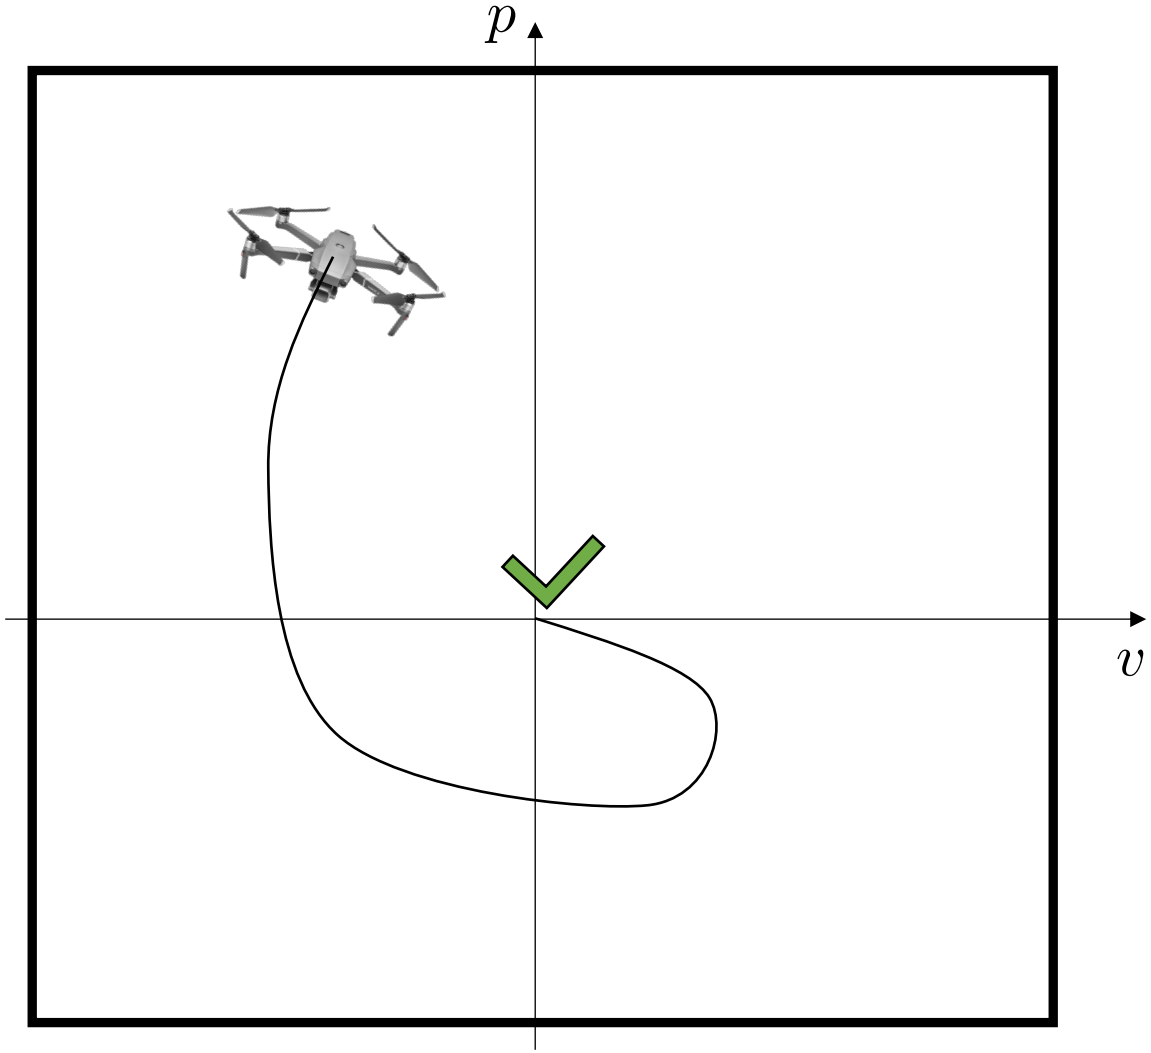
What is different in safety-critical systems? Constraints




What is different in safety-critical systems? Constraints




What is different in safety-critical systems? Constraints




What is different in safety-critical systems? Constraints




What is different in safety-critical systems? Constraints



What is different in safety-critical systems? Constraints




What is different in safety-critical systems? Constraints



Driving the drone to the origin is **impossible** due to inertia and input saturation


What is different in safety-critical systems? Constraints



Driving the drone to the origin is **impossible** due to inertia and input saturation

The drone can be driven to the origin only from a **subset** of the feasible set

What is different in safety-critical systems? Constraints





Driving the drone to the origin is **impossible** due to inertia and input saturation

The drone can be driven to the origin only from a **subset** of the feasible set

Key Message: We need to approximate the value function only over a subset of the feasible set

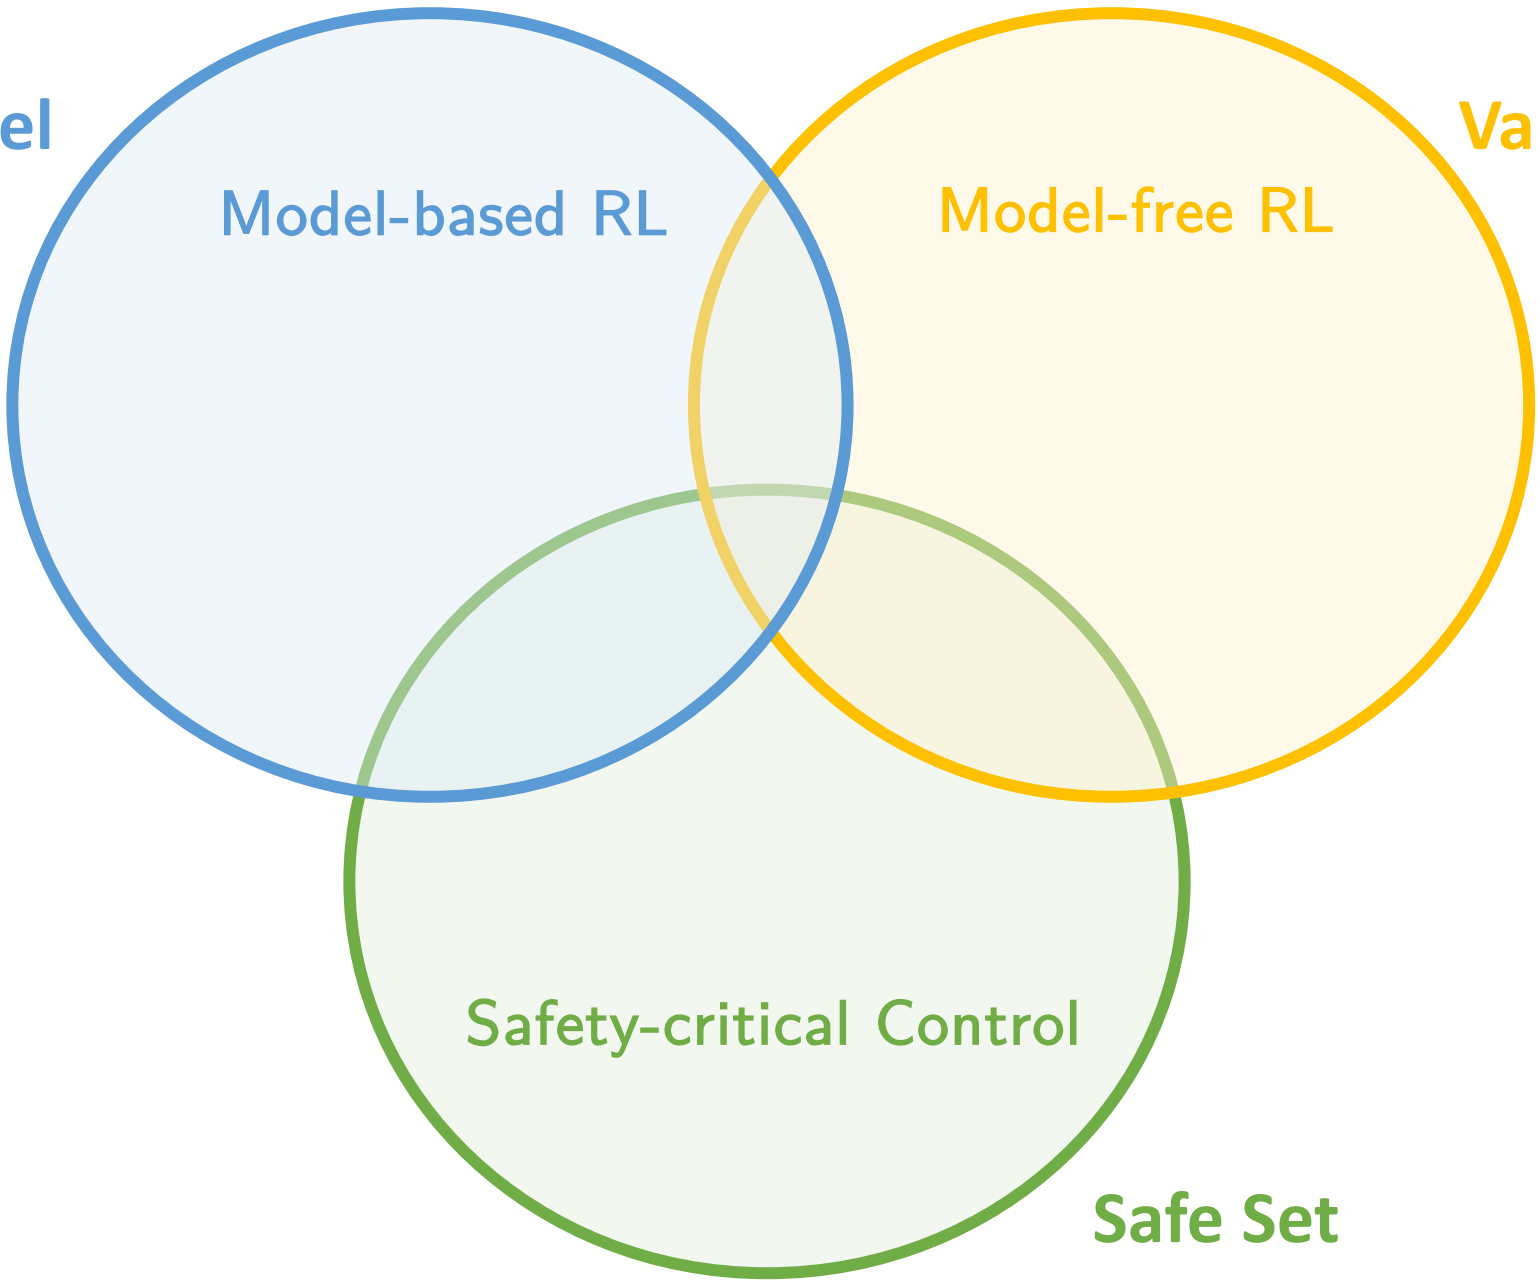
Computation of Safe Sets in the Control



Three key components to learn

Prediction Model

Value Function



Model-based RL

Model-free RL


Safety-critical Control

Safe Set

Three key components to learn

Prediction Model

Value Function



Model-based RL


Model-free RL

Safety-critical Control

Data Efficient Learning!


Safe Set

Learning Model Predictive Controller



Learning Model Predictive Controller


At time t of lap j solve the following Constrained Finite Time Optimal Control Problem (CFTOCP)



Learning Model Predictive Controller

At time t of lap j solve the following Constrained Finite Time Optimal Control Problem (CFTOCP)

$$J_{0 \rightarrow N}^{\text{LMPC},j}(x(t)) = \min_{u_t, \dots, u_{N-1}} \sum_{k=0}^{N-1} h(x_k, u_k) + V^{j-1}(x_N, \boldsymbol{x})$$




Learning Model Predictive Controller

At time t of lap j solve the following Constrained Finite Time Optimal Control Problem (CFTOCP)

$$J_{0 \rightarrow N}^{\text{LMPC},j}(x(t)) = \min_{u_t, \dots, u_{N-1}} \sum_{k=0}^{N-1} h(x_k, u_k) + V^{j-1}(x_N, \mathbf{x})$$

Value Function



Learning Model Predictive Controller

At time t of lap j solve the following Constrained Finite Time Optimal Control Problem (CFTOCP)


$$J_{0 \rightarrow N}^{\text{LMPC},j}(x(t)) = \min_{u_t, \dots, u_{N-1}} \sum_{k=0}^{N-1} h(x_k, u_k) + V^{j-1}(x_N, \mathbf{x})$$

s.t.

$$x_{k+1} = A_k x_k + B_k u_k + C_k,$$

$$x_t = x(t),$$

Value Function



Learning Model Predictive Controller

At time t of lap j solve the following Constrained Finite Time Optimal Control Problem (CFTOCP)

$$J_{0 \rightarrow N}^{\text{LMPC},j}(x(t)) = \min_{u_t, \dots, u_{N-1}} \sum_{k=0}^{N-1} h(x_k, u_k) + V^{j-1}(x_N, \mathbf{x})$$


s.t.

$$x_{k+1} = A_k x_k + B_k u_k + C_k,$$

$$x_t = x(t),$$

Value Function

Prediction
Model



Learning Model Predictive Controller

At time t of lap j solve the following Constrained Finite Time Optimal Control Problem (CFTOCP)

$$J_{0 \rightarrow N}^{\text{LMPC},j}(x(t)) = \min_{u_t, \dots, u_{N-1}} \sum_{k=0}^{N-1} h(x_k, u_k) + V^{j-1}(x_N, \mathbf{x})$$

s.t.


$$x_{k+1} = A_k x_k + B_k u_k + C_k,$$

$$x_t = x(t),$$

$$x_k \in \mathcal{X}, u_k \in \mathcal{U}, \forall k \in [0, \dots, N-1]$$

Prediction
Model

Value Function



Learning Model Predictive Controller

At time t of lap j solve the following Constrained Finite Time Optimal Control Problem (CFTOCP)

$$J_{0 \rightarrow N}^{\text{LMPC},j}(x(t)) = \min_{u_t, \dots, u_{N-1}} \sum_{k=0}^{N-1} h(x_k, u_k) + V^{j-1}(x_N, \mathbf{x})$$

s.t.

$$x_{k+1} = A_k x_k + B_k u_k + C_k,$$


$$x_t = x(t),$$

$$x_k \in \mathcal{X}, u_k \in \mathcal{U}, \forall k \in [0, \dots, N-1]$$

$$x_N \in \mathcal{CS}^{j-1}(\mathbf{x}),$$

Prediction
Model

Value Function



Learning Model Predictive Controller

At time t of lap j solve the following Constrained Finite Time Optimal Control Problem (CFTOCP)

$$J_{0 \rightarrow N}^{\text{LMPC},j}(x(t)) = \min_{u_t, \dots, u_{N-1}} \sum_{k=0}^{N-1} h(x_k, u_k) + V^{j-1}(x_N, \boldsymbol{x})$$

s.t.

$$x_{k+1} = A_k x_k + B_k u_k + C_k,$$

$$x_t = x(t),$$


$$x_k \in \mathcal{X}, u_k \in \mathcal{U}, \forall k \in [0, \dots, N-1]$$

$$x_N \in \mathcal{CS}^{j-1}(\boldsymbol{x}),$$

Value Function

Safe Region

Prediction Model



Learning Model Predictive Controller

At time t of lap j solve the following Constrained Finite Time Optimal Control Problem (CFTOCP)

$$J_{0 \rightarrow N}^{\text{LMPC},j}(x(t)) = \min_{u_t, \dots, u_{N-1}} \sum_{k=0}^{N-1} h(x_k, u_k) + V^{j-1}(x_N, \boldsymbol{x})$$

s.t.


$$x_{k+1} = A_k x_k + B_k u_k + C_k,$$

$$x_t = x(t),$$

Prediction
Model

$$x_k \in \mathcal{X}, u_k \in \mathcal{U}, \forall k \in [0, \dots, N-1]$$

$$x_N \in \mathcal{CS}^{j-1}(\boldsymbol{x}),$$



Learning Model Predictive Controller

At time t of lap j solve the following Constrained Finite Time Optimal Control Problem (CFTOCP)

$$J_{0 \rightarrow N}^{\text{LMPC},j}(x(t)) = \min_{u_t, \dots, u_{N-1}} \sum_{k=0}^{N-1} h(x_k, u_k) + V^{j-1}(x_N, \boldsymbol{x})$$

s.t.

$$x_{k+1} = A_k x_k + B_k u_k + C_k,$$


$$x_t = x(t),$$

Prediction
Model

$$x_k \in \mathcal{X}, u_k \in \mathcal{U}, \forall k \in [0, \dots, N-1]$$

$$x_N \in \mathcal{CS}^{j-1}(\boldsymbol{x}),$$

In this topic area you will learn how to leverage DNN to estimate system dynamics



System ID in Autonomous Racing

- ▶ Nonlinear Dynamical System,

$$\ddot{x} = \dot{y}\dot{\psi} + \frac{1}{m} \sum_i F_{x_i}$$

$$\ddot{y} = -\dot{x}\dot{\psi} + \frac{1}{m} \sum_i F_{y_i}$$

$$\ddot{\psi} = \frac{1}{I_z} (a(F_{y_{1,2}}) - b(F_{y_{2,3}}) + c(-F_{x_{1,3}} + F_{x_{2,4}}))$$

$$\dot{X} = \dot{x} \cos \psi - \dot{y} \sin \psi, \quad \dot{Y} = \dot{x} \sin \psi + \dot{y} \cos \psi$$

System ID in Autonomous Racing

- ▶ Nonlinear Dynamical System,

$$\ddot{x} = \dot{y}\dot{\psi} + \frac{1}{m} \sum_i F_{x_i}$$

$$\ddot{y} = -\dot{x}\dot{\psi} + \frac{1}{m} \sum_i F_{y_i}$$

$$\ddot{\psi} = \frac{1}{I_z} (a(F_{y_{1,2}}) - b(F_{y_{2,3}}) + c(-F_{x_{1,3}} + F_{x_{2,4}}))$$

$$\dot{X} = \dot{x} \cos \psi - \dot{y} \sin \psi, \quad \dot{Y} = \dot{x} \sin \psi + \dot{y} \cos \psi$$

Kinematic Equations

System ID in Autonomous Racing

- ▶ Nonlinear Dynamical System,

$$\ddot{x} = \dot{y}\dot{\psi} + \frac{1}{m} \sum_i F_{x_i}$$

$$\ddot{y} = -\dot{x}\dot{\psi} + \frac{1}{m} \sum_i F_{y_i}$$

$$\ddot{\psi} = \frac{1}{I_z} (a(F_{y_{1,2}}) - b(F_{y_{2,3}}) + c(-F_{x_{1,3}} + F_{x_{2,4}}))$$

$$\dot{X} = \dot{x} \cos \psi - \dot{y} \sin \psi, \quad \dot{Y} = \dot{x} \sin \psi + \dot{y} \cos \psi$$

Kinematic Equations

- ▶ Identifying the Dynamical System

$$x_{k+1|t}^j = \begin{bmatrix} \dot{x}_{k+1|t} \\ \dot{y}_{k+1|t} \\ \ddot{\psi} \\ \psi_{k+1|t} \\ X_{k+1|t} \\ Y_{k+1|t} \end{bmatrix} = \begin{bmatrix} \dot{x}_{k+1|t} \\ \dot{y}_{k+1|t} \\ \ddot{\psi} \\ \psi_{k+1|t} \\ X_{k+1|t} \\ Y_{k+1|t} \end{bmatrix} x_{k|t}^j + \begin{bmatrix} \dot{x}_{k+1|t} \\ \dot{y}_{k+1|t} \\ \ddot{\psi} \\ \psi_{k+1|t} \\ X_{k+1|t} \\ Y_{k+1|t} \end{bmatrix} \begin{bmatrix} u_{k|t}^j \\ 1 \end{bmatrix}$$

Linearization around predicted trajectory

System ID in Autonomous Racing

- ▶ Nonlinear Dynamical System,

$$\begin{aligned} \ddot{x} &= \dot{y}\dot{\psi} + \frac{1}{m} \sum_i F_{x_i} \\ \ddot{y} &= -\dot{x}\dot{\psi} + \frac{1}{m} \sum_i F_{y_i} \\ \ddot{\psi} &= \frac{1}{I_z} (a(F_{y_{1,2}}) - b(F_{y_{2,3}}) + c(-F_{x_{1,3}} + F_{x_{2,4}})) \\ \dot{X} &= \dot{x} \cos \psi - \dot{y} \sin \psi, \quad \dot{Y} = \dot{x} \sin \psi + \dot{y} \cos \psi \end{aligned}$$

Dynamic Equations

Kinematic Equations

- ▶ Identifying the Dynamical System

$$x_{k+1|t}^j = \begin{bmatrix} \dot{x}_{k+1|t} \\ \dot{y}_{k+1|t} \\ \ddot{\psi} \\ \psi_{k+1|t} \\ X_{k+1|t} \\ Y_{k+1|t} \end{bmatrix} = \begin{bmatrix} \text{Linearized Kinematics} \\ \text{Linearized Kinematics} \\ \text{Linearized Kinematics} \end{bmatrix} x_{k|t}^j + \begin{bmatrix} \text{Linearized Kinematics} \\ \text{Linearized Kinematics} \\ \text{Linearized Kinematics} \end{bmatrix} \begin{bmatrix} u_{k|t}^j \\ 1 \end{bmatrix}$$

Linearization around predicted trajectory

System ID in Autonomous Racing

► Nonlinear Dynamical System,

$$\begin{aligned} \ddot{x} &= \dot{y}\dot{\psi} + \frac{1}{m} \sum_i F_{x_i} \\ \ddot{y} &= -\dot{x}\dot{\psi} + \frac{1}{m} \sum_i F_{y_i} \\ \ddot{\psi} &= \frac{1}{I_z} (a(F_{y_{1,2}}) - b(F_{y_{2,3}}) + c(-F_{x_{1,3}} + F_{x_{2,4}})) \\ \dot{X} &= \dot{x} \cos \psi - \dot{y} \sin \psi, \quad \dot{Y} = \dot{x} \sin \psi + \dot{y} \cos \psi \end{aligned}$$

Dynamic Equations

Kinematic Equations

► Identifying the Dynamical System

Local Linear Regression

$$x_{k+1|t}^j = \begin{bmatrix} \dot{x}_{k+1|t} \\ \dot{y}_{k+1|t} \\ \ddot{\psi} \\ \psi_{k+1|t} \\ X_{k+1|t} \\ Y_{k+1|t} \end{bmatrix} = \begin{bmatrix} \arg \min_{i,s} \sum K(x_{k|t}^j - x_s^i) || \Lambda_y \begin{bmatrix} x_{k|t}^j \\ u_{k|t}^j \\ 1 \end{bmatrix} - y_{s+1}^i ||, \forall y \in \{\dot{x}, \dot{y}, \ddot{\psi}\} \\ \text{Linearized Kinematics} \\ \text{Linearized Kinematics} \\ \text{Linearized Kinematics} \end{bmatrix} x_{k|t}^j + \begin{bmatrix} \text{Linearized Kinematics} \\ \text{Linearized Kinematics} \\ \text{Linearized Kinematics} \end{bmatrix} \begin{bmatrix} u_{k|t}^j \\ 1 \end{bmatrix}$$

Linearization around predicted trajectory

Learning Model Predictive Controller

At time t of iteration j solve the following Constrained Finite Time Optimal Control Problem (CFTOCP)

$$J_{0 \rightarrow N}^{\text{LMPC},j}(x_t) = \min_{u_t, \dots, u_{N-1}} \sum_{k=0}^{N-1} h(x_k, u_k) + V^{j-1}(x_N, \boldsymbol{x})$$

s.t.


$$x_{k+1} = A_k x_k + B_k u_k + C_k,$$

$$x_t = x_t,$$

$$x_k \in \mathcal{X}, u_k \in \mathcal{U}, \forall k \in [0, \dots, N-1]$$

$$x_N \in \mathcal{CS}^{j-1}(\boldsymbol{x}),$$

Prediction
Model



Learning Model Predictive Controller

At time t of iteration j solve the following Constrained Finite Time Optimal Control Problem (CFTOCP)

$$J_{0 \rightarrow N}^{\text{LMPC},j}(x_t) = \min_{u_t, \dots, u_{N-1}} \sum_{k=0}^{N-1} h(x_k, u_k) + V^{j-1}(x_N, \boldsymbol{x})$$

s.t.

$$x_{k+1} = A_k x_k + B_k u_k + C_k,$$

$$x_t = x_t,$$

$$x_k \in \mathcal{X}, u_k \in \mathcal{U}, \forall k \in [0, \dots, N-1]$$

$$x_N \in \mathcal{CS}^{j-1}(\boldsymbol{x}),$$

Safe Set


Value Function



Learning Model Predictive Controller full-size vehicle experiments

Credits: Siddharth Nair, Nitin Kapania and Ugo Rosolia

Lap Time



The **control policy** is constructed using **~1k** data points (last 2 laps)

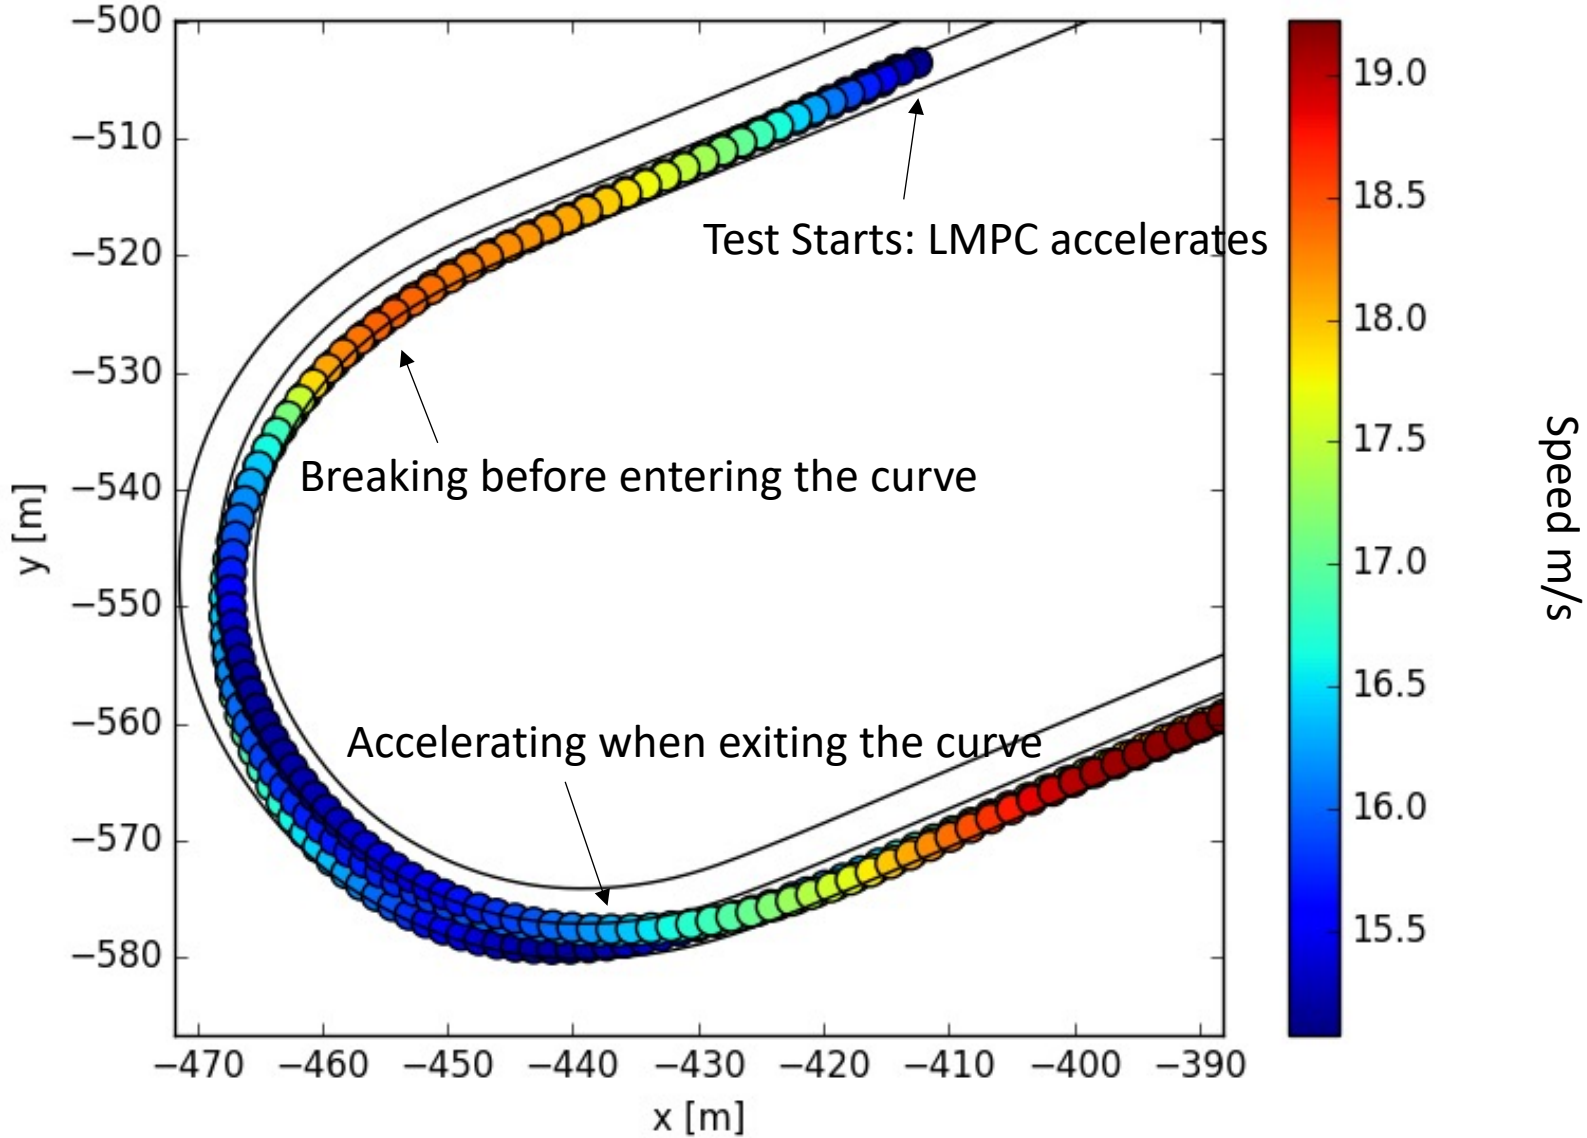
The **control action** is computed using **~100** data points




Learning Model Predictive Controller full-size vehicle experiments

Credits: Siddharth Nair, Nitin Kapania and Ugo Rosolia

Velocity Profile at Convergence (Curve 1)




Velocity Profile at Convergence (Chicane)



The key components


- ▶ Predicted trajectory given by **prediction model**
- ▶ Predicted cost estimated by **value function**
- ▶ Safe region estimated by the **safe set**




What is next?

- ▶ Partial Observability


- ▶ Multi-agent systems



- ▶ Hierarchy + Learning

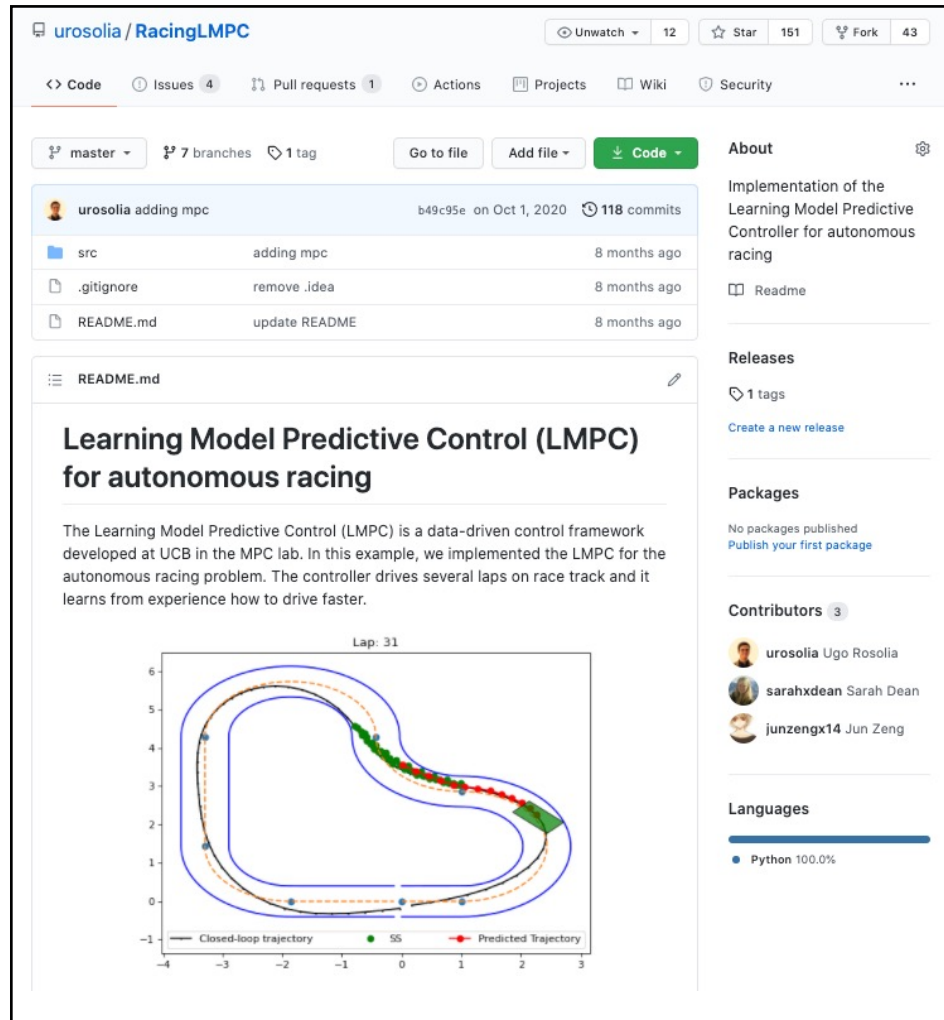


- ▶ Optimize over strategies, not trajectories



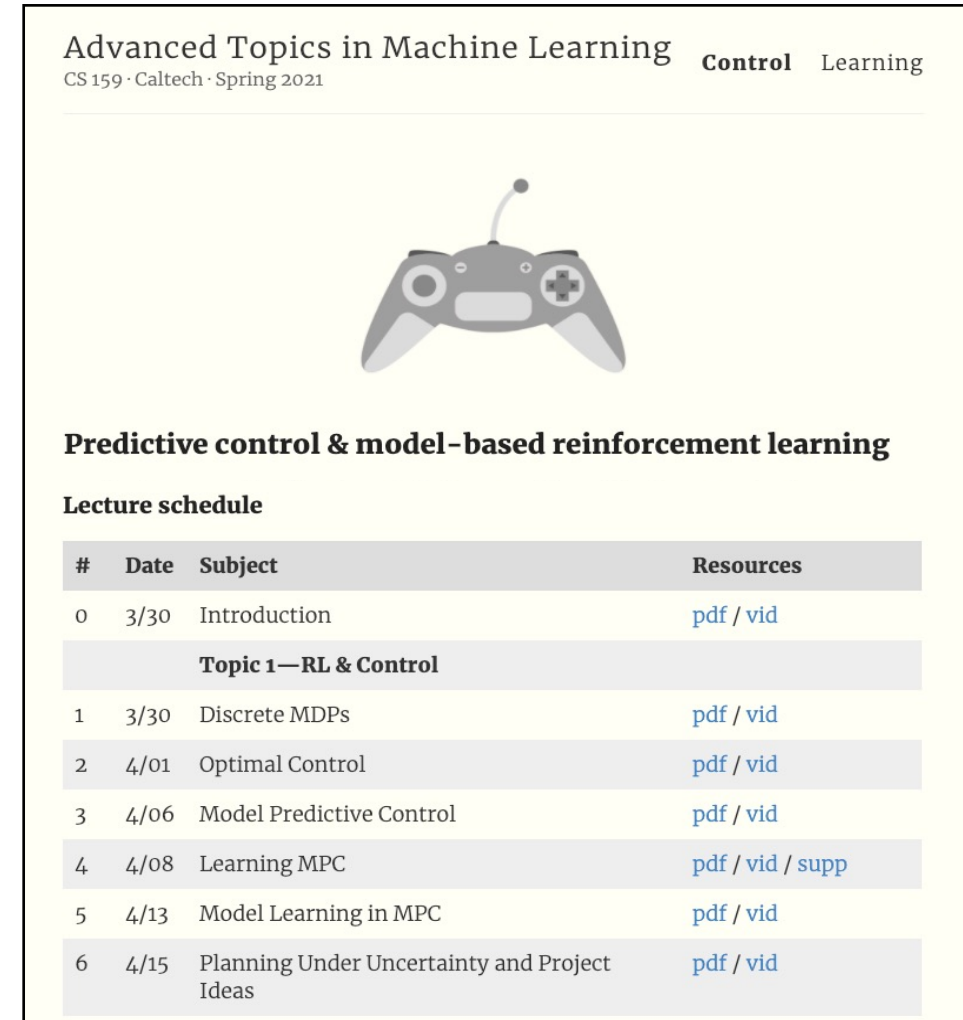
Thanks! Questions?

Code available online



The screenshot shows the GitHub repository page for 'urosolia / RacingLMPC'. The repository has 12 Unwatch, 151 Stars, and 43 Forks. It features a 'Code' button and a 'README.md' file. The README file is titled 'Learning Model Predictive Control (LMPC) for autonomous racing' and includes a plot of a race track trajectory. The plot shows a closed-loop trajectory (blue line) and a predicted trajectory (red line) with green dots representing sensor data (SS). The plot is labeled 'Lap: 31' and has axes ranging from -4 to 3 on the x-axis and -1 to 6 on the y-axis. The legend indicates 'Closed-loop trajectory' (blue line), 'SS' (green dots), and 'Predicted Trajectory' (red line).

Course material online



The screenshot shows the course page for 'Advanced Topics in Machine Learning' (CS 159) at Caltech, Spring 2021. The page features a controller icon and the text 'Control Learning'. The main heading is 'Predictive control & model-based reinforcement learning'. Below this is a 'Lecture schedule' table.

#	Date	Subject	Resources
0	3/30	Introduction	pdf / vid
Topic 1—RL & Control			
1	3/30	Discrete MDPs	pdf / vid
2	4/01	Optimal Control	pdf / vid
3	4/06	Model Predictive Control	pdf / vid
4	4/08	Learning MPC	pdf / vid / supp
5	4/13	Model Learning in MPC	pdf / vid
6	4/15	Planning Under Uncertainty and Project Ideas	pdf / vid