Dealing with complex cost functions: Trajectory library

Instructor: Chris Mavrogiannis

TAs: Kay Ke, Gilwoo Lee, Matt Schmittle

*Slides based on or adapted from Sanjiban Choudhury

Evolution of controllers

	Uses model	Stability Guarantee	Minimize Cost
PID	No	No	No
Pure Pursuit	Circular arcs	Yes - with assumptions	No
Lyapunov	Non-linear	Yes	No
LQR	Linear	Yes	Quadratic
MPC	Non-linear	Yes	Yes

Evolution of controllers

Increasing complexity of problem ....

Heuristic Rules

PID

Trajectory Optimization

iLQR

Recap: Model predictive control (MPC)

$$\min_{u_{t+1},\ldots u_{t+H}} \sum_{k=t}^{t+H-1} J(x_k, u_{k+1})$$

$$x_{k+1} = f(x_k, u_{k+1})$$

$$g(x_k, u_{k+1}) \leq 0$$

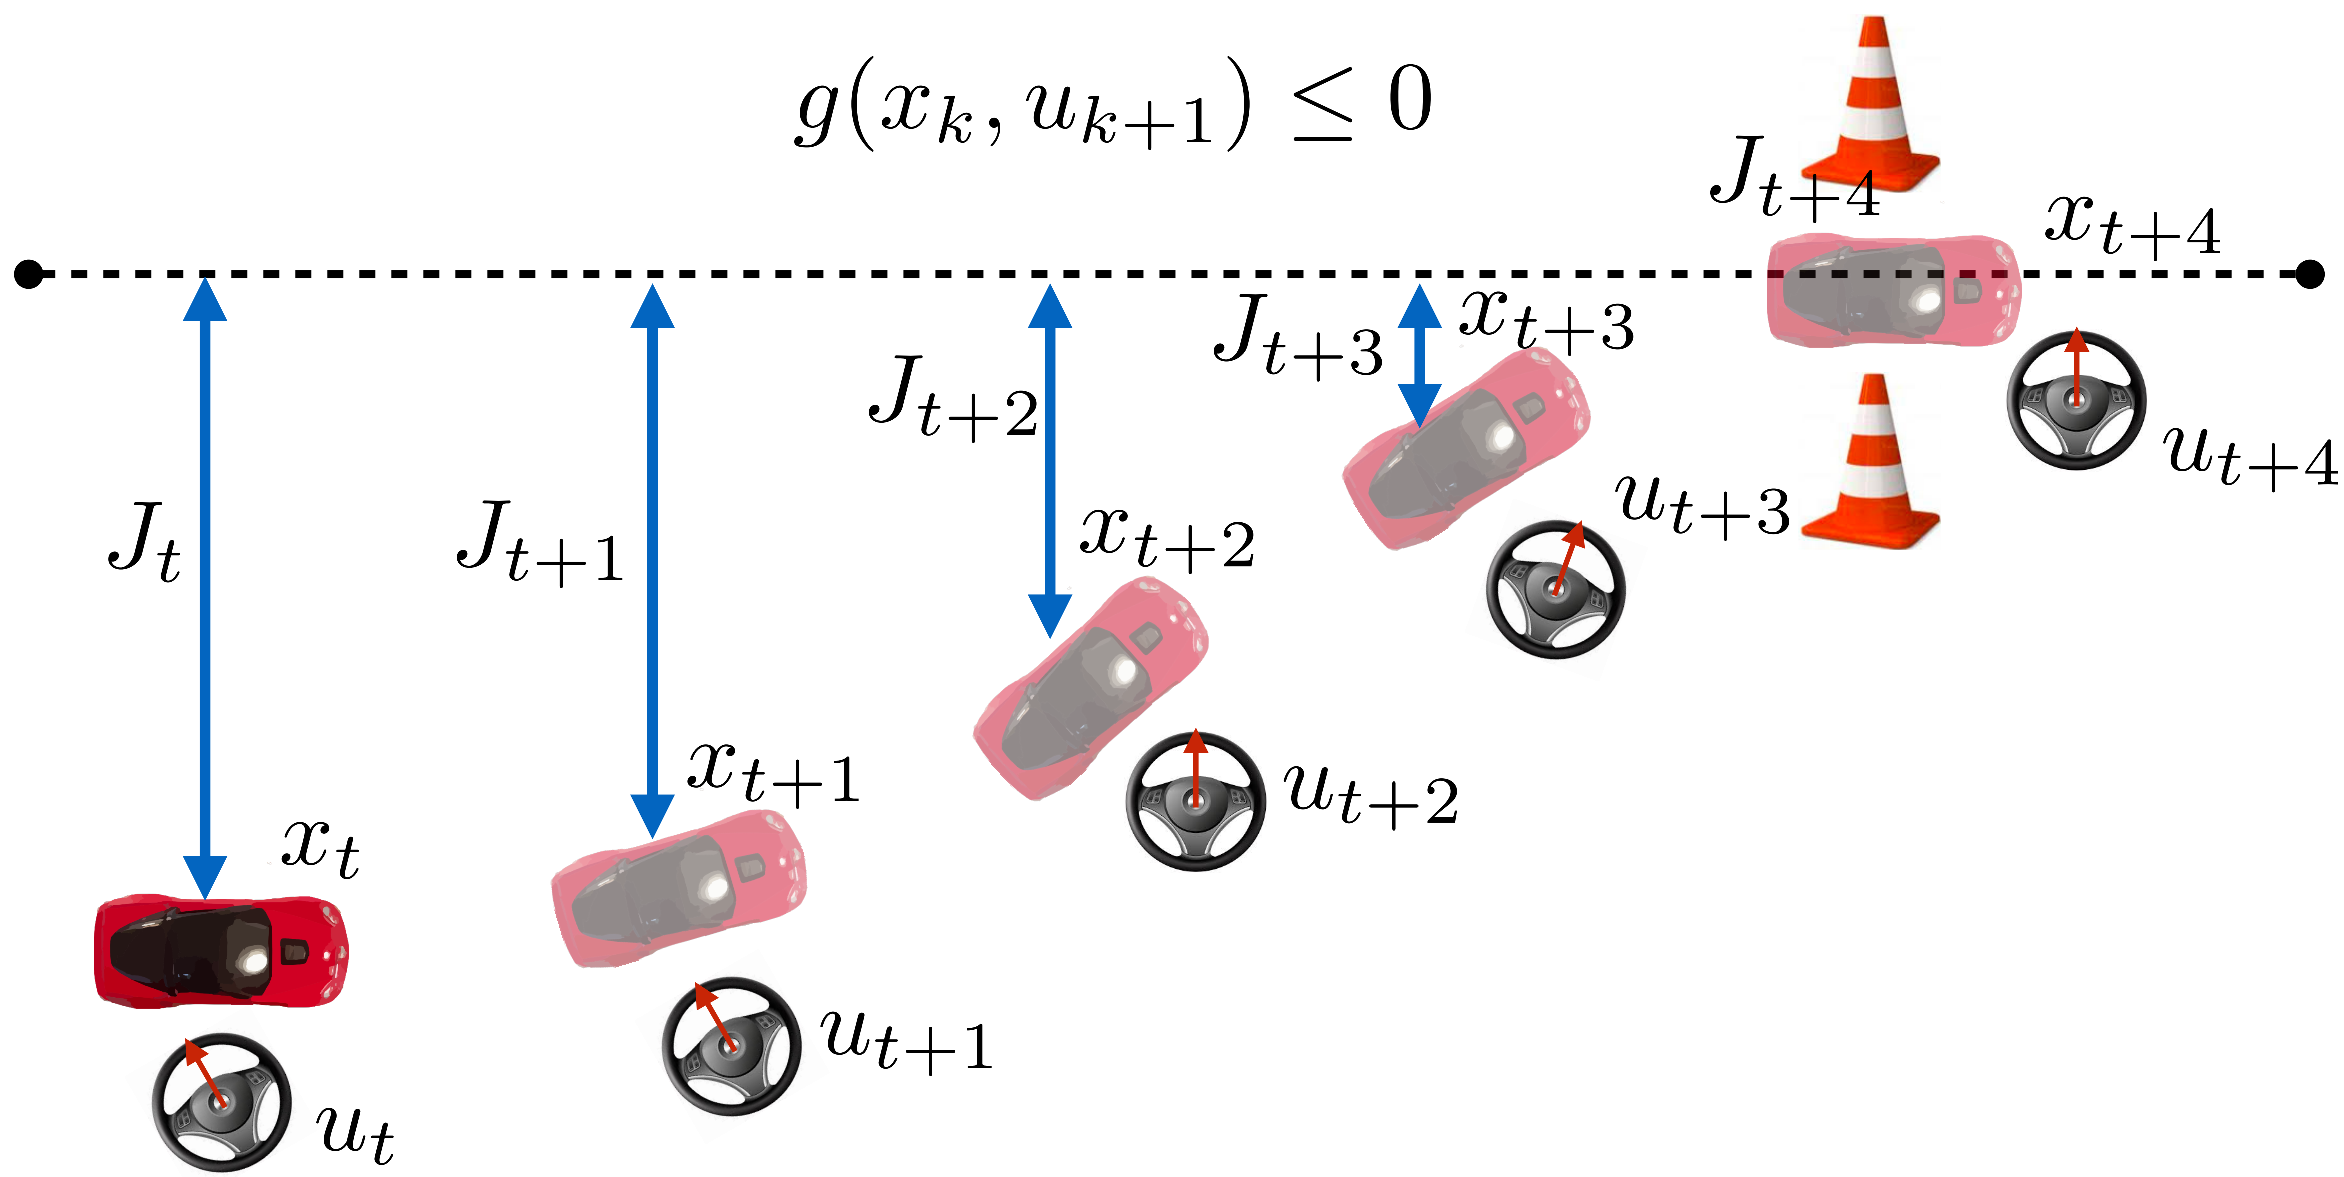

How are the controls executed?

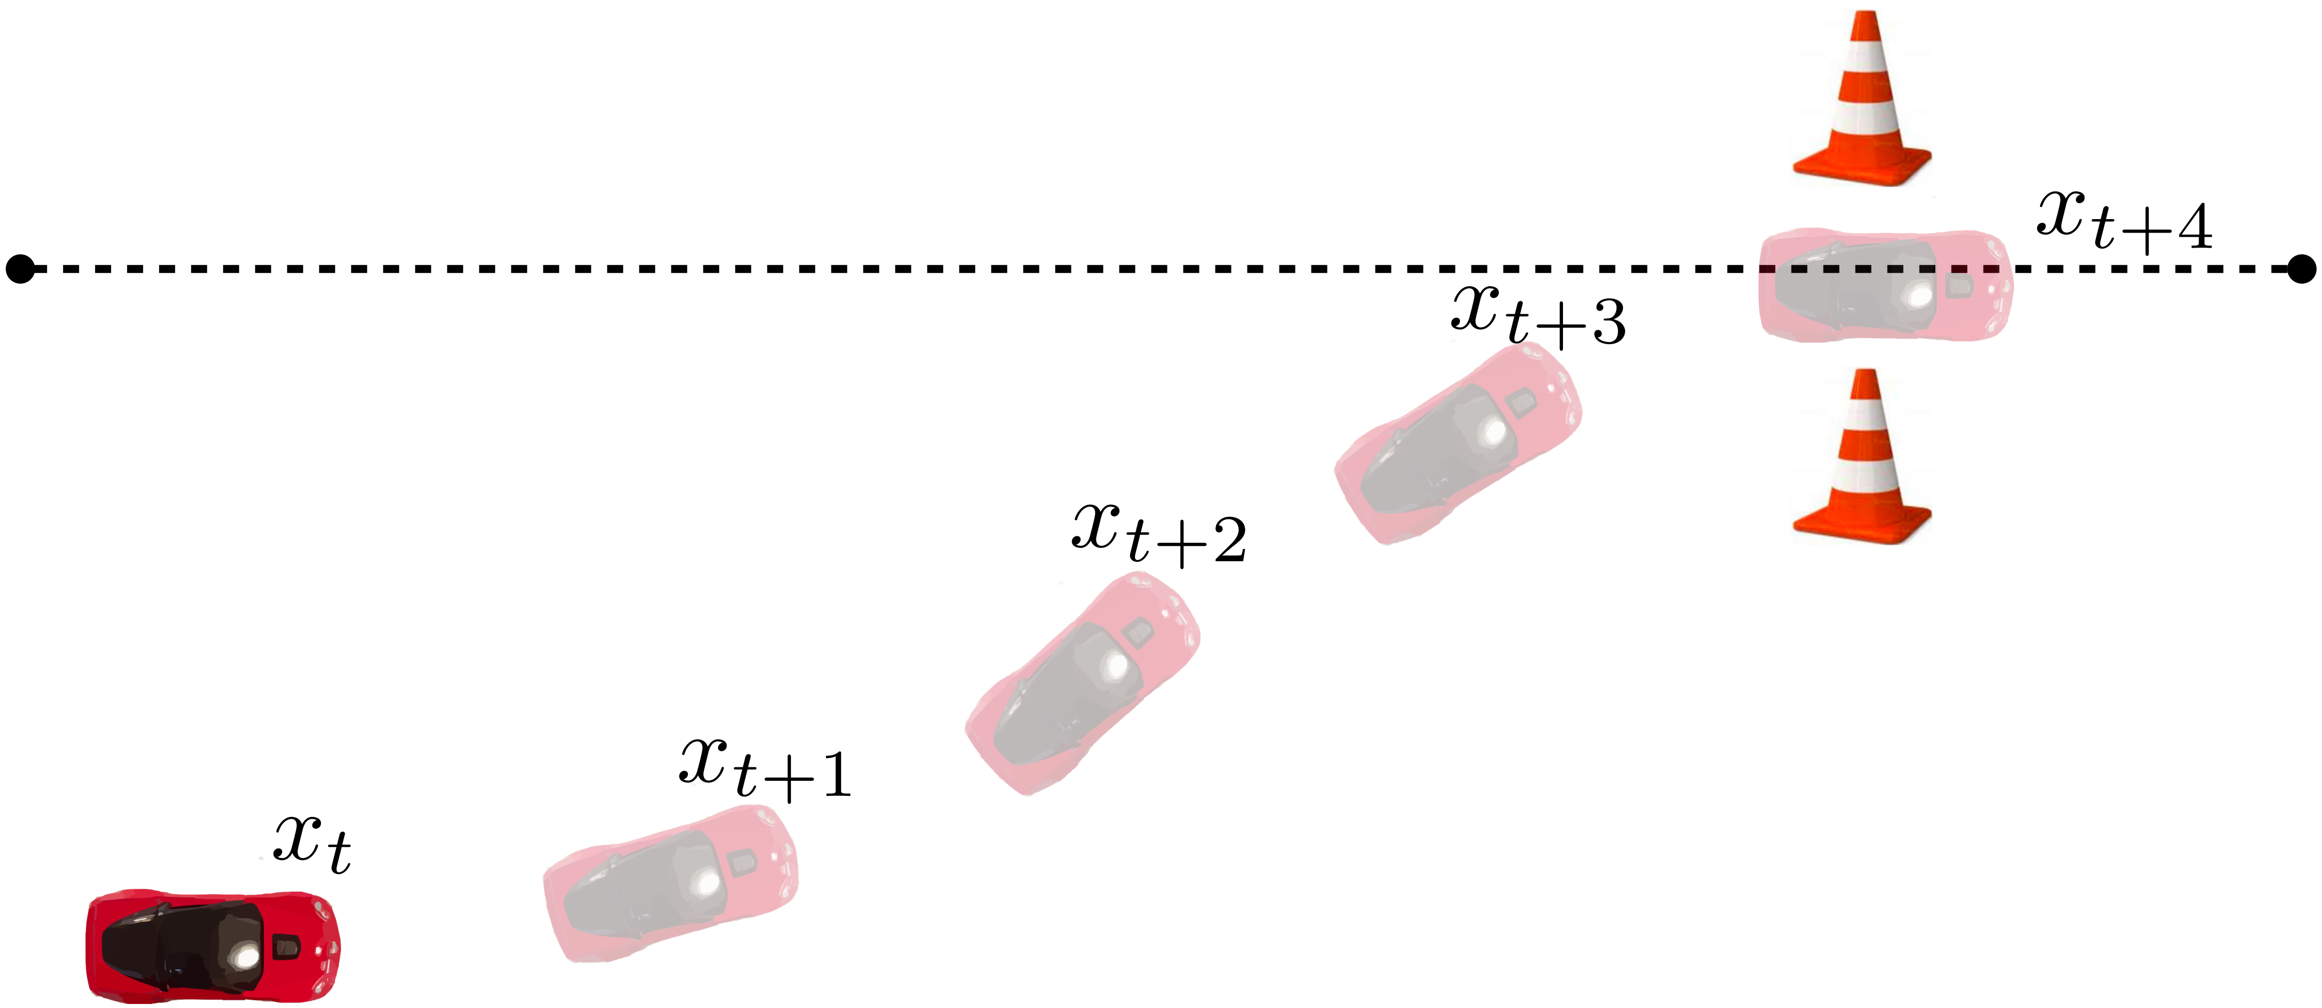

Step 1: Solve optimization problem to a horizon

How are the controls executed?

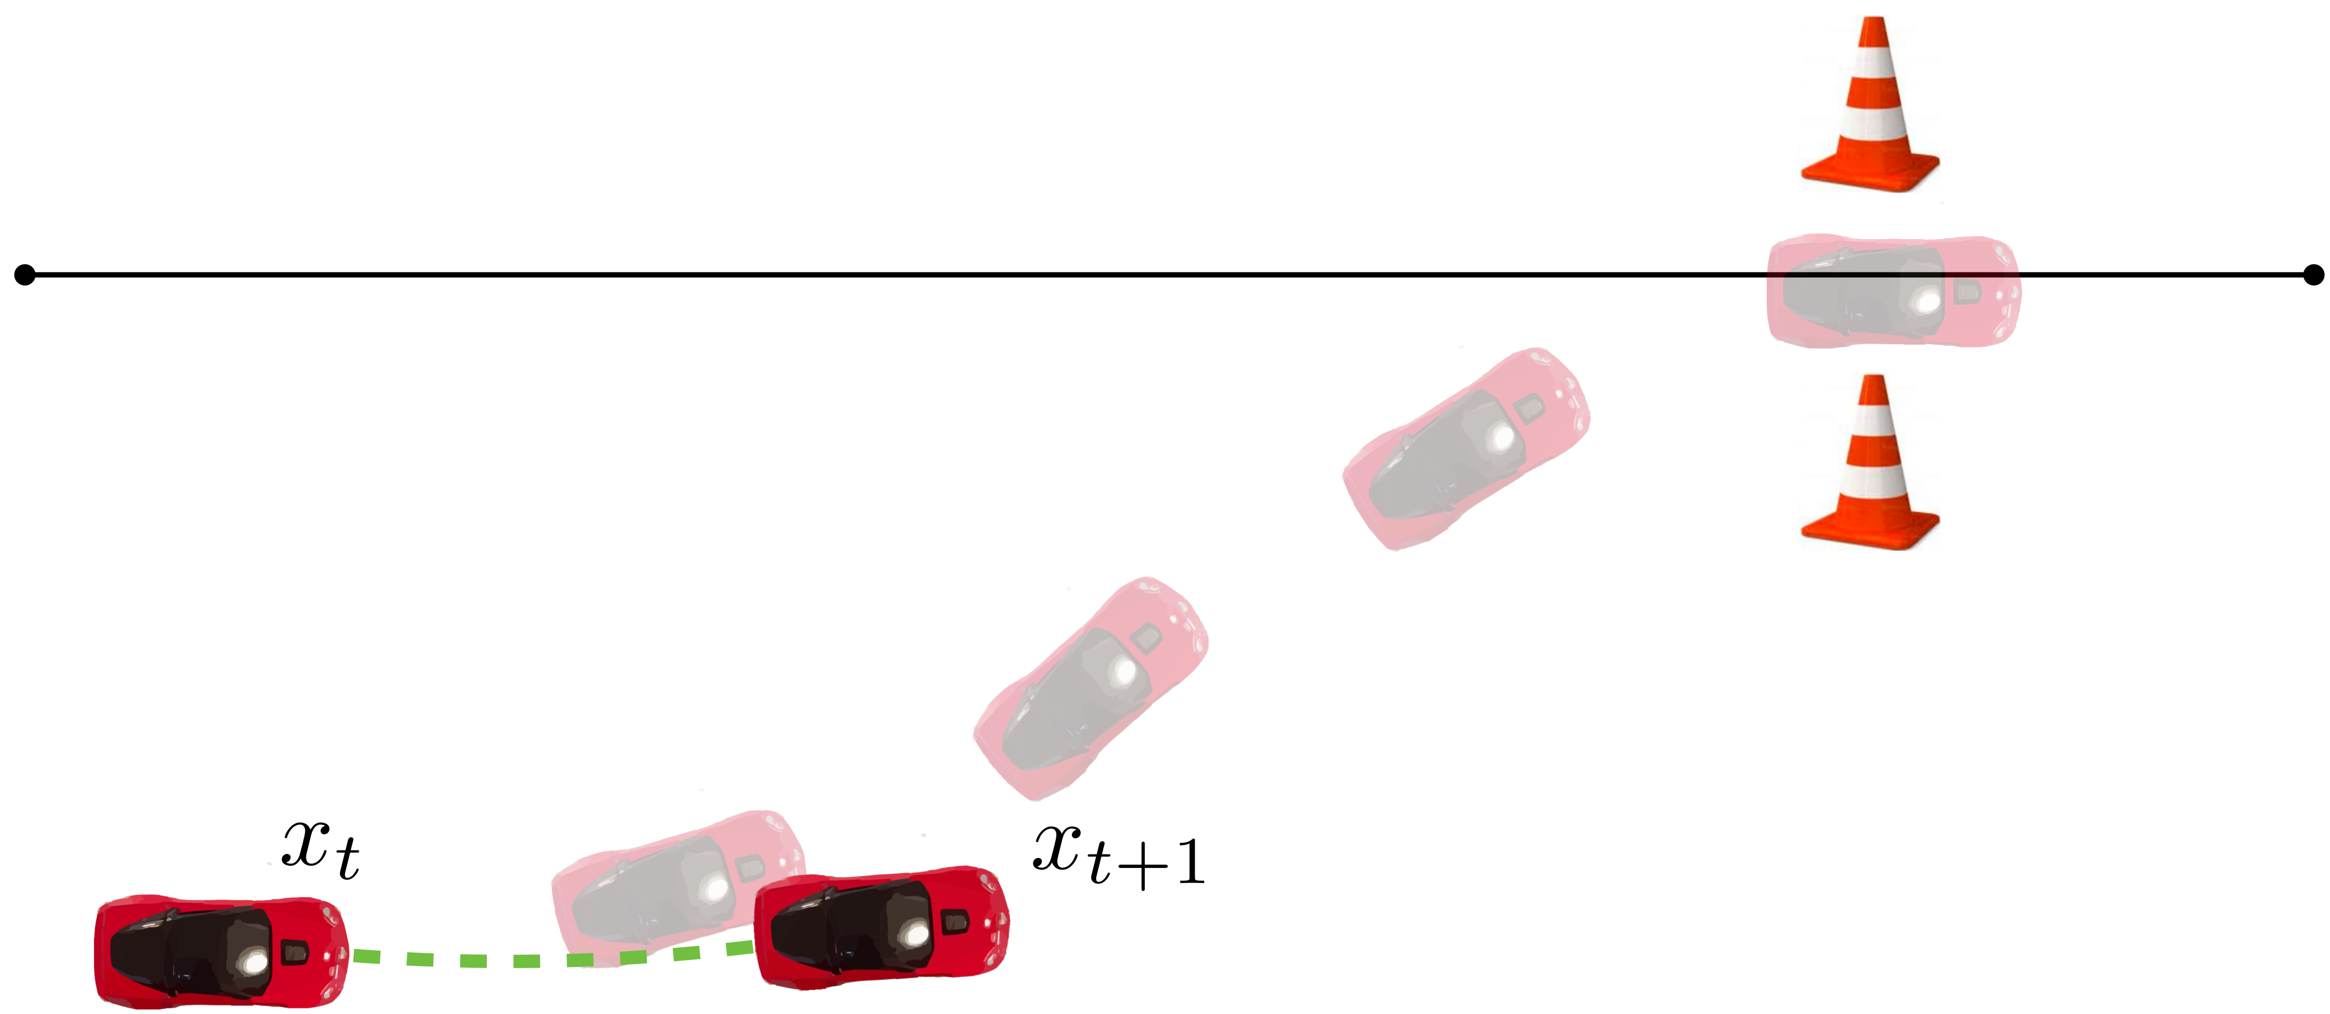

Step 1: Solve optimization problem to a horizon

Step 2: Execute the first control

How are the controls executed?

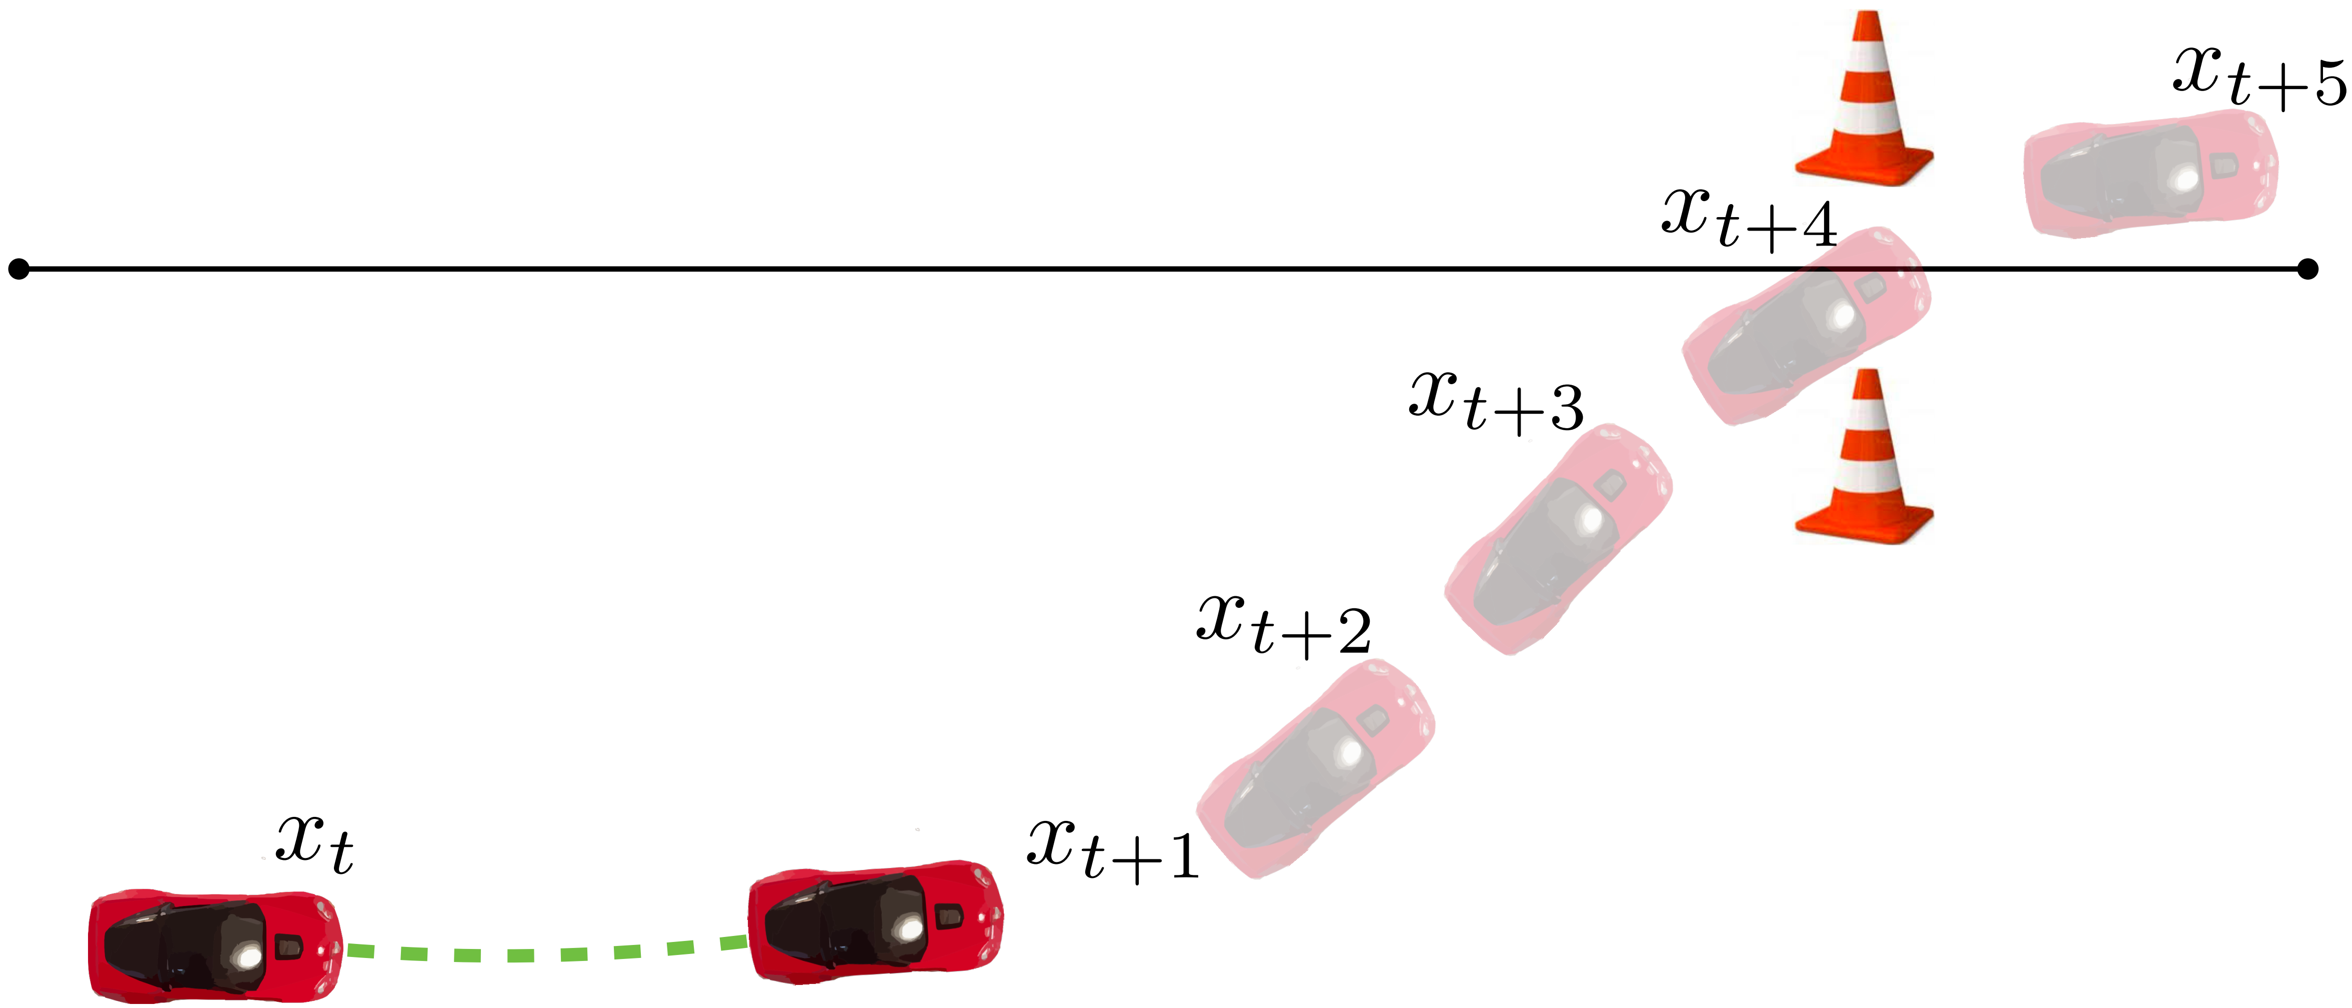

Step 1: Solve optimization problem to a horizon

Step 2: Execute the first control

Step 3: Repeat!

What is the cost? What constraints?

$$\min_{u_{t+1},\dots u_{t+H}} \sum_{k=t}^{t+H-1} \boxed{\text{Cost}}$$

(plan till horizon H)

$$x_{k+1} = f(x_k, u_k)$$

$$\boxed{\text{Constraints}}$$

$$g(x_k, u_{k+1}) \leq 0$$

What is the cost? What constraints?

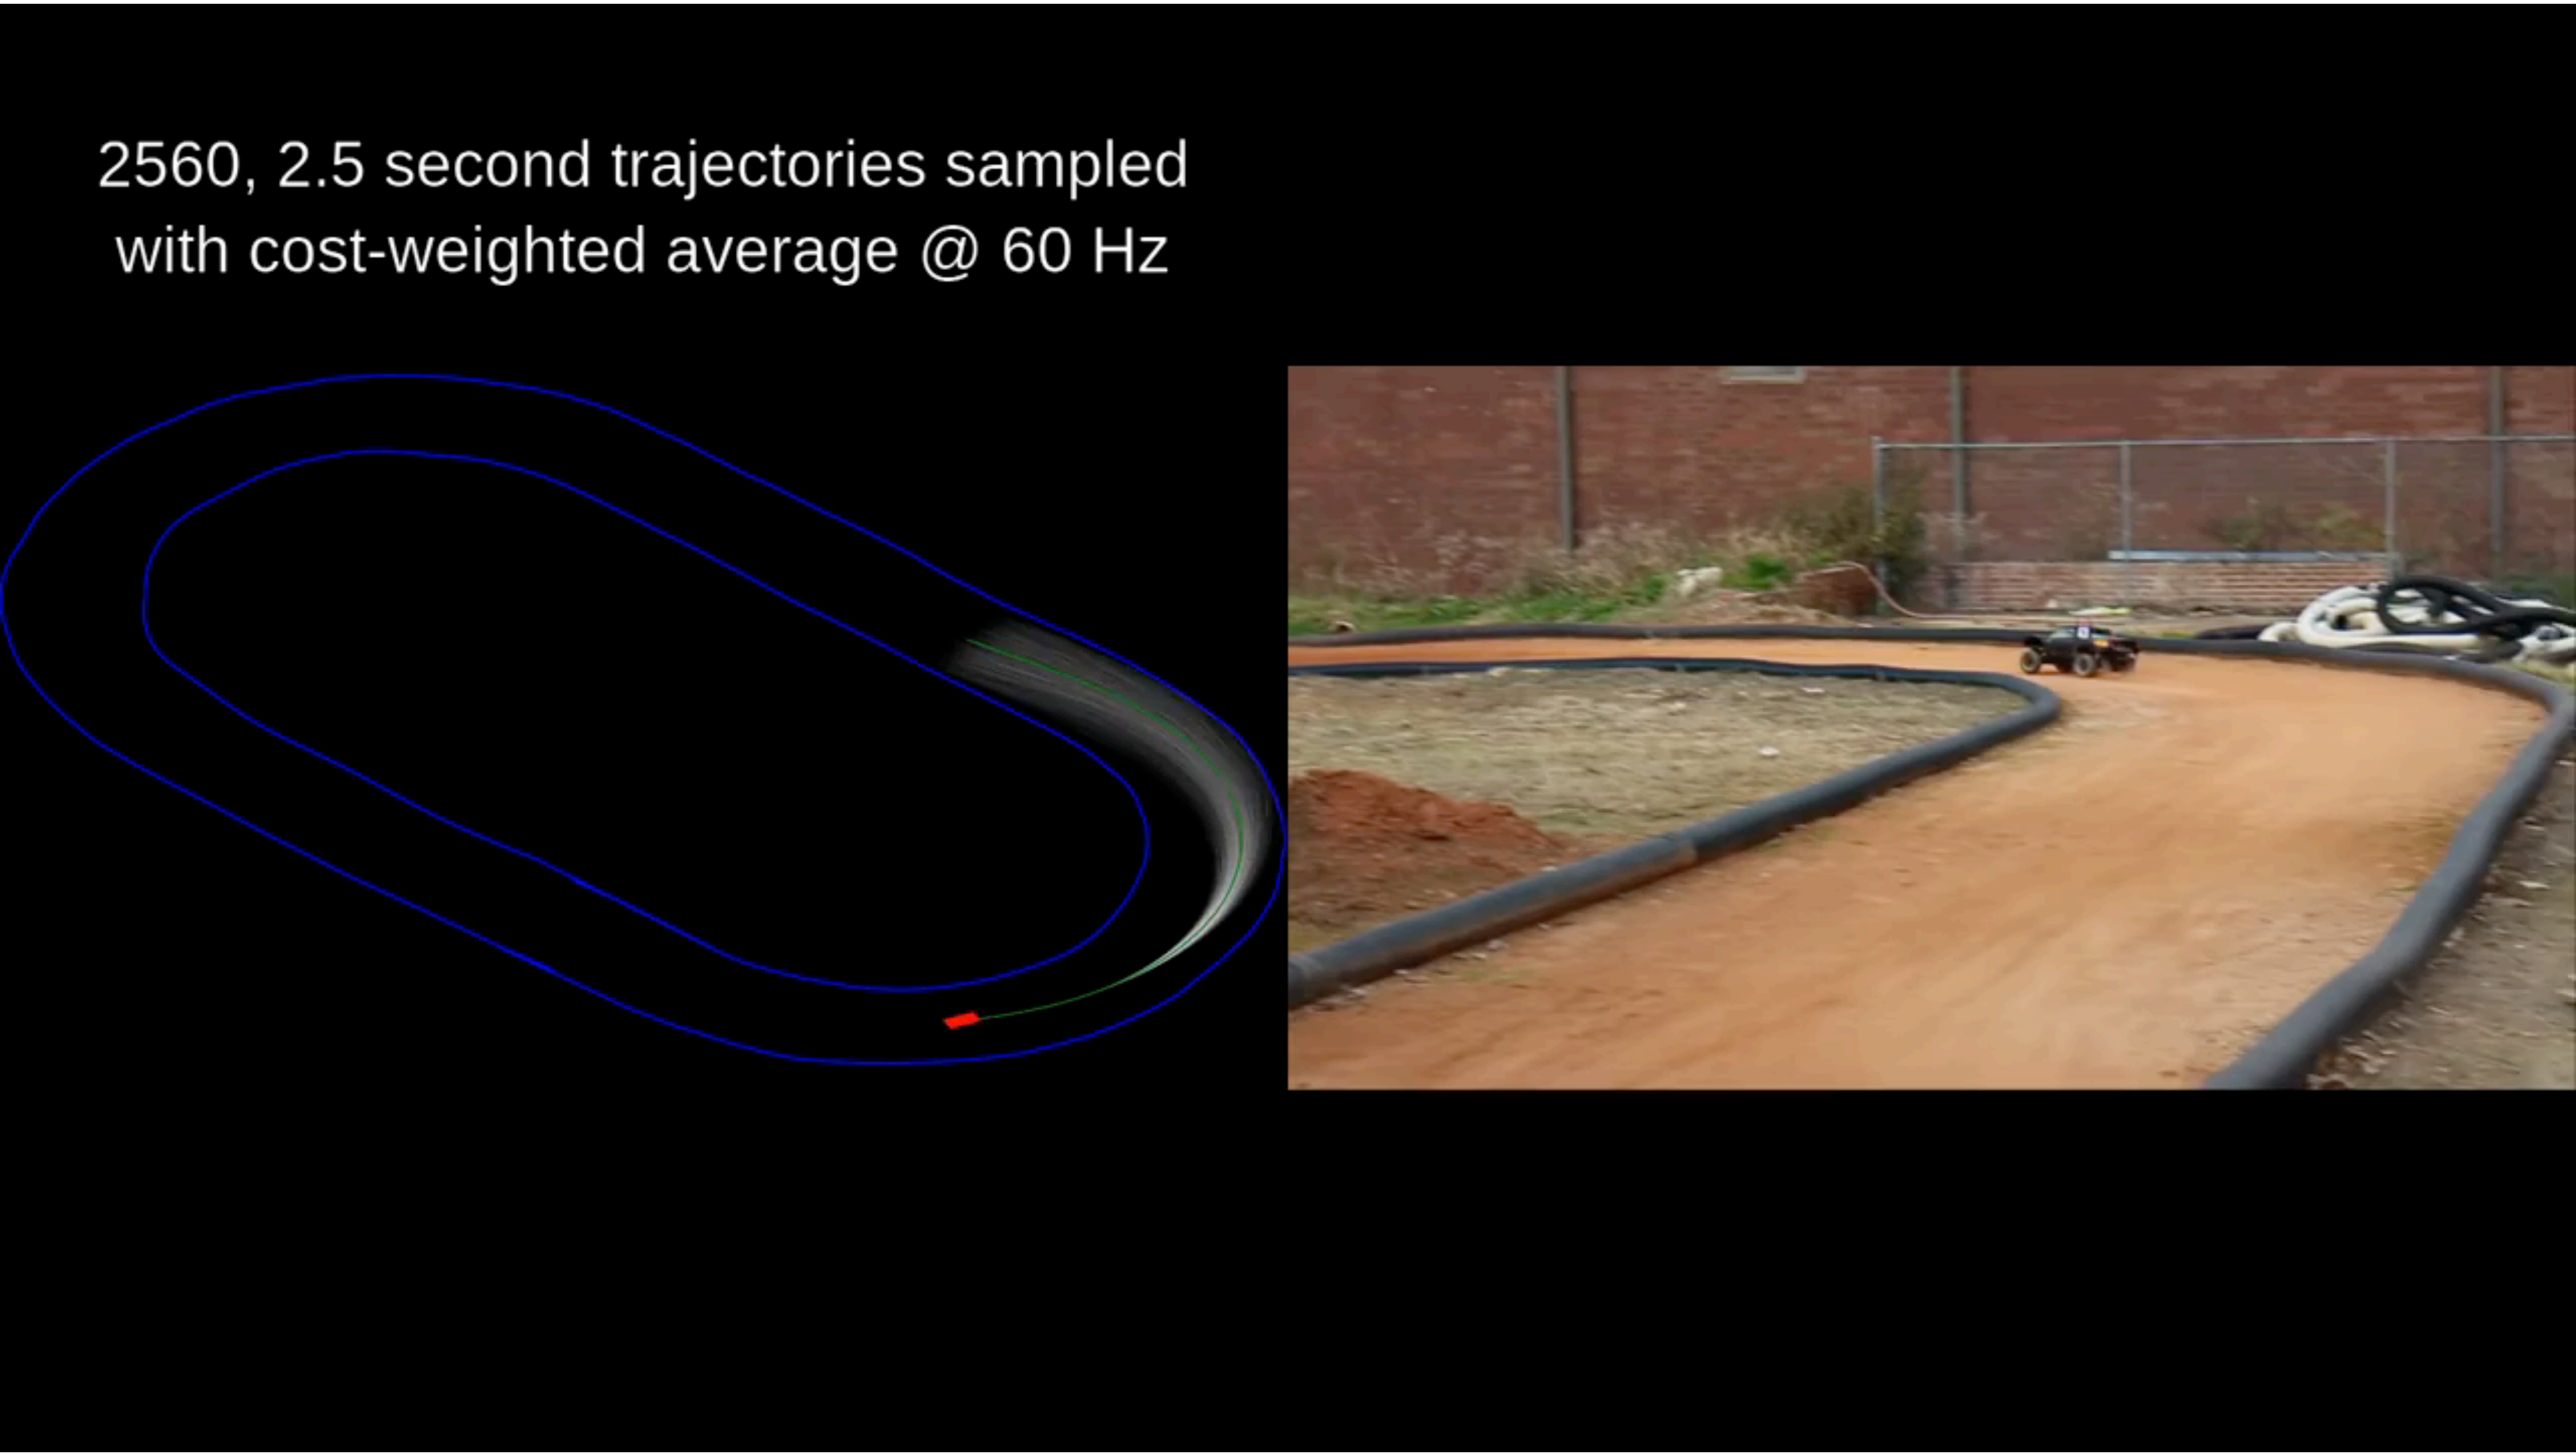

Examples of complex cost functions

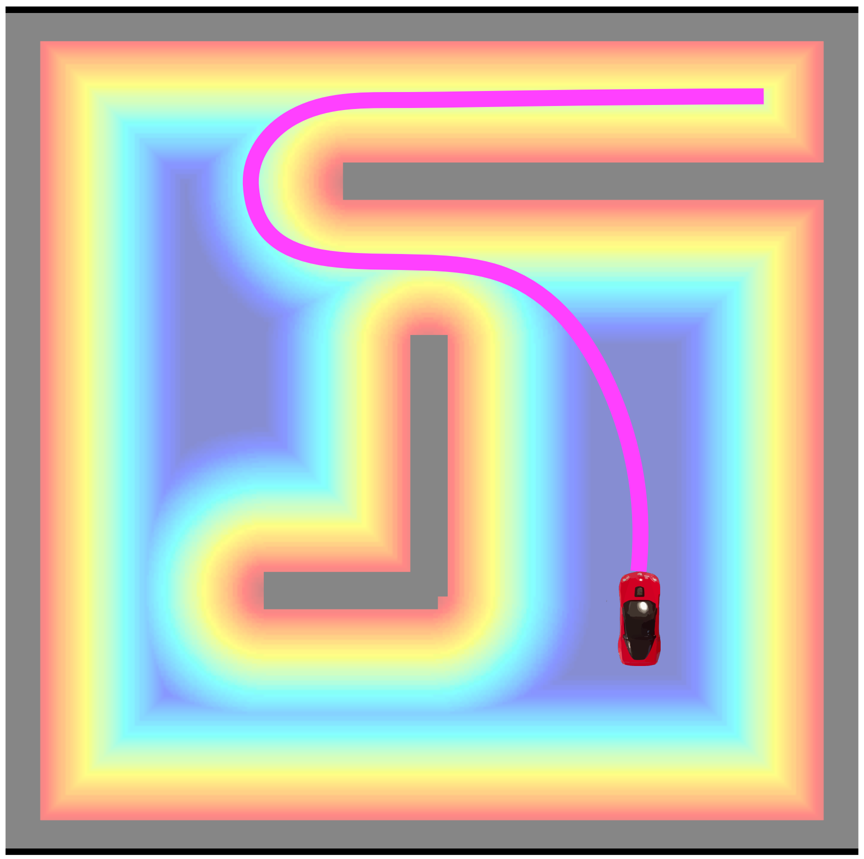

Proximity to obstacles

(Ratliff et al. 2009)

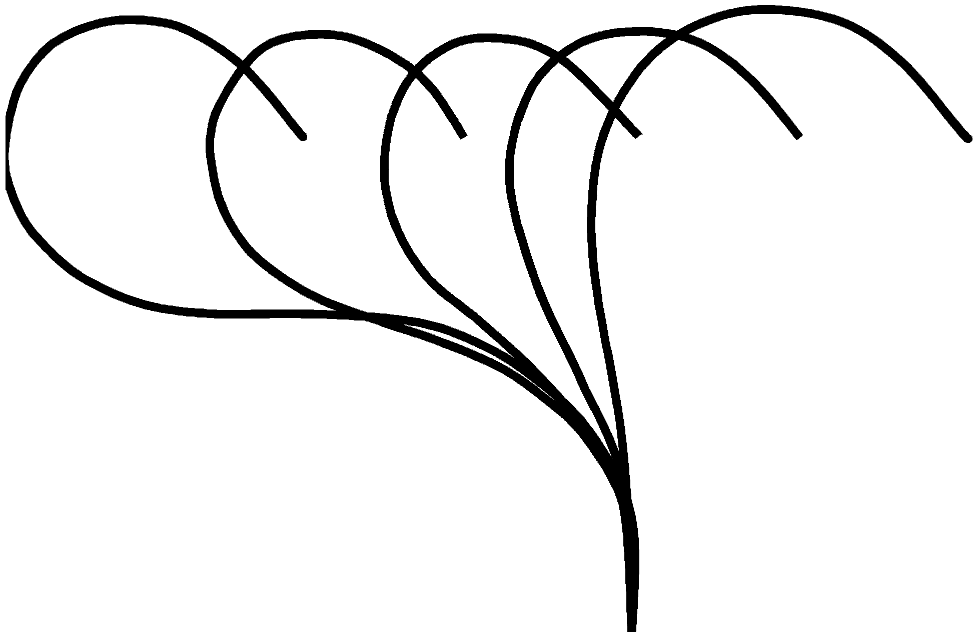

Curvature

(Kelly et al. 2003)

Problem: Costs can be non-convex (terrible local minima)

Instead: Create a gigantic library of paths

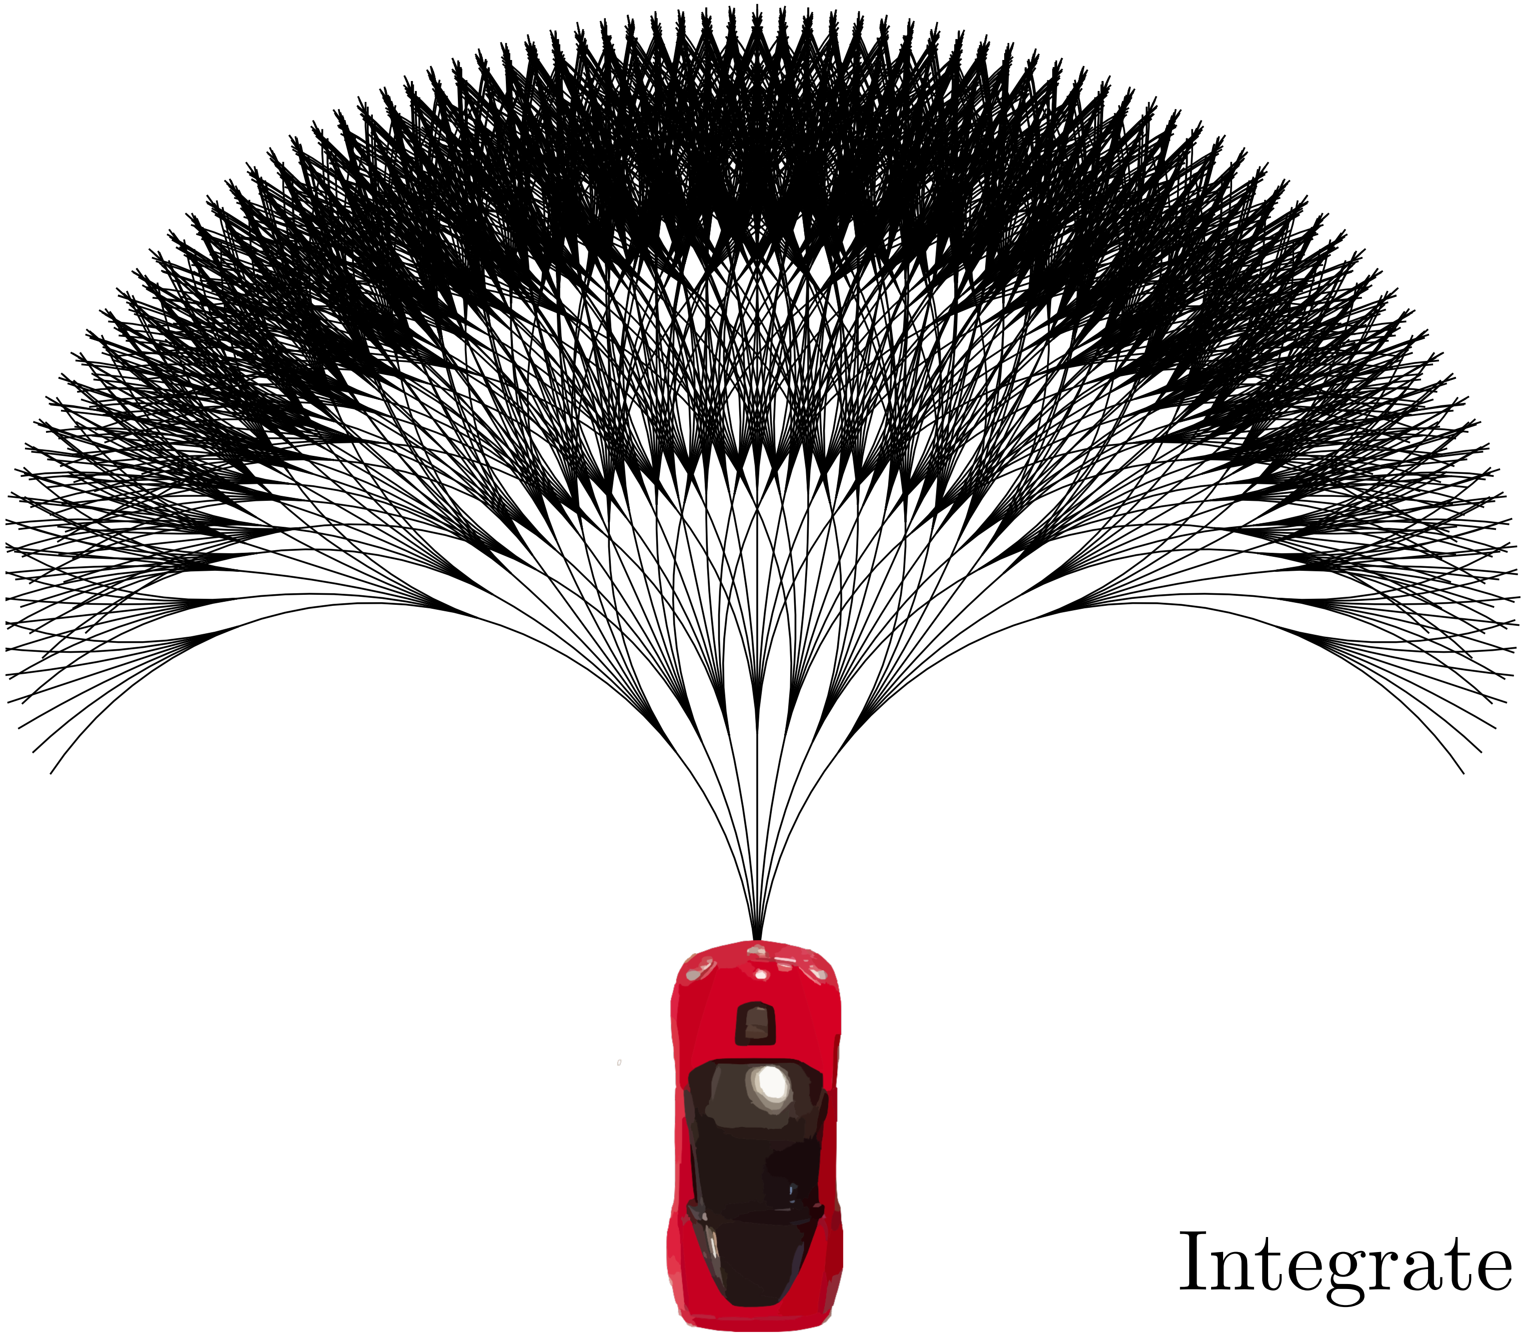

Discretize steering angle into d bins (set d=7)

Branch every t sec for N*t horizon (N=5)

Get 7^5 = 16,807 control trajectories

$$u(t)$$

Integrate dynamics to get $x(t)$

How will we use it for MPC?

1. Iterate (2-3) for every path in library.
2. Compute constraint. If violated, chuck out path
3. Compute cost.
4. Pick path with least cost.
5. Execute first control action. Robot moves. Replan.

Problem: Library too big!!

Solution: Subsample library

But what is the right sampling strategy?

Strategy 1: Uniformly subsample

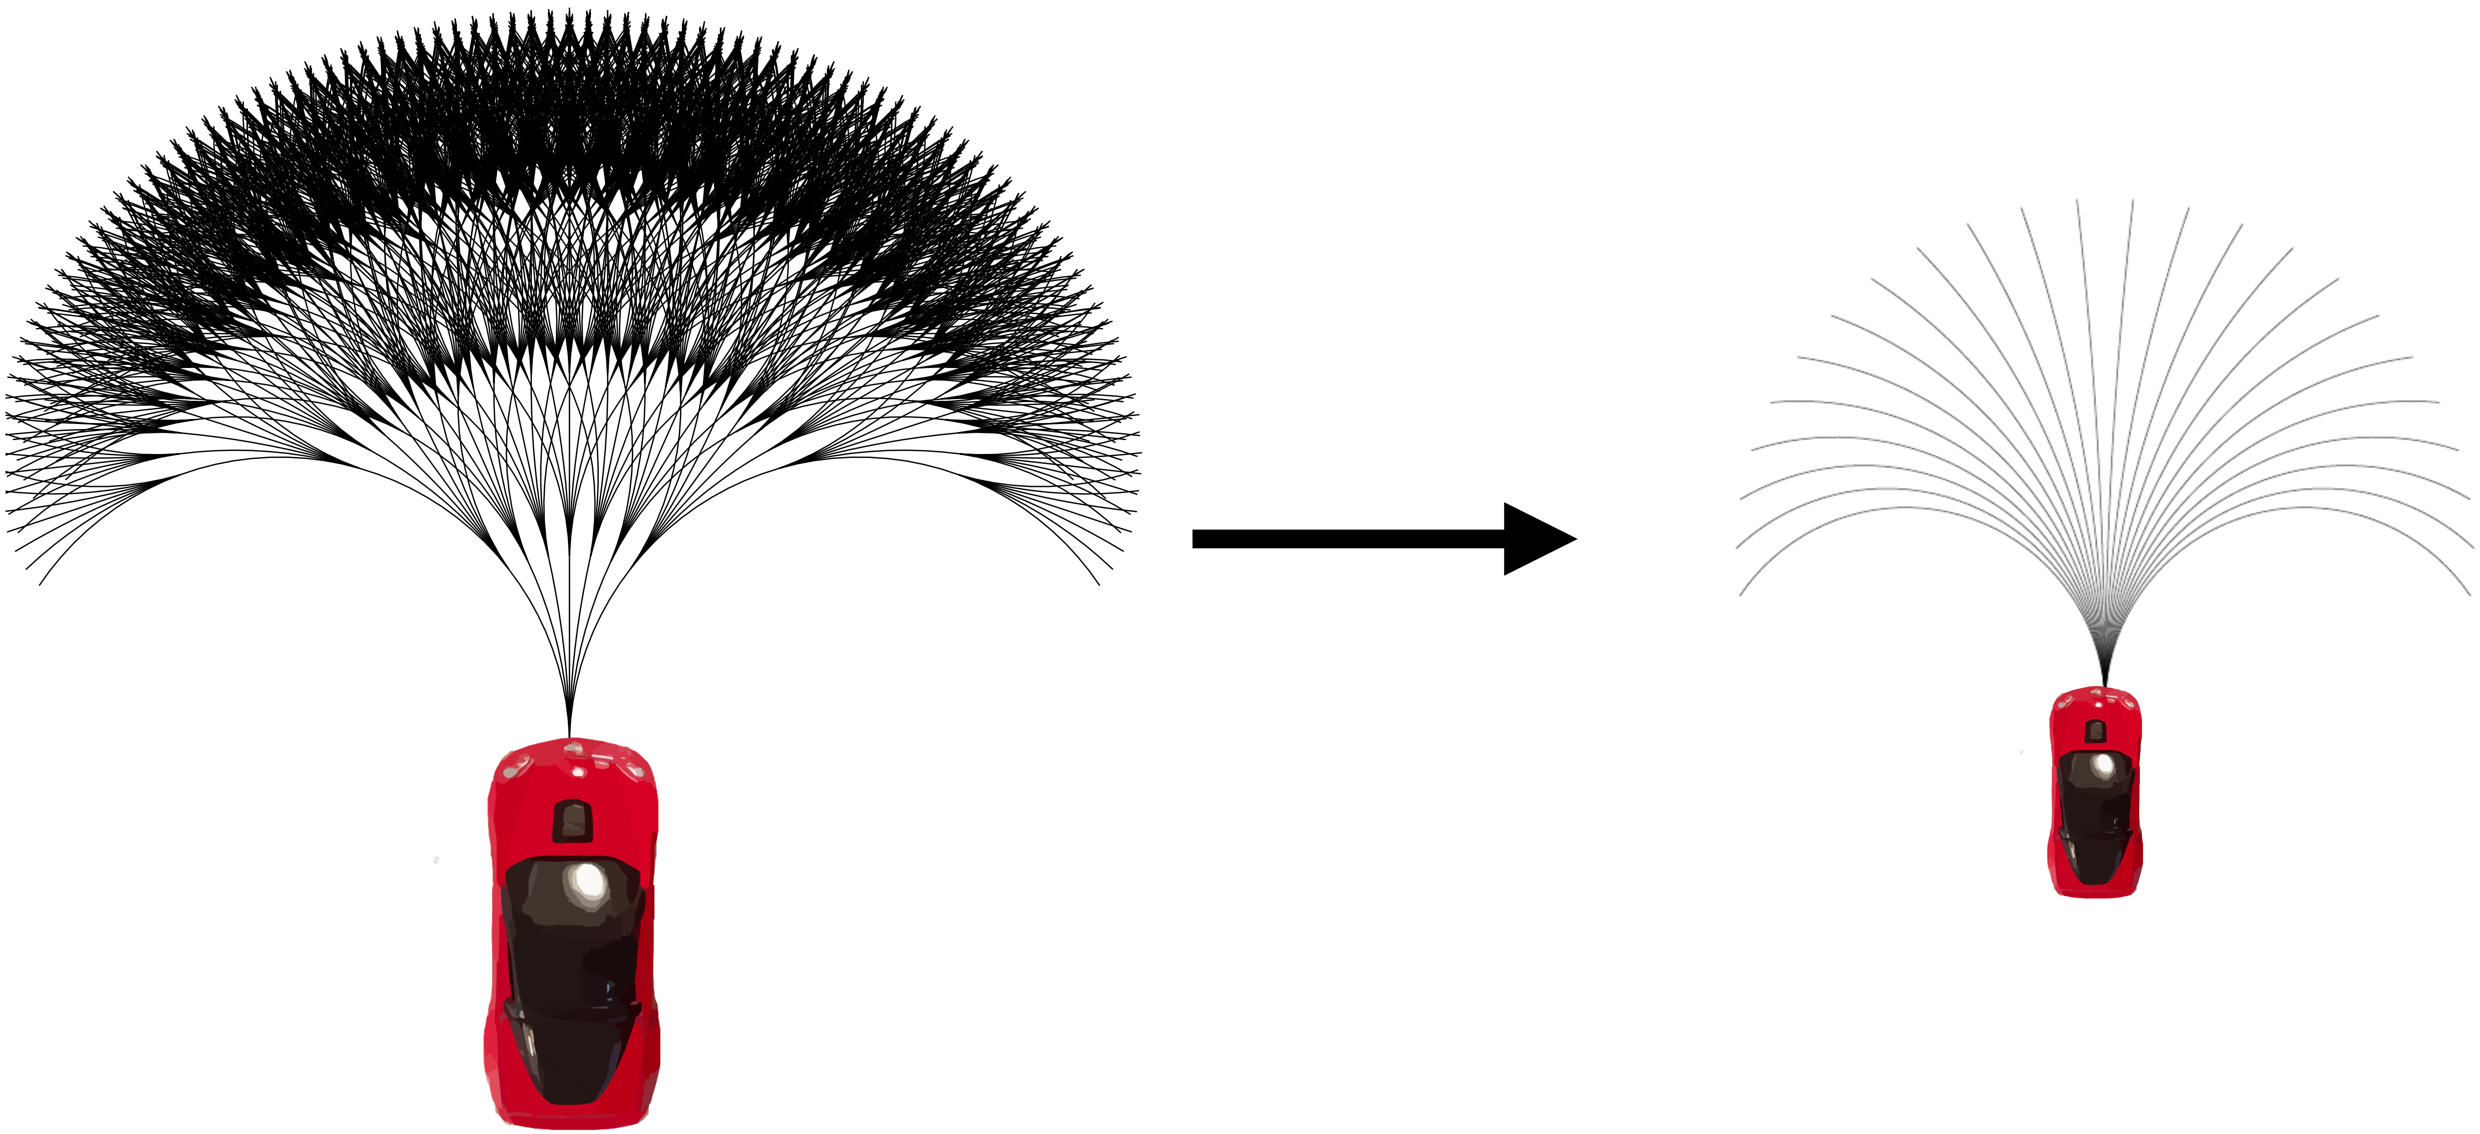

Why is this not a good library?

Strategy 2: Get more diversity

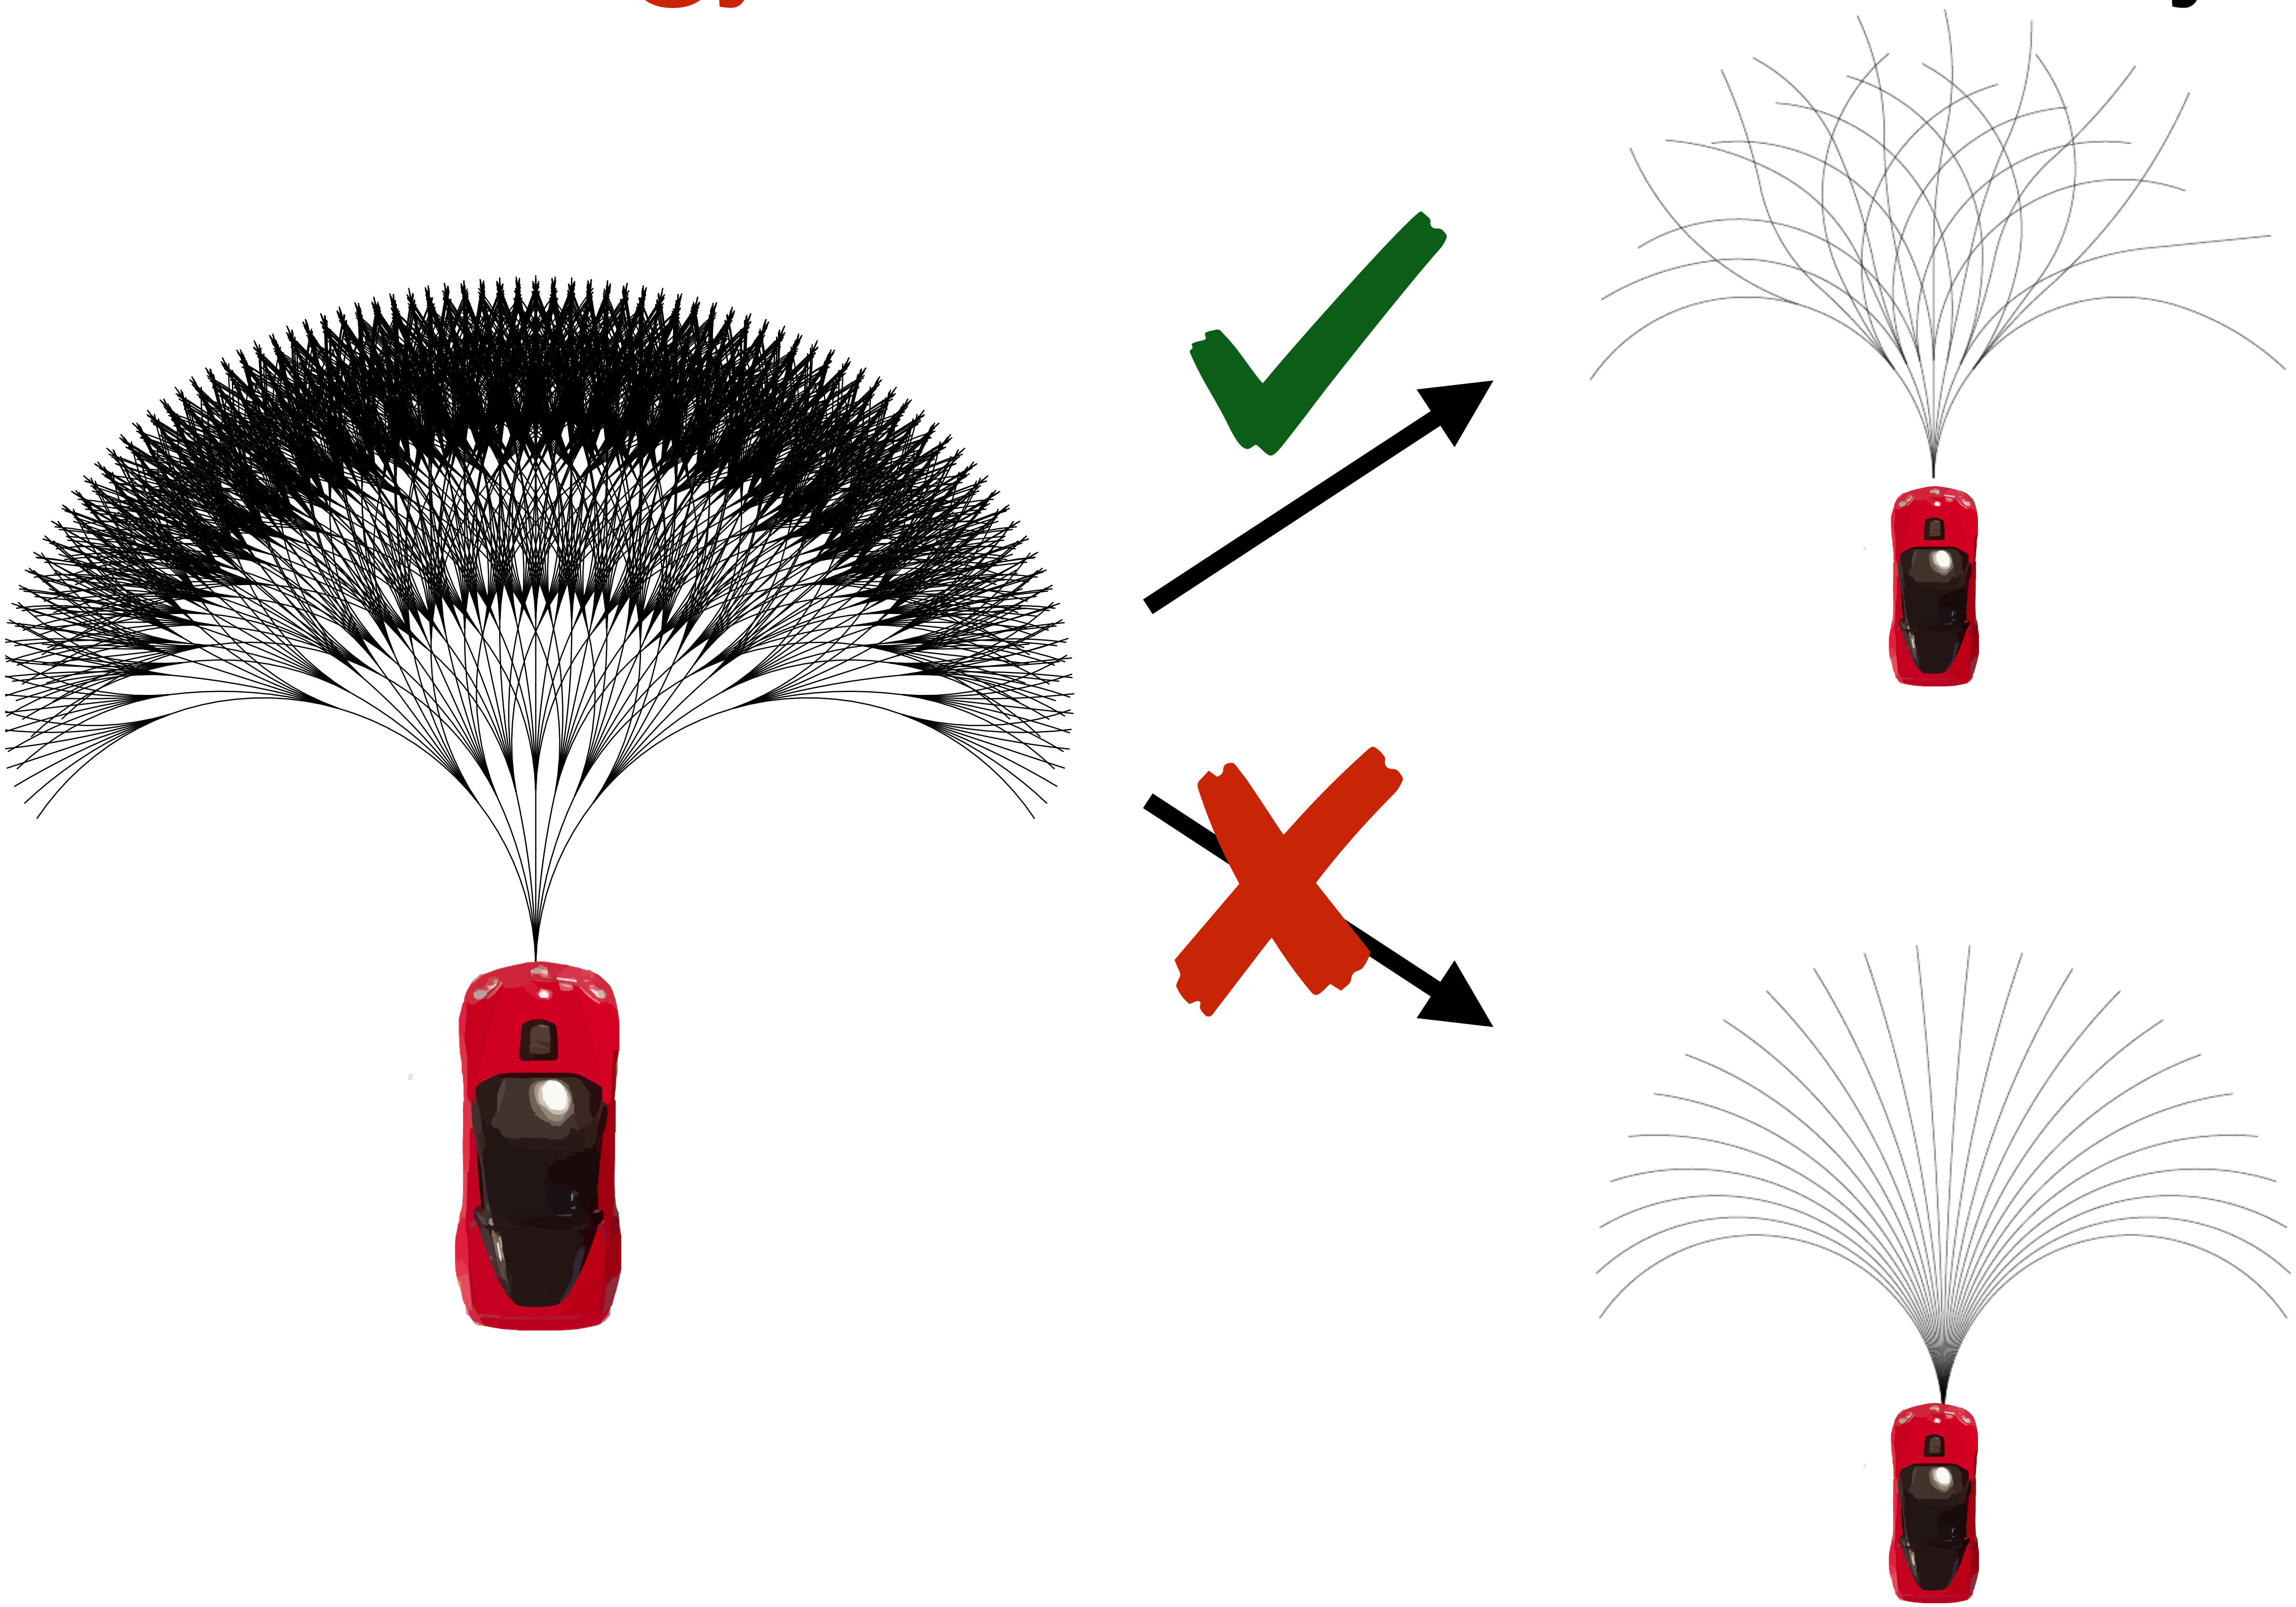

Proposed problem formulation

Given a dense path set D

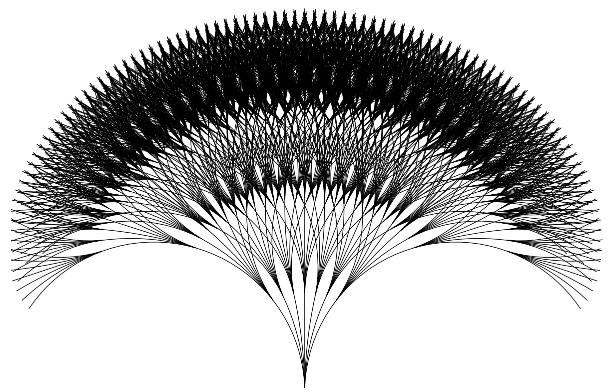

Find a sparse subset $S \subset D$

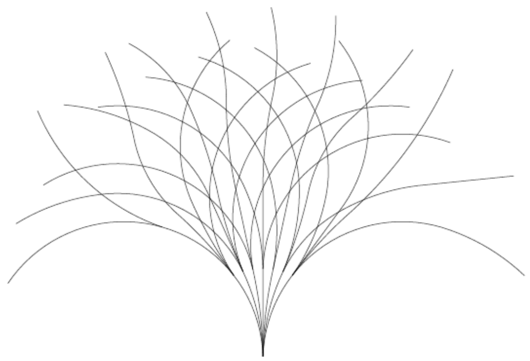

$$\arg\max_{S \subset D} \textsc{Coverage}(S, D)$$

What is coverage? Why do we want coverage?

You want to cover the space well so that you maximize the likelihood of finding a collision-free path.

Why do we want coverage?

Let's say the dense set has an optimal path $\xi^D = \arg\min_{\xi \in D} J(\xi)$

We want to make sure the sparse set has a path ξ^S "close to" ξ^D

such that $J(\xi^S) \leq J(\xi^D) + \epsilon$

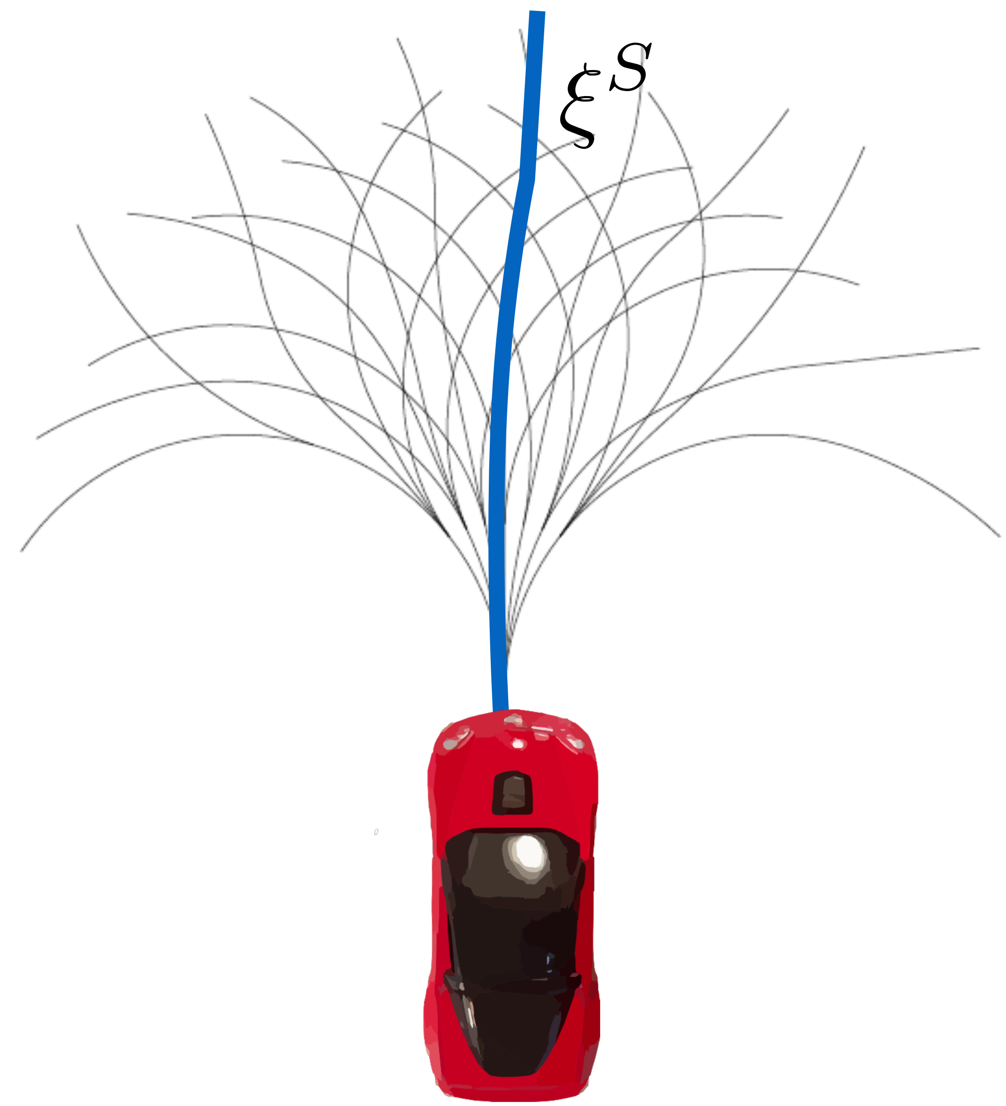

Lipschitz continuity: Smoothness assumption

A Lipschitz continuous function satisfies the following inequality

$$\forall x, y, \; \exists L \geq 0 \text{ s.t. } \quad |f(x) - f(y)| \leq L||x - y||$$

(Close points have the same function value)

The cost functions we consider are Lipschitz continuous

$$|J(\xi_1) - J(\xi_2)| \leq L d(\xi_1, \xi_2)$$

(Close trajectories have the same cost value)

Defining a "closeness" between paths

A metric is a generalization of distance

$$d(\xi_1, \xi_2)$$

Path 1 Path2

We will choose a metric in the space of paths

1. Non-negative
2. Triangle inequality
3. Identity
4. Symmetry

Conditions that a metric satisfies

There are plenty of metrics to choose from!

Defining a “closeness” between paths

A set with a metric.

Project paths on to a metric space where each dot is a path

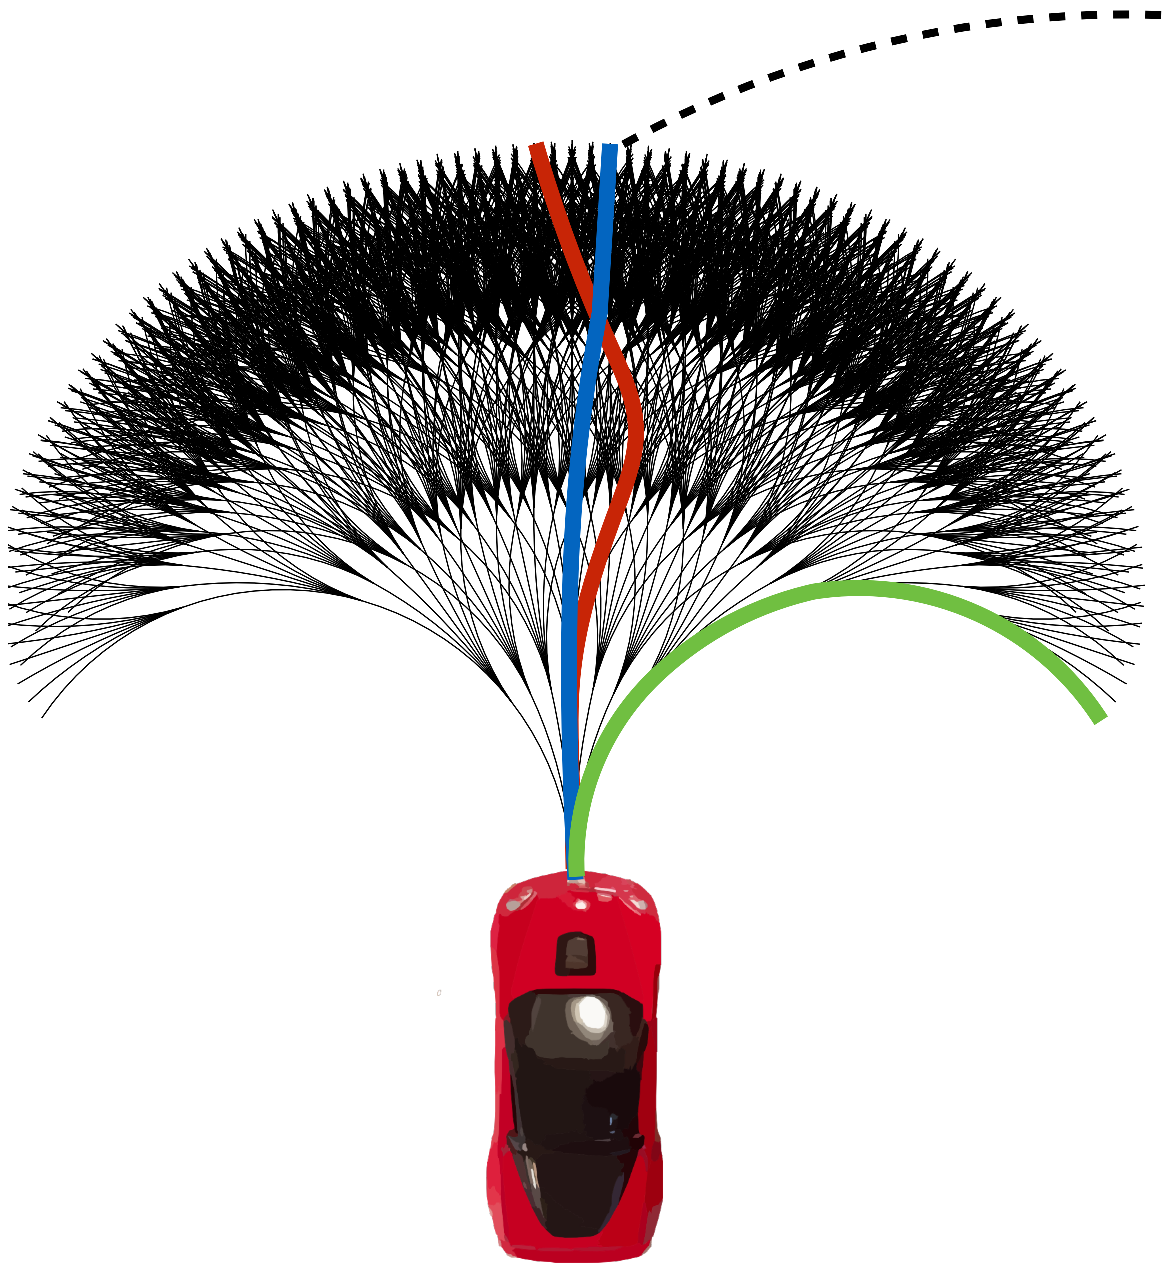

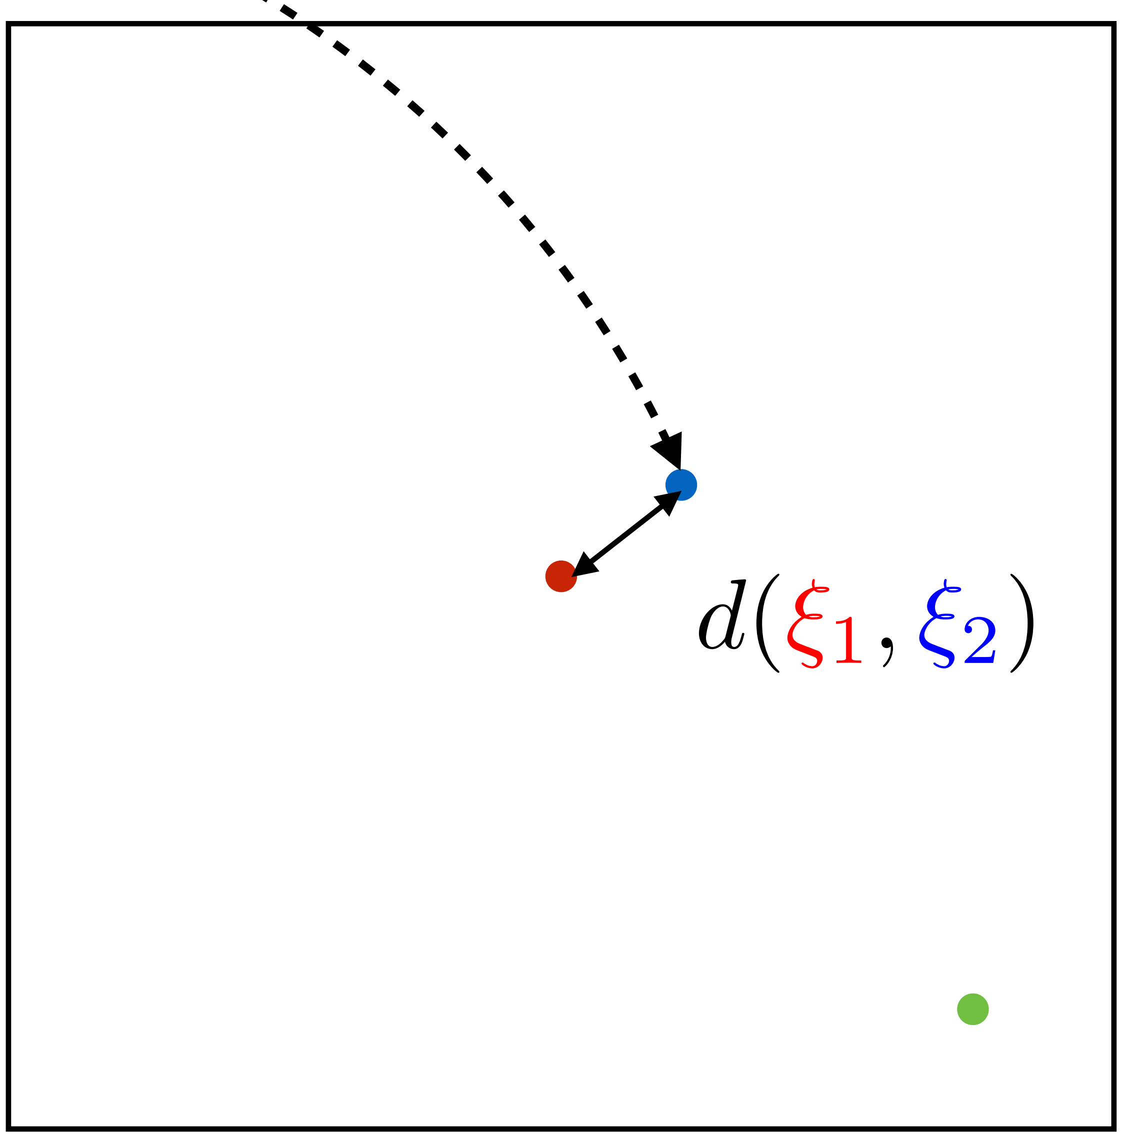

Defining a "closeness" between paths

Every dot is a path - neighboring dots are close

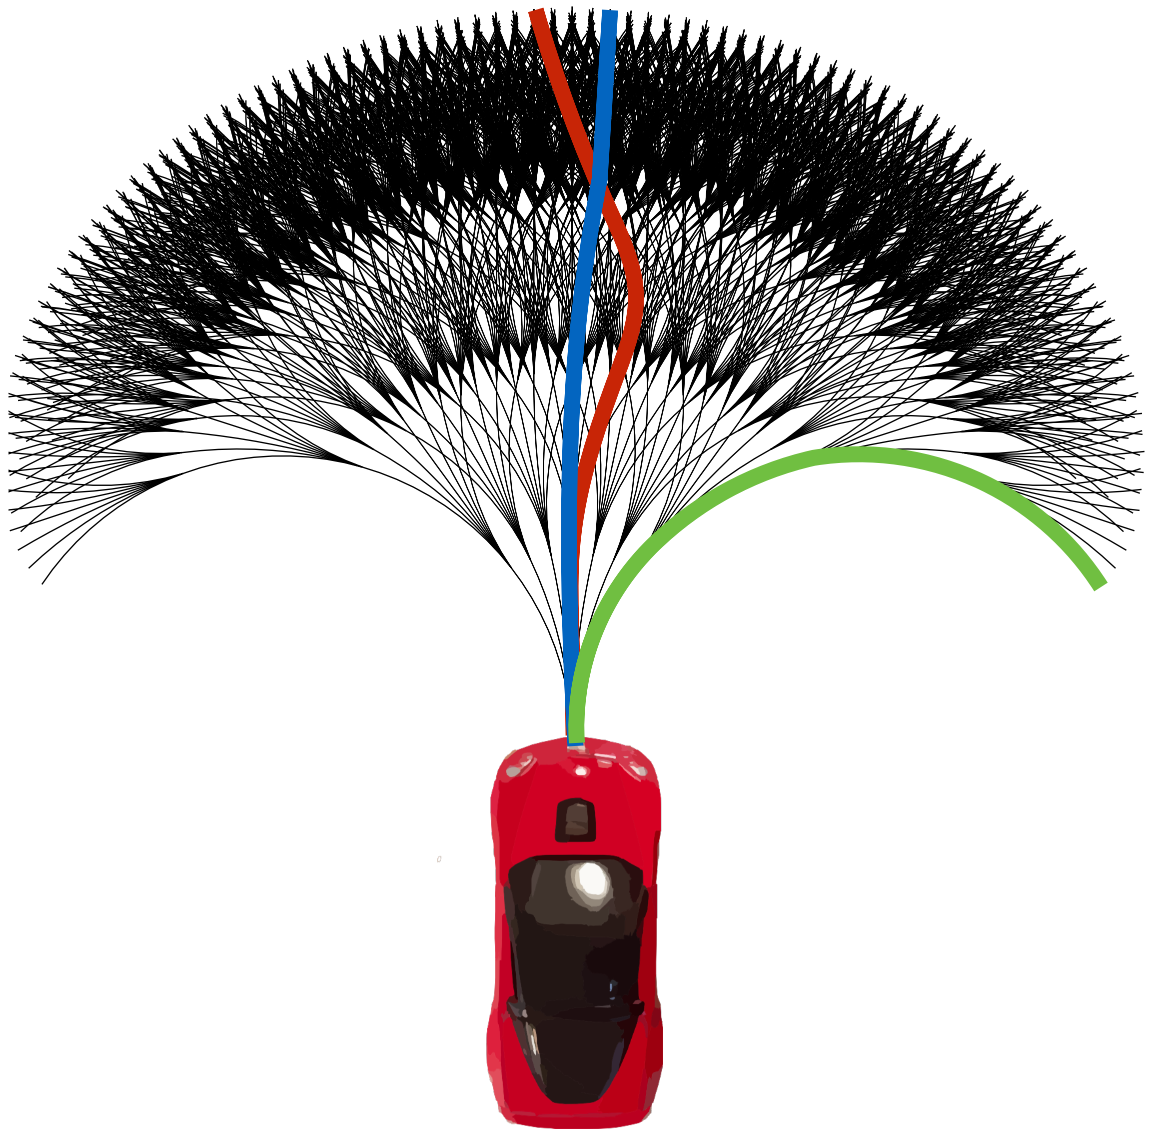

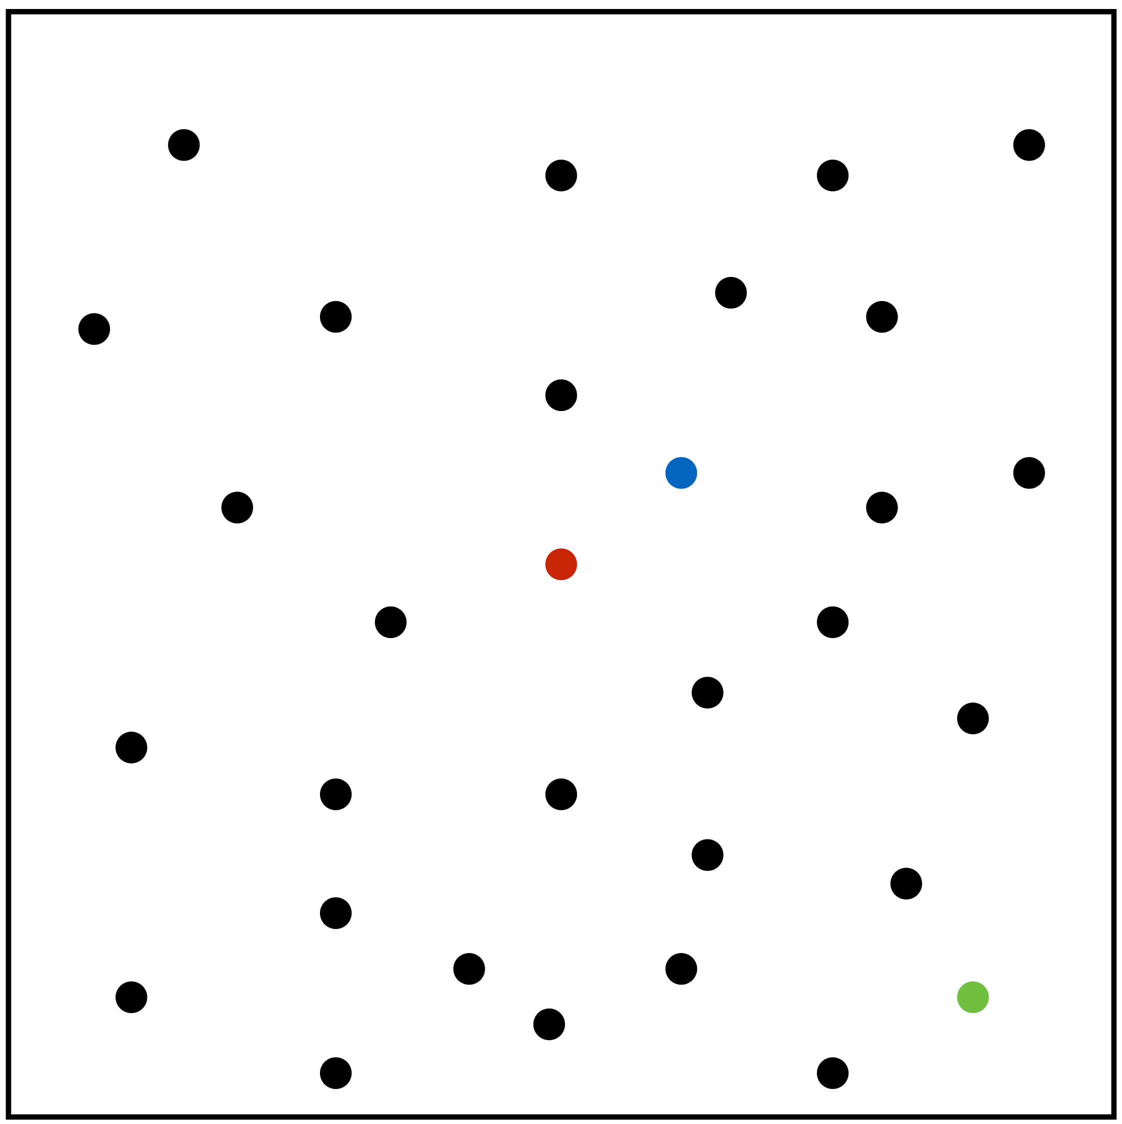

Objective: "Cover" as many black dots

Which library does a better job?

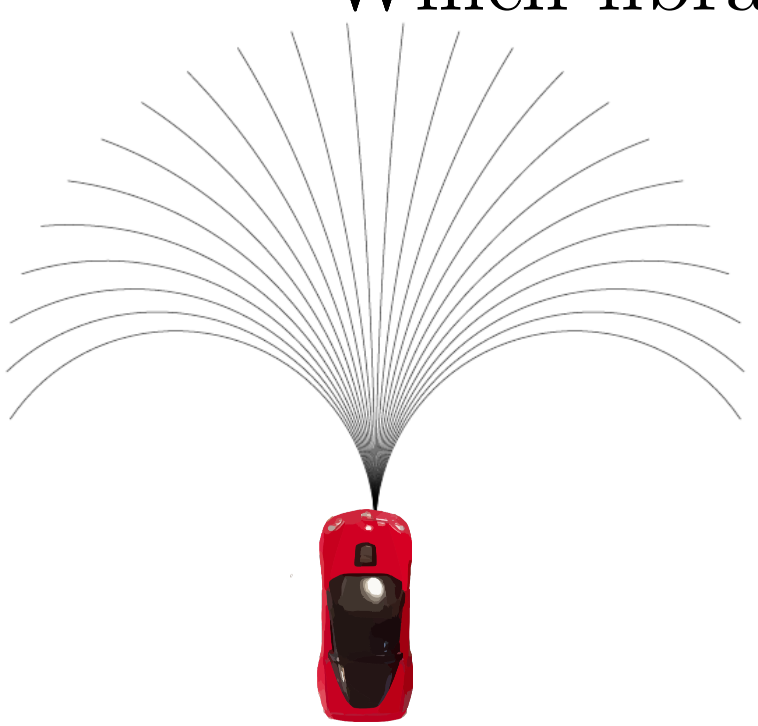

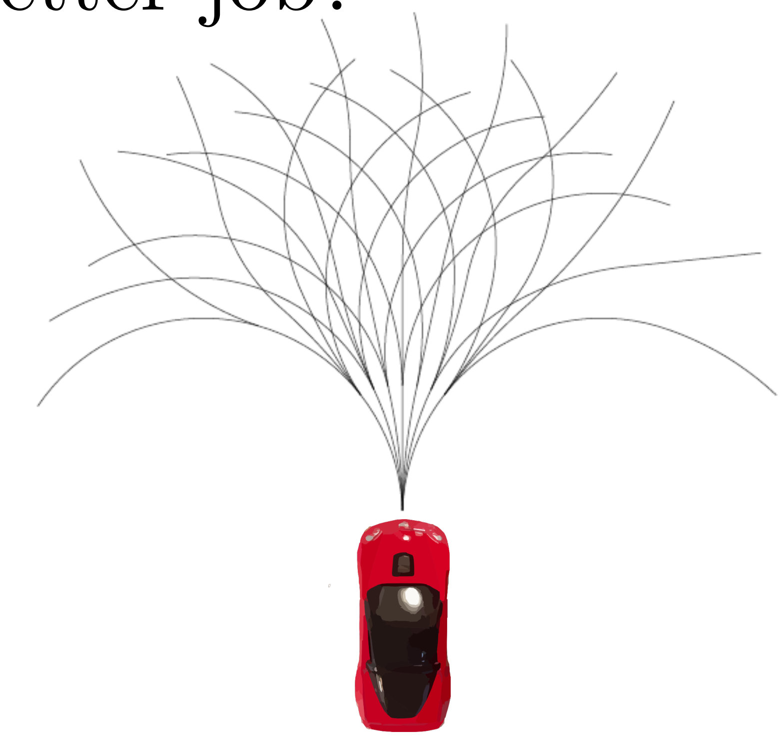

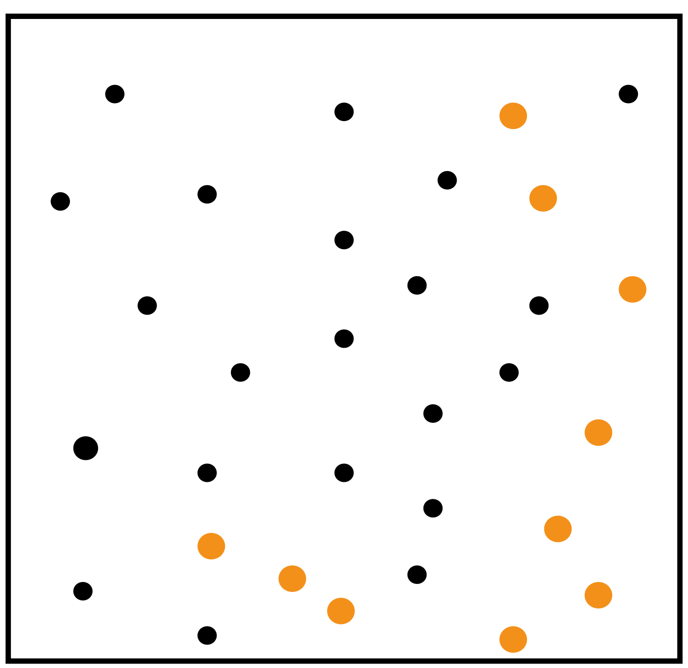

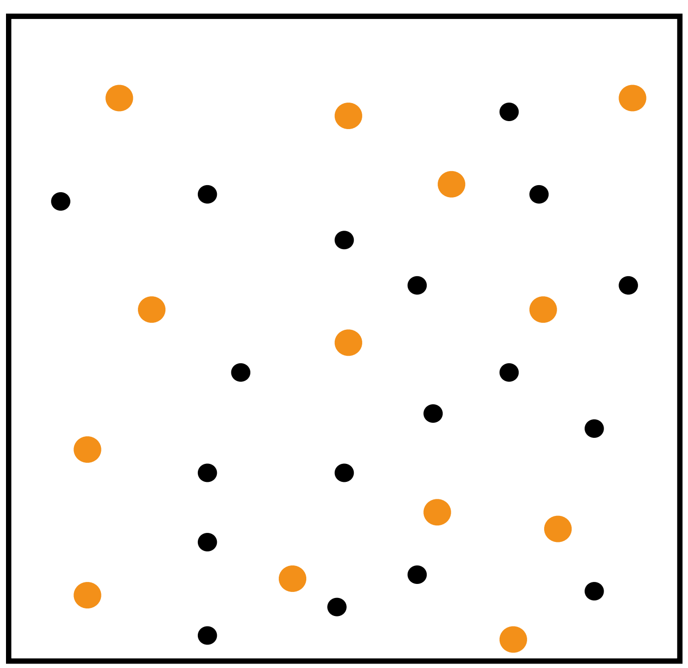

Remember, this is a space of paths!

Formalizing through dispersion

Dispersion is the radius of the largest ball around a point in D that does not have a point in S

$$Dispersion = \max_{\xi_1 \in D} \min_{\xi_2 \in S} d(\xi_1, \xi_2)$$

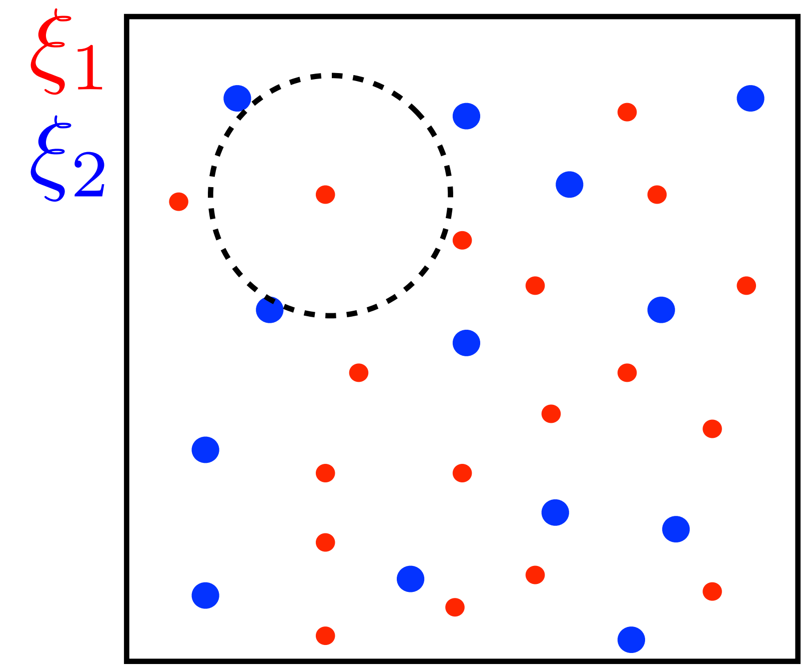

$$\textsc{Coverage}(S, D) \equiv - \max_{\xi_1 \in D} \min_{\xi_2 \in S} d(\xi_1, \xi_2)$$

The larger this ball, the worse the coverage of the space, so we want to minimize it!

Problem: Minimize dispersion

Given a dense path set D

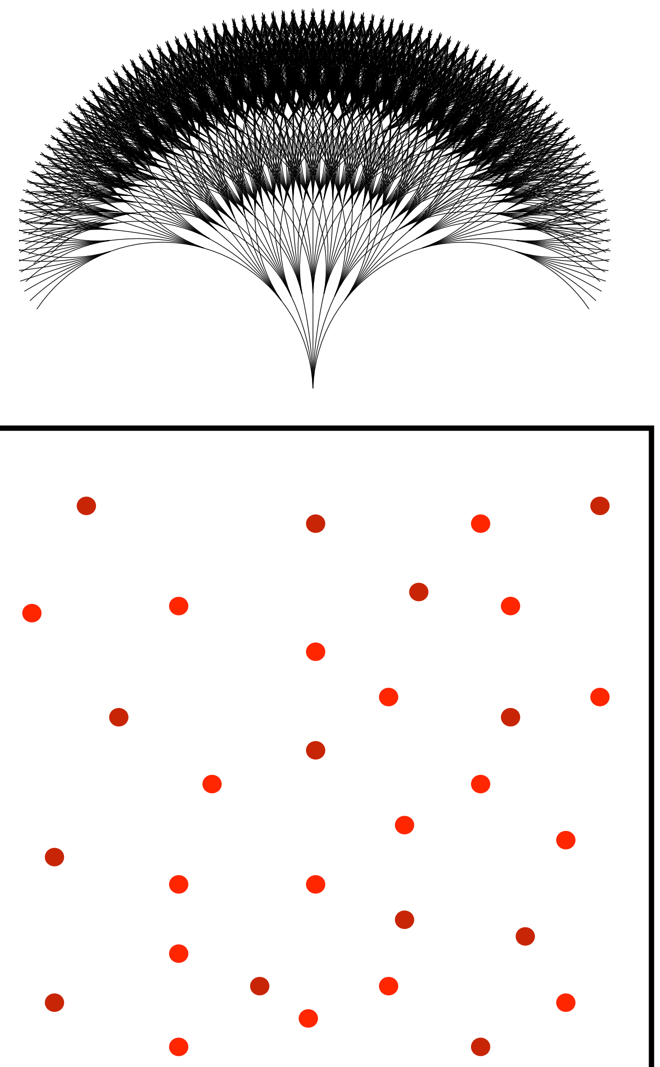

Find a sparse subset $S \subset D$

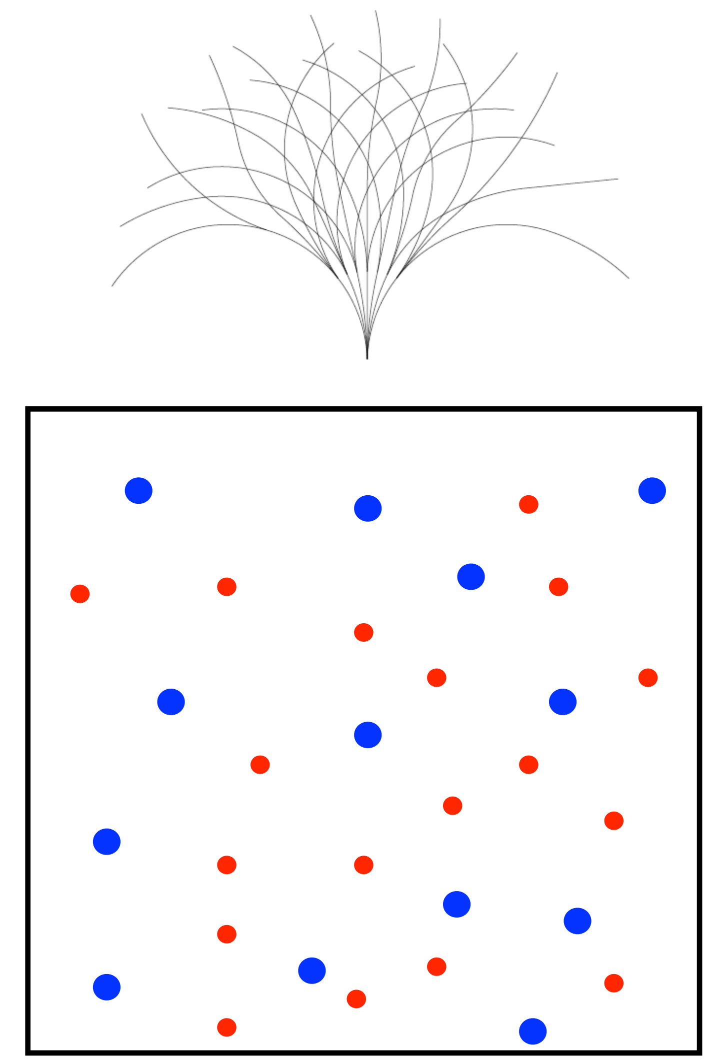

$$\arg\min_{S \subset D} \left[\max_{\xi_1 \in D} \min_{\xi_2 \in S} d(\xi_1, \xi_2) \right]$$

Questions

1. What distance metric should we use?

2. How do we solve the combinatorially hard problem of dispersion minimization?

Choosing a metric: Hausdorff distance

$$d(\xi_1, \xi_2) = \max\left(\max_{p \in \xi_1} \min_{q \in \xi_2} ||p - q||,\right.$$

$$\left.\max_{p \in \xi_2} \min_{q \in \xi_1} ||p - q||\right)$$

Find the p from which the **minimum** distance to all q is **maximal**.

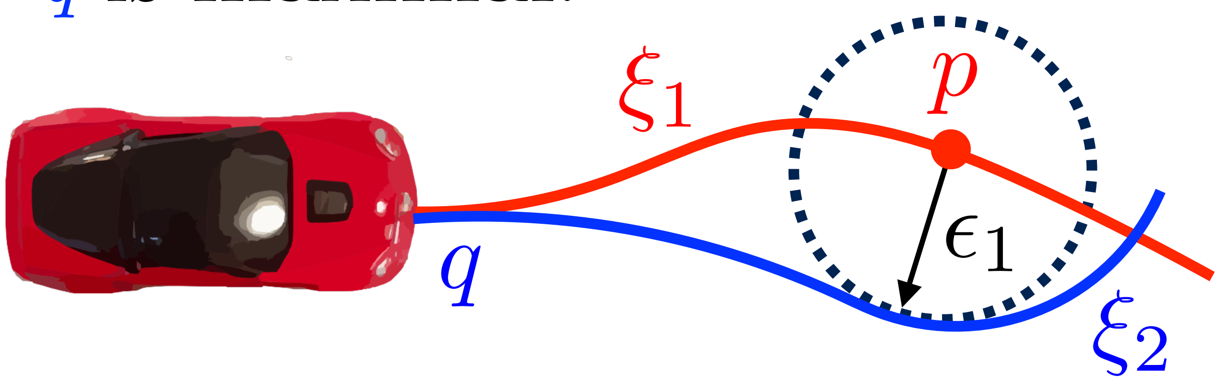

$$\max(\epsilon_1, \epsilon_2)$$

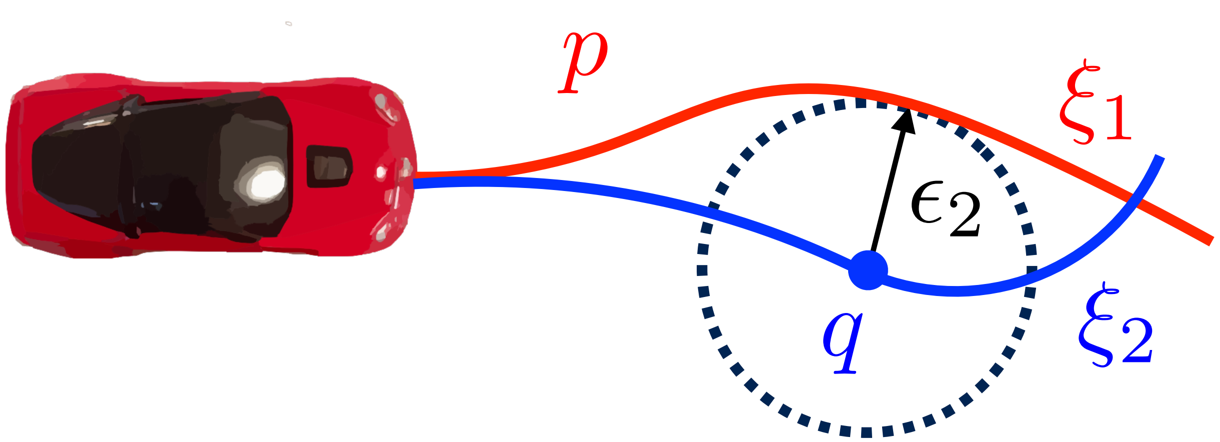

Find the q from which the **minimum** distance to all p is **maximal**.

Trajectory library in action

Extra Credits

What is wrong with control space?

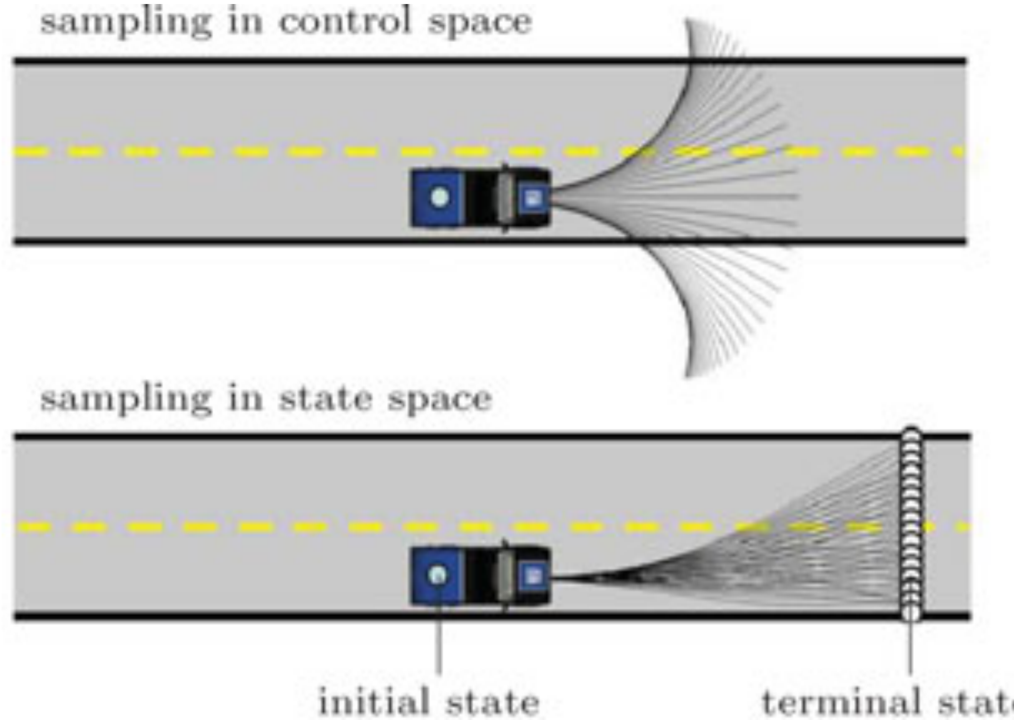

Uniform sampling in control space is non-uniform in state space.

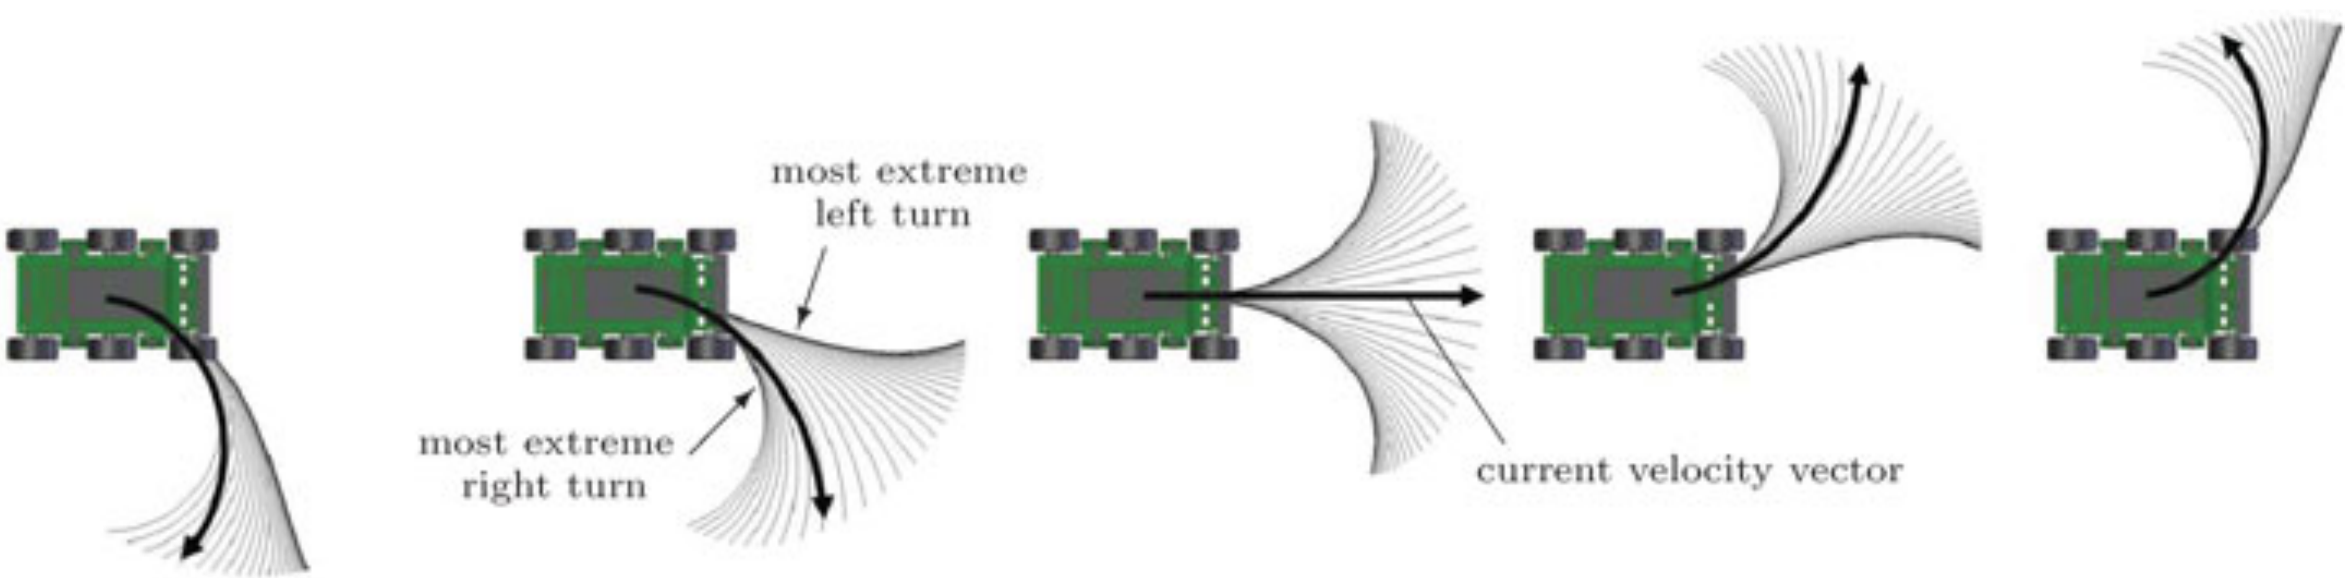

T.Howard "State Space Sampling of Feasible Motions for High-Performance Mobile Robot Navigation in Complex Environments" JFR, 2008

Why should we care about state space?

$$u(t) \longrightarrow \int_0^t f(x(t), u(t))dt \longrightarrow x(t)$$

Control trajectory

Dynamics

State trajectory

Space in which we sample

Space in which cost functions defined

State space sampling

Do optimization to different terminal points

Store these trajectories

Load path set dependent on the global cost

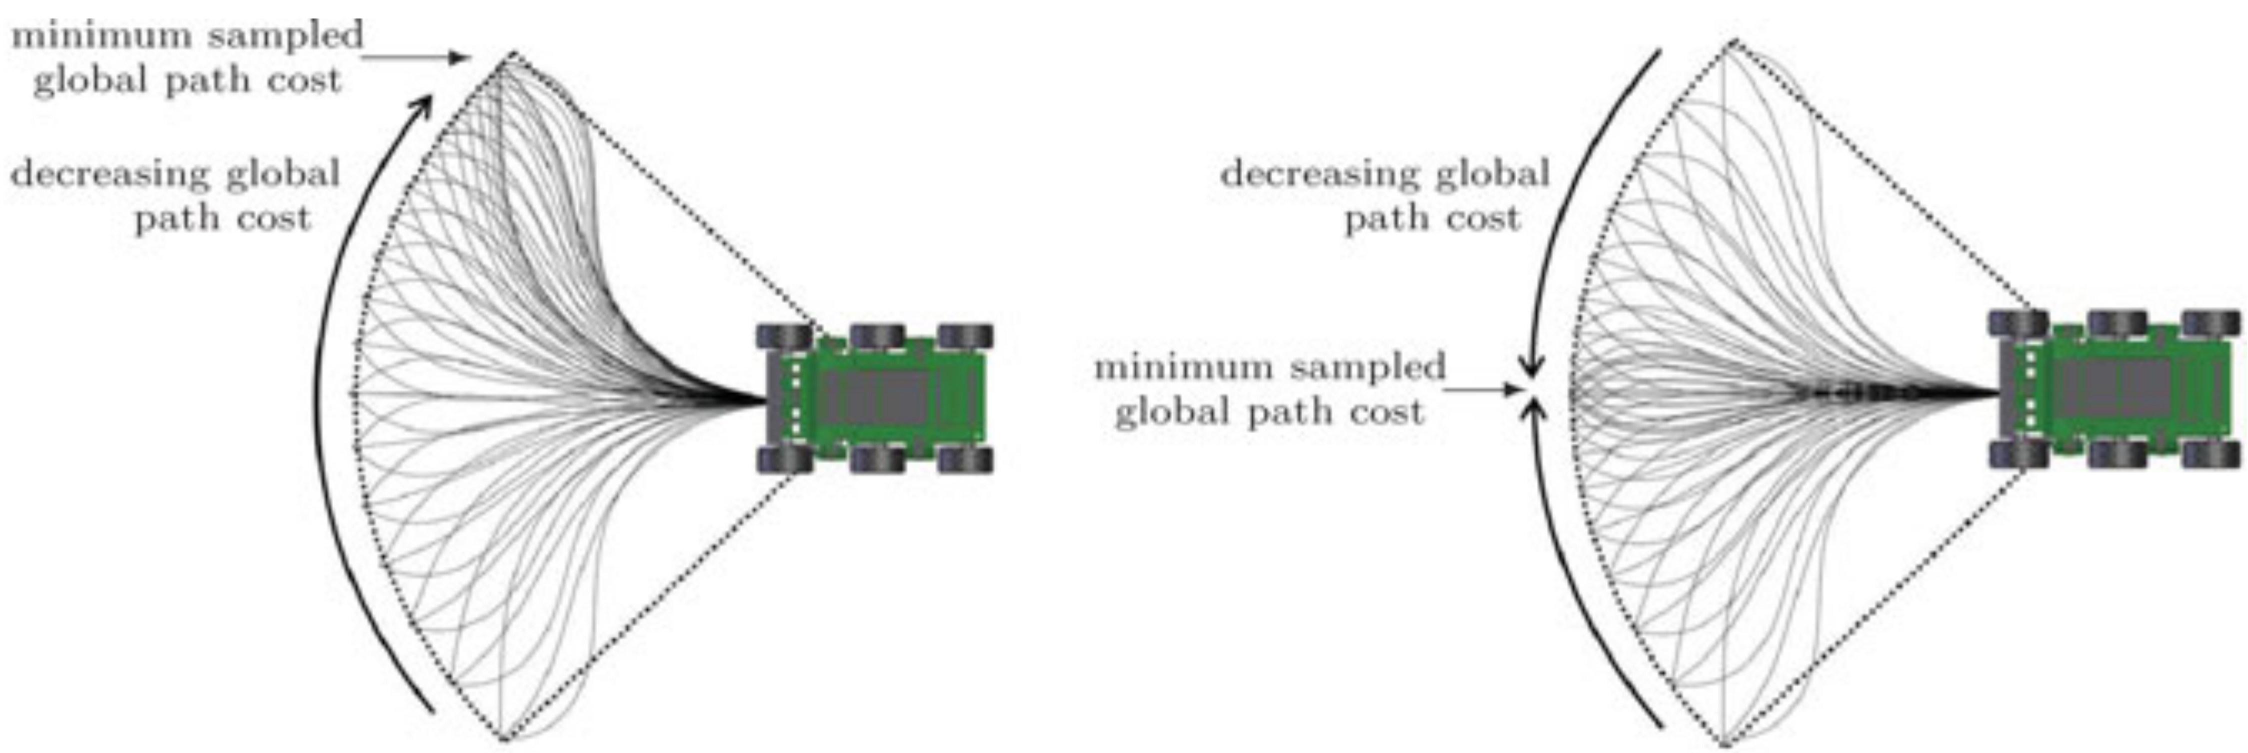

(a) Minimum global cost to the right of the vehicle (b) Minimum global cost in front of the vehicle

T.Howard "State Space Sampling of Feasible Motions for High-Performance Mobile Robot Navigation in Complex Environments" JFR, 2008

Using learning to select libraries

Instead of selecting a static library, we can select a contextual one

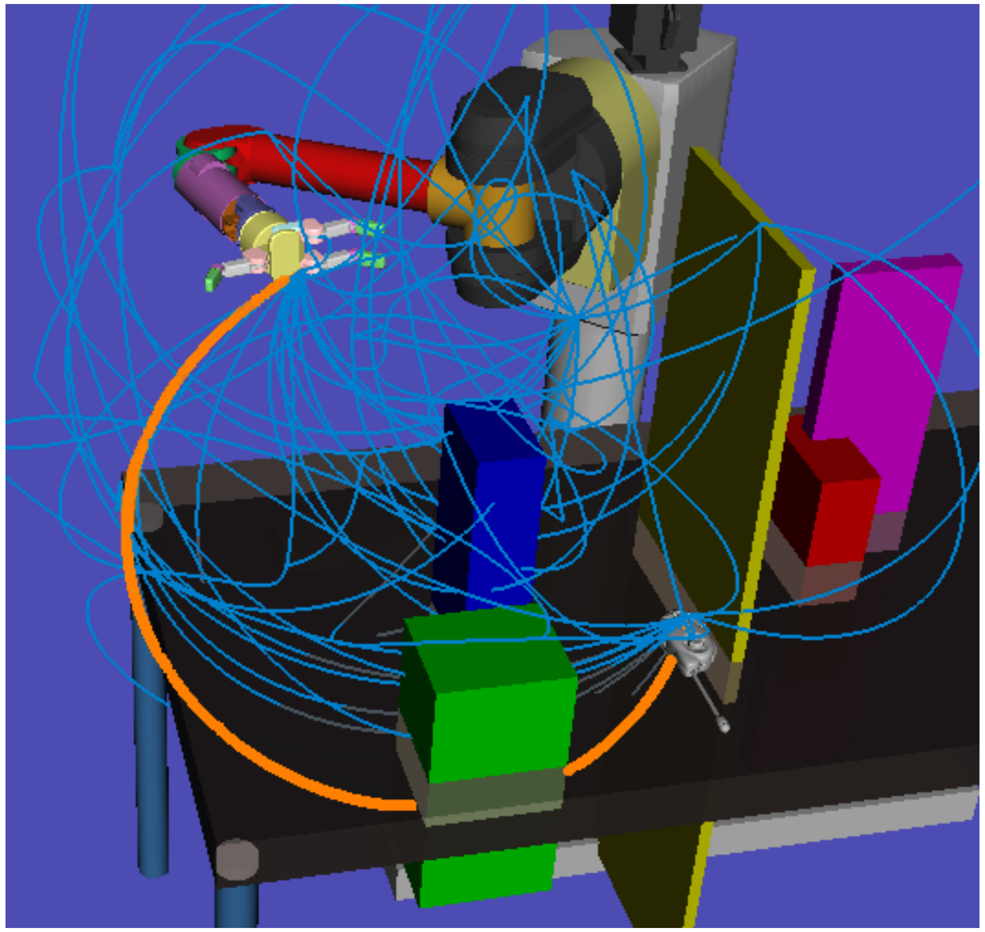

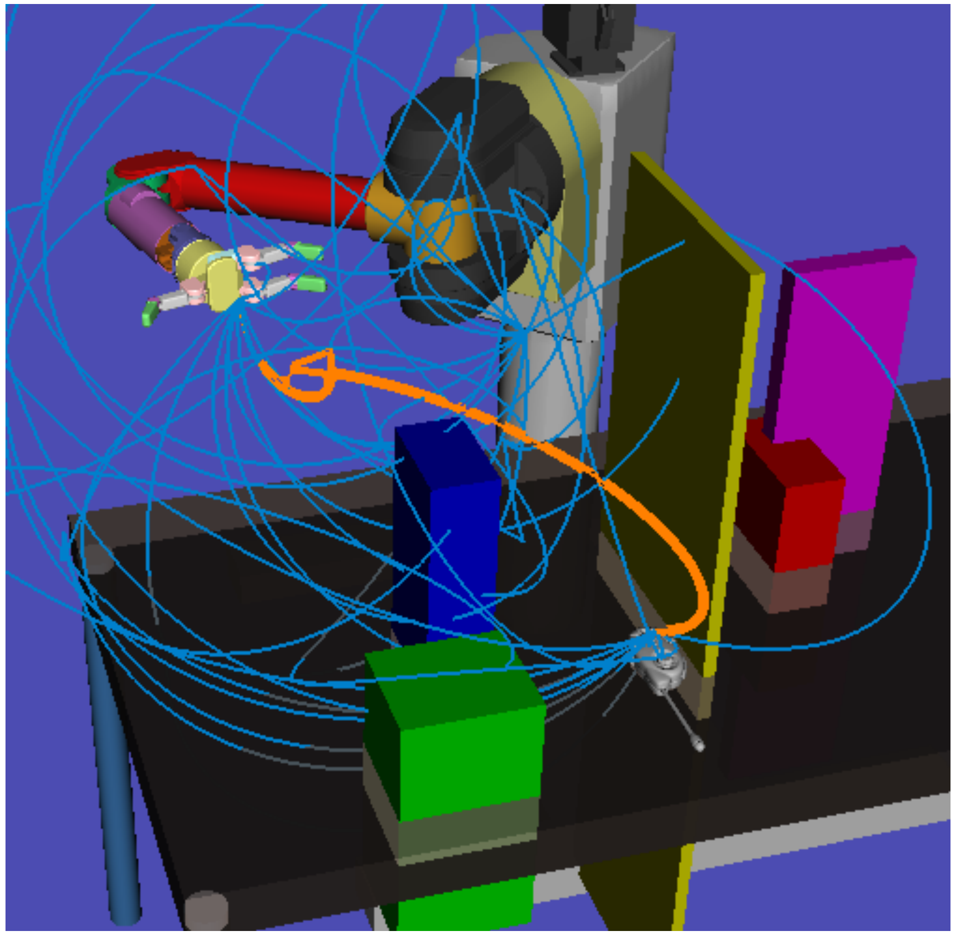

D.Dey et al "Contextual Sequence Prediction with Application to Control Library Optimization", 2013

Guaranteeing safety with libraries

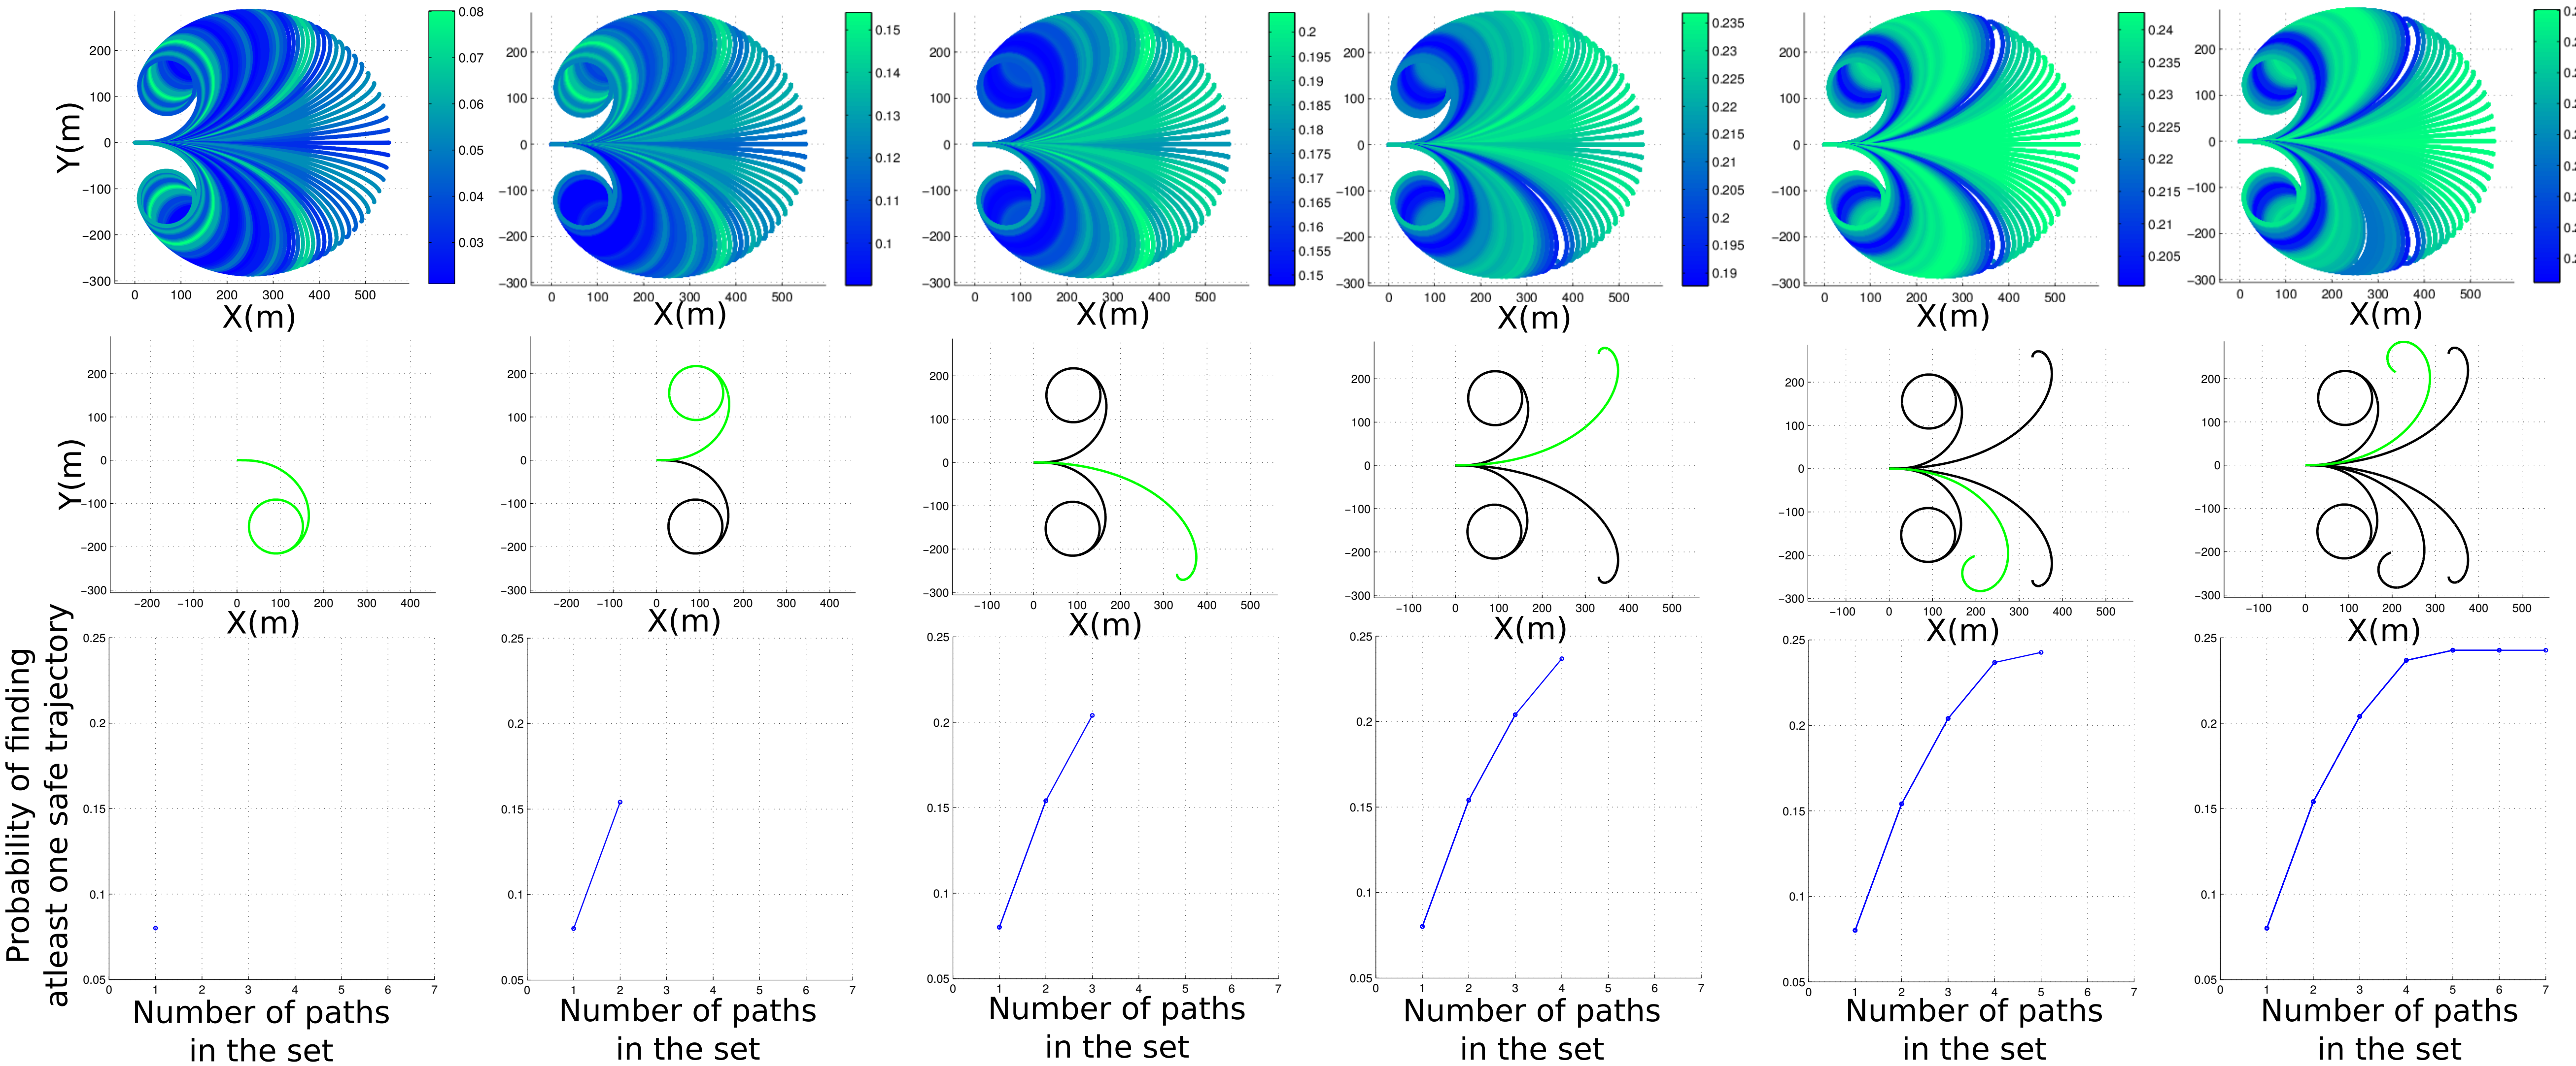

S.Arora "Emergency Maneuver Library – Ensuring Safe Navigation in Partially Known Environments", 2015