Path Planning in Unknown Environment by Optimal Transport on Graph

Haomin Zhou

School of Mathematics, Georgia Tech

Collaborators: Magnus Egerstedt (ECE, GT), Haoyan Zhai (Math, GT)

Partially Supported by NSF, ONR

Optimal Path In Dynamical Environment

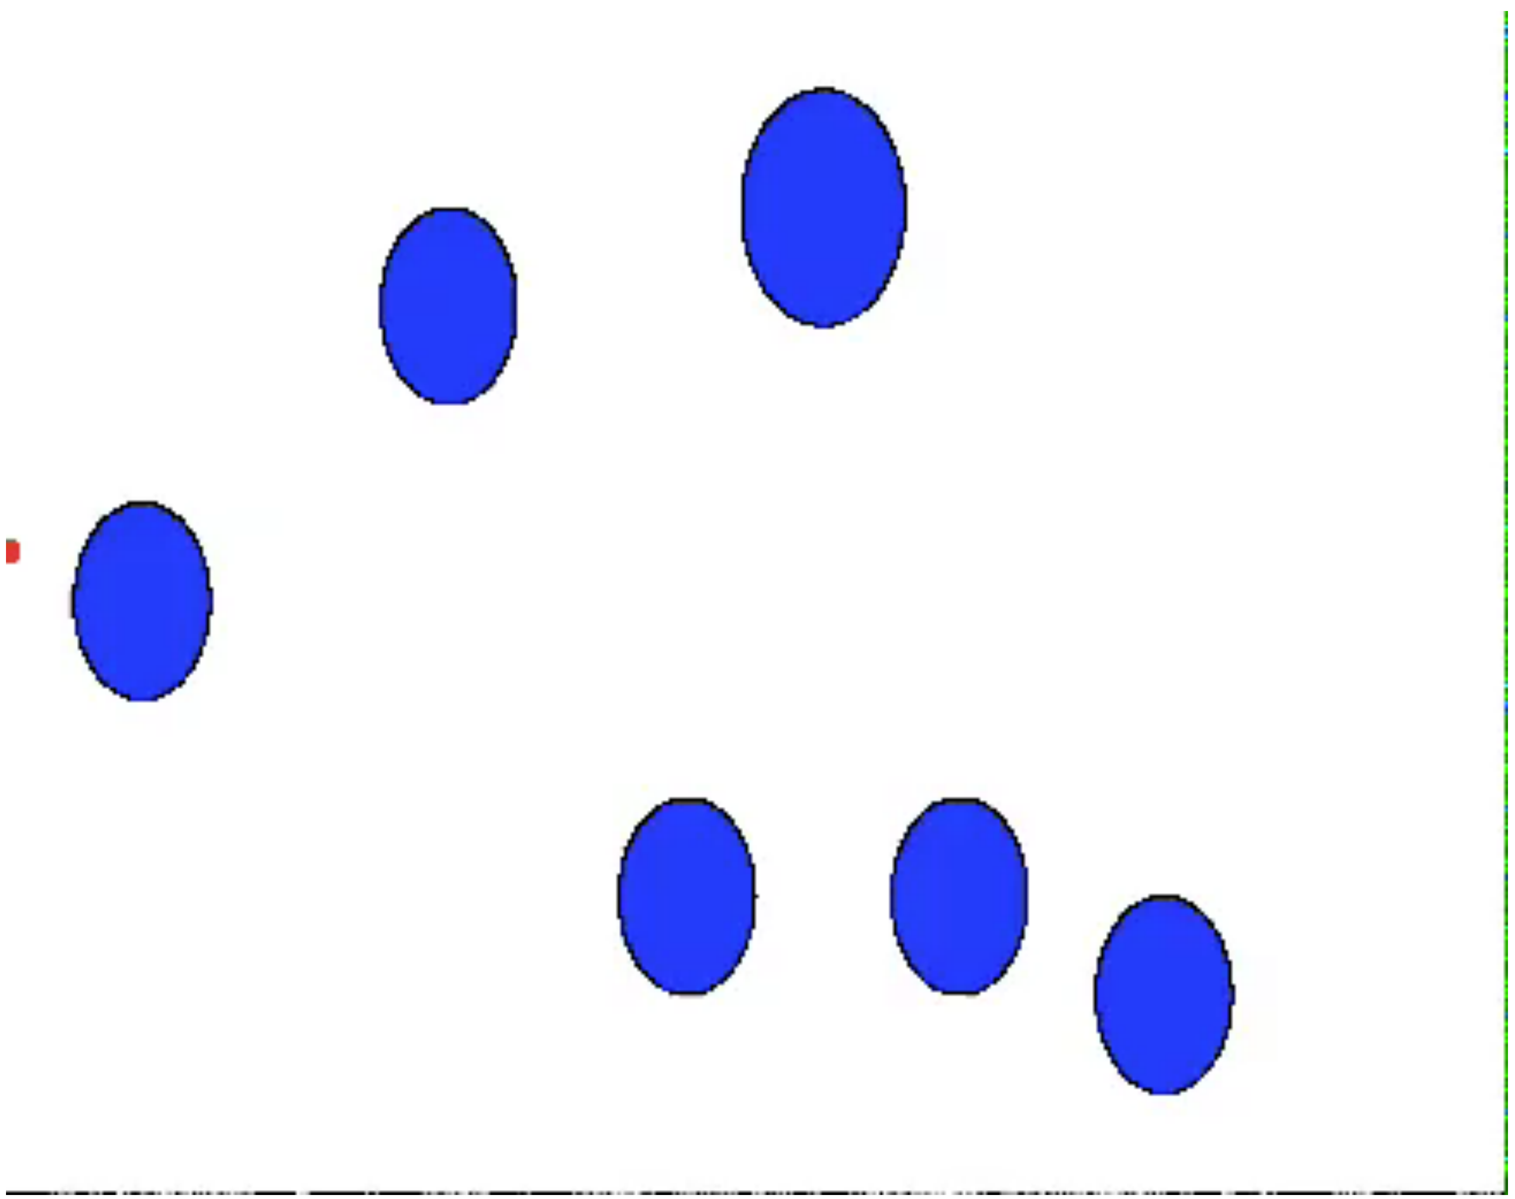

Method of Evolving Junctions (MEJ) (*Automatica* 2017, IJRR 2017, with Chow-Egerstedt-Li-Lu,).

Multi-Agent System

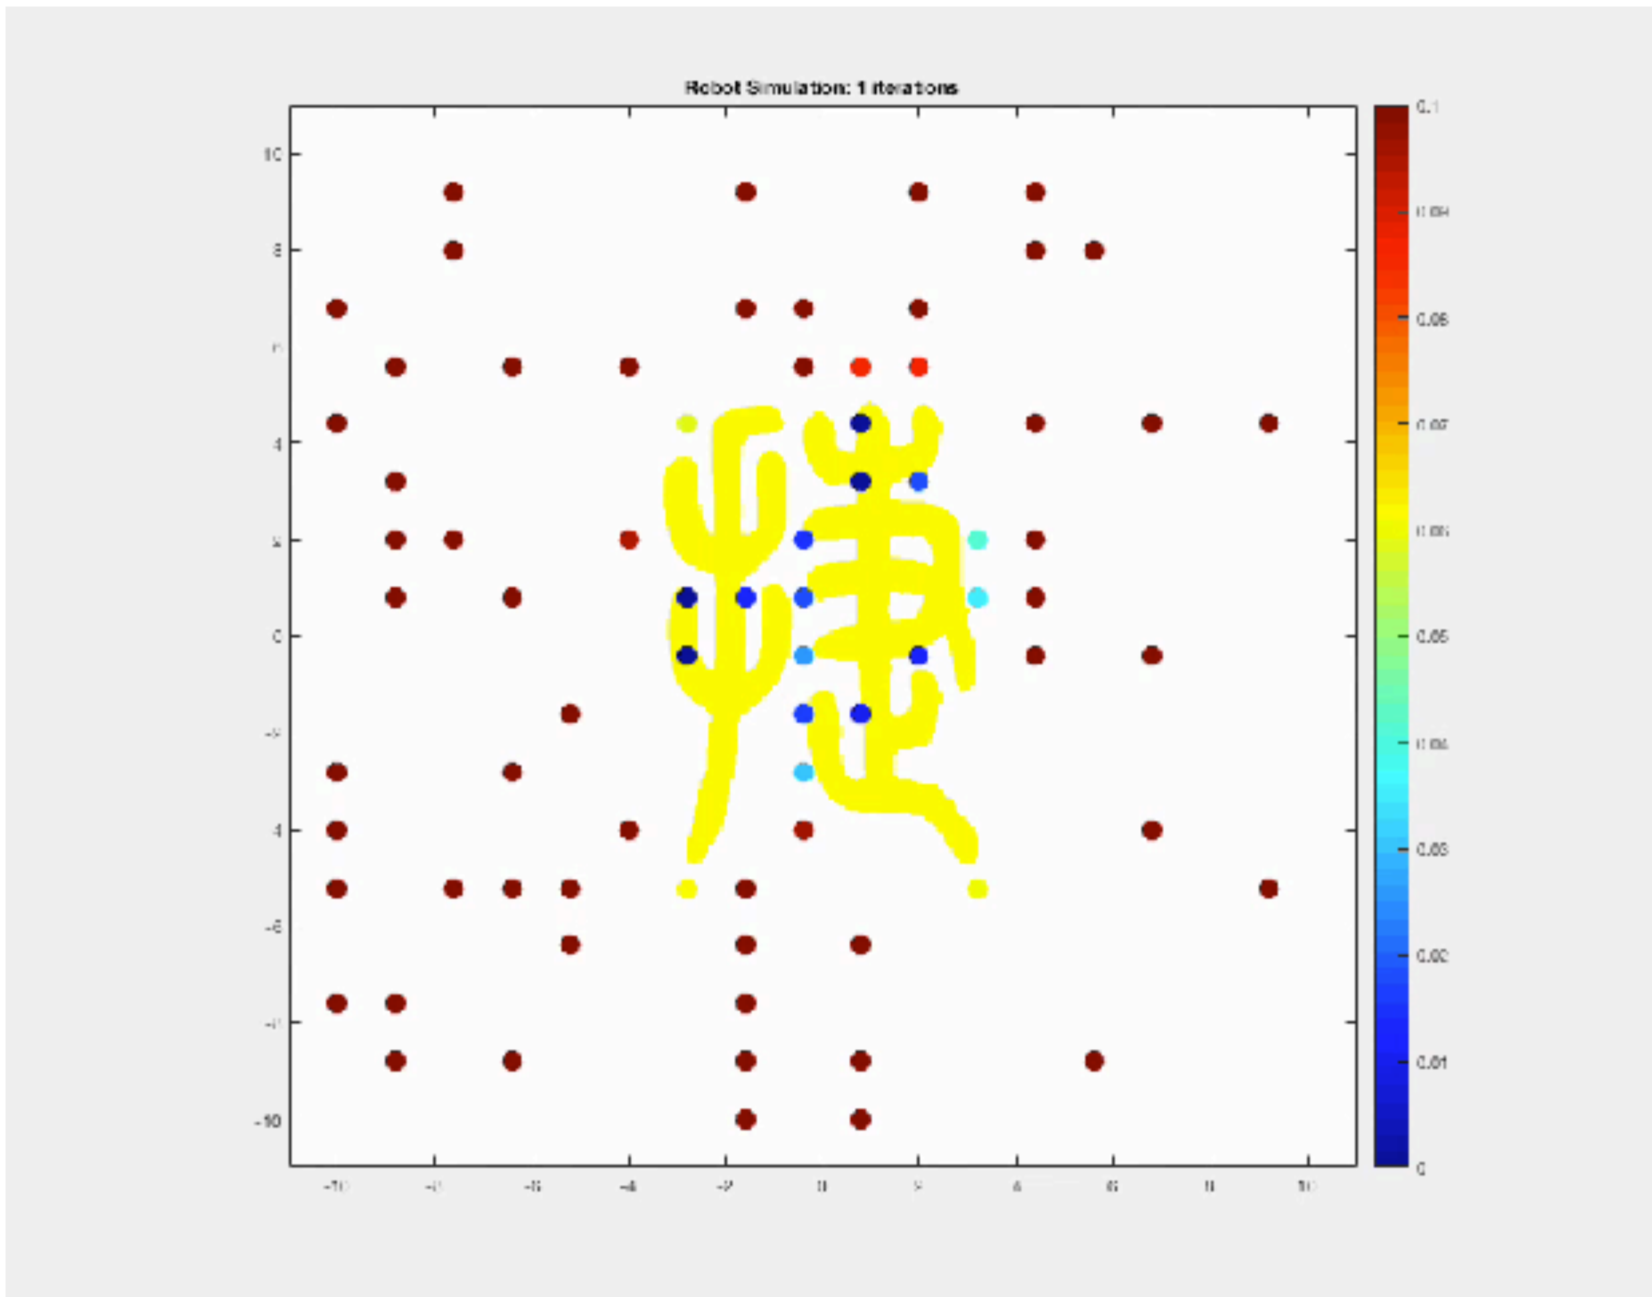

The robots have limited detection ranges.
Path generated by Intermittent Diffusion (with Egerstedt-Frederick).

Path Exploration in Unknown Environment

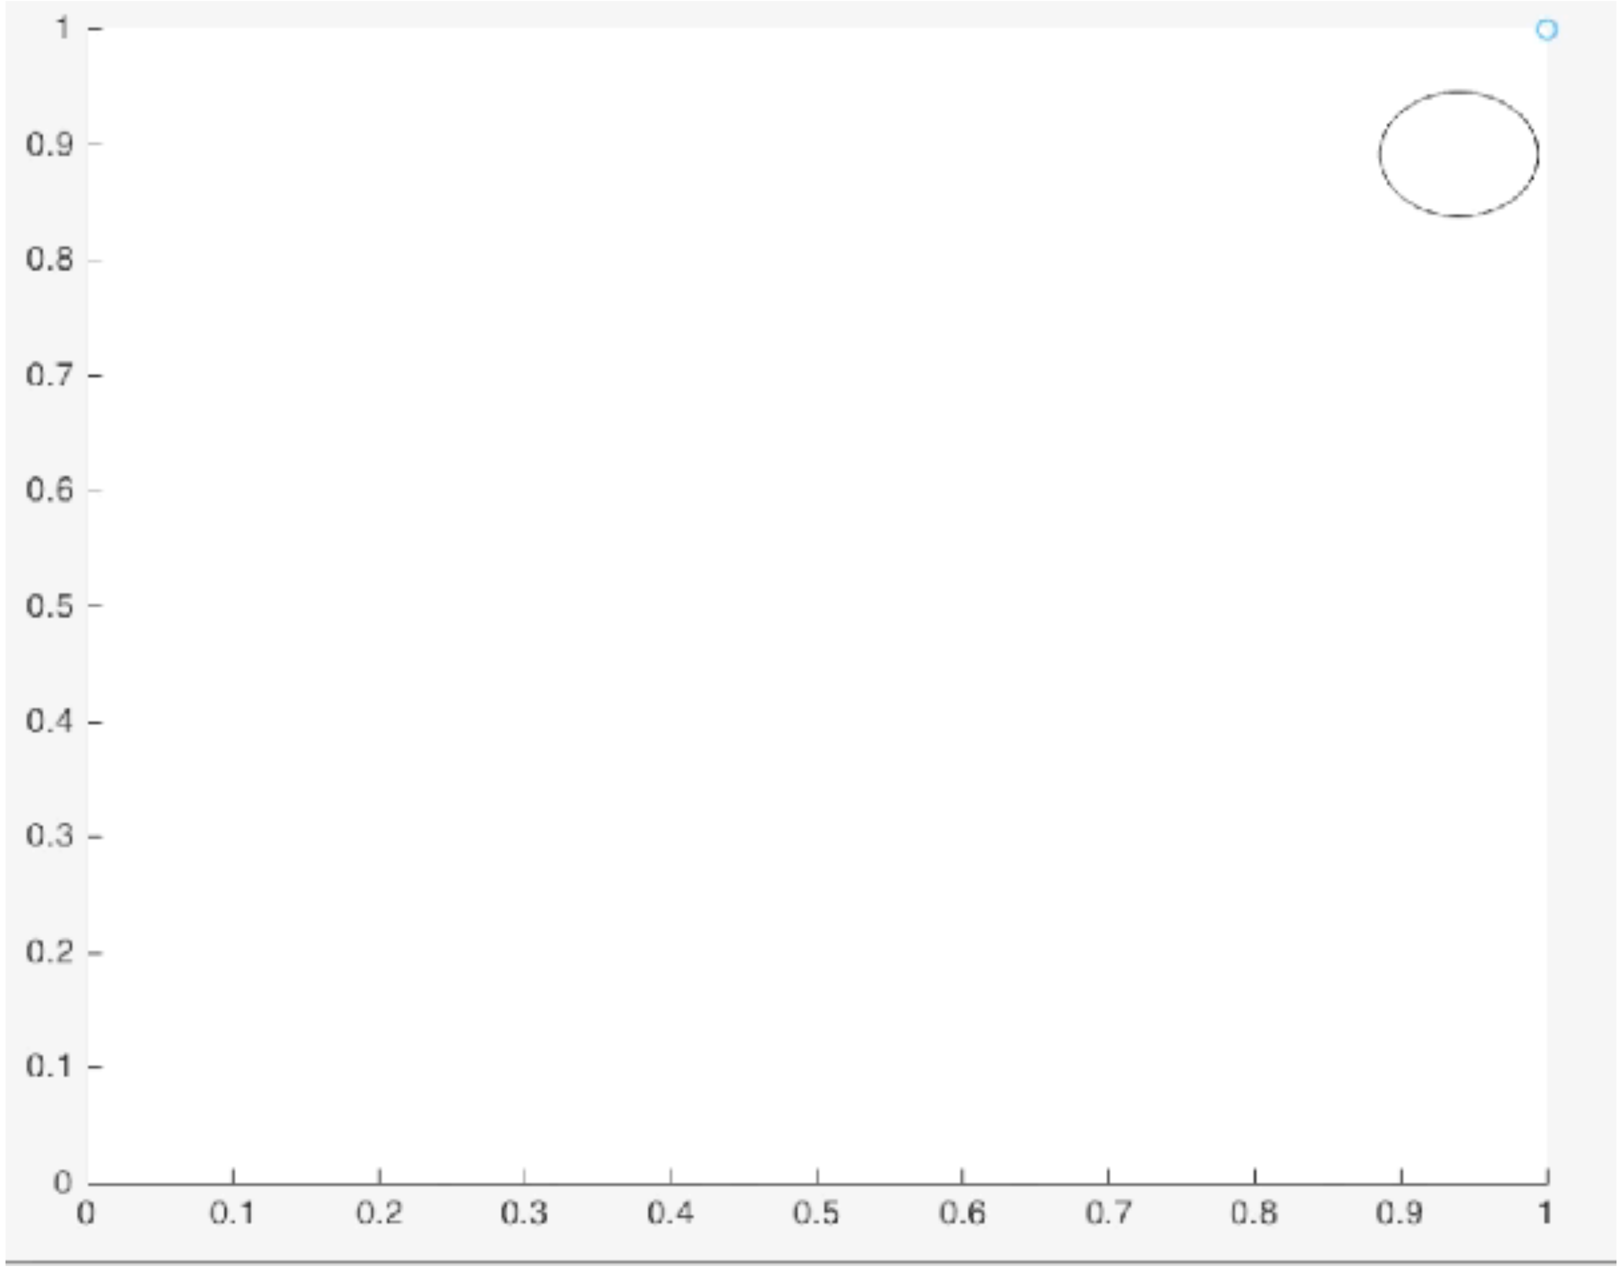

The robot has a limited detection range.

Path Exploration in Unknown Environment

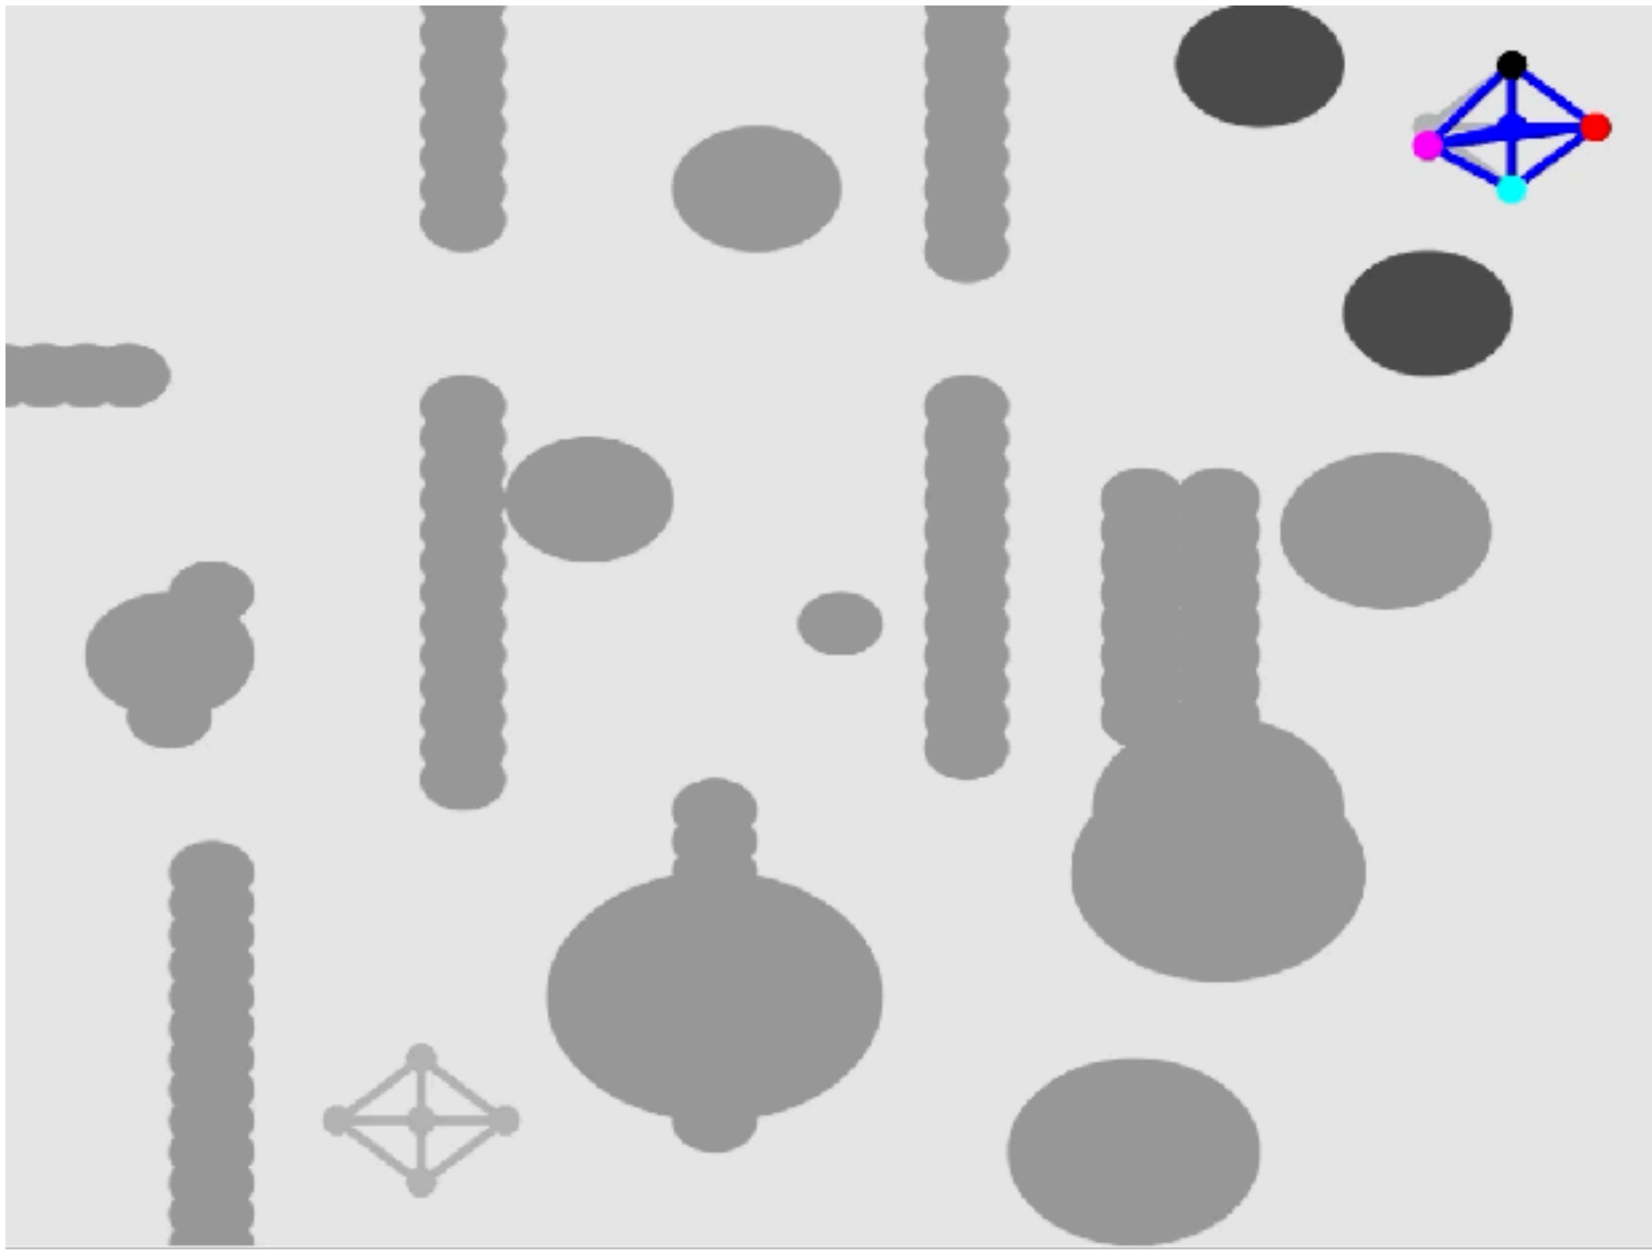

The robots have limited detection ranges.

Outline

- Path planning in unknown environments
- Optimal transport on finite graphs
- General control with unknown constraints

Path Planning in Unknown Environments

Problem: Find a continuous curve $\gamma(t)$ in $\Omega \subset \mathbb{R}^d$ such that

$$\gamma(0) = x_0, \gamma(T) = x_f,$$
$$\phi(\gamma(t)) \geq 0 \text{ for all } t \in [0, T],$$
$$\hat{\psi}(\gamma(t), \gamma, t) \geq 0 \text{ for all } t \in [0, T],$$

$\phi(x) \leq 0$ are the known constraints,
$\psi(x) \leq 0$ are the unknown obstacles.

$$\hat{\psi}(x, t, \gamma) = \begin{cases} \psi(x) & \text{if } d(x, \gamma(\tau)) \leq R \text{ for some } \tau \leq t \\ 0 & \text{otherwise} \end{cases}$$

Constraints are expressed in terms of level set functions.

Challenges

- Local traps and replanning,
- Narrow pathways,
- Collisions,
- Communications,
- Computational cost in higher dimensions.

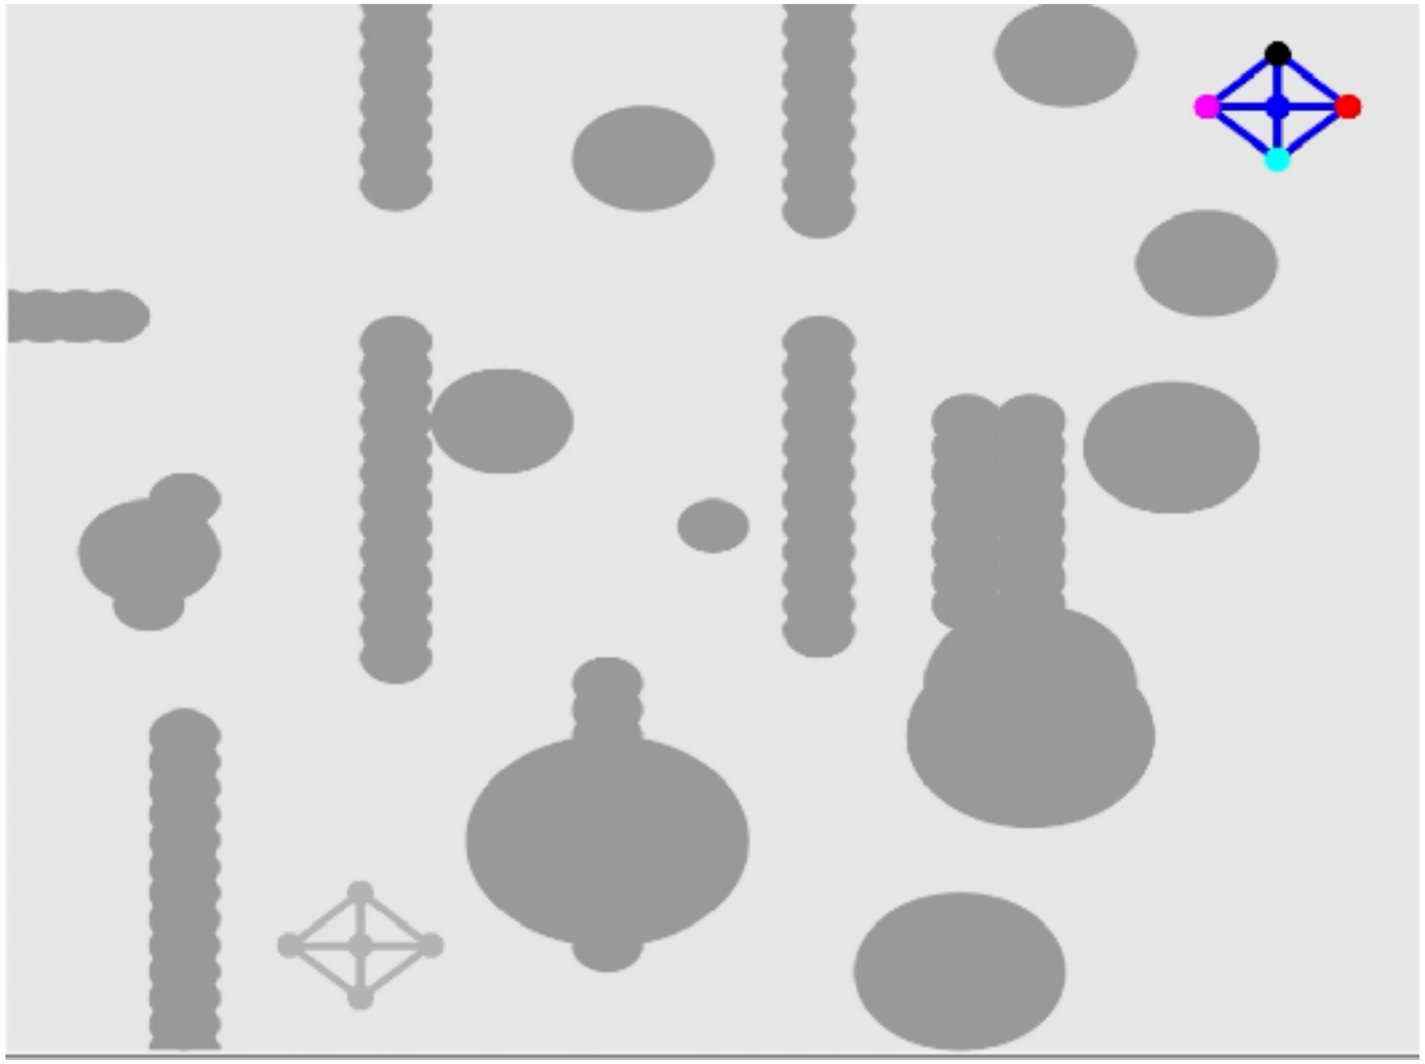

Existing methods

- Bug family: Bug0, Bug1, Bug2, TangentBug, DistBug, …
- Probabilistic Road Map (PRM),
- Rapid-growing Random Tree (RRT), RRT* (dynamical version),
- Artificial Potential Field (APF),
- Graph based methods (Dijkstra style): A*, D, D*, focus-D*, D*-lite and more,
- Genetic algorithm, Neural network, fuzzy logic, fast marching tree, and many more.

The convergence for many of the methods, if exists, is asymptotic.

Our Algorithm

Idea: potential guided, tree based 2-layer iterations.

3 main steps:

- Tree generating,
- Path finding,
- Environment updating.

Potential is used to ensure convergence, and Trees are used to control the computation cost in high dimensions.

Properties of Our Algorithm

Proposition

There exists a unique path from initial to target configurations over the generated graph G. And if the path is denoted by $\{x_i\}_{i=0}^{q} \subset V$ with $x_0 \to x_1 \to \cdots \to x_q = x_f$, where x_i is the ancestor of x_{i+1}.

We use back tracing to find the path:

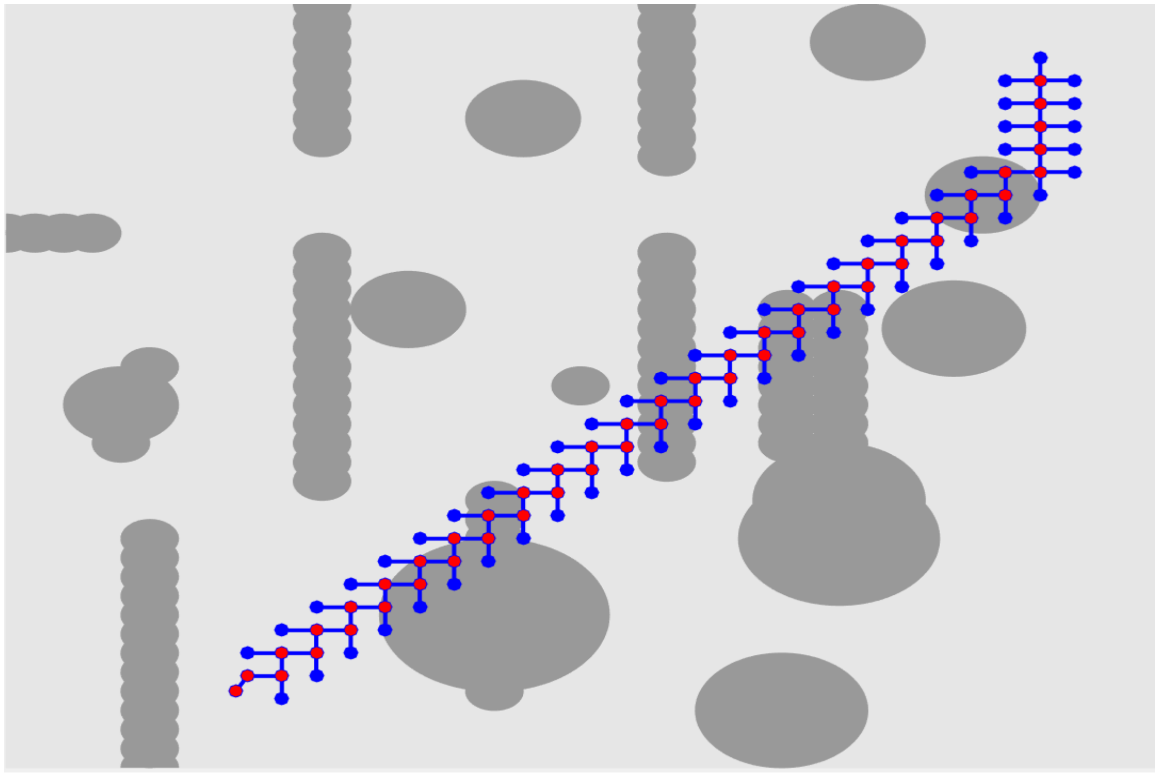

Convergence Analysis

Theorem

Assuming that

$$\sup_{\gamma \in \Gamma} \inf_{t \in [0,T]} \sup_{r \geq 0} \{r : B(\gamma(t), r) \cap \mathcal{O} = \emptyset\} = L > 0,$$

and

$$l < \frac{2L}{\sqrt{n}},$$

where n *is the dimension of* Ω *and* l *is the step size of the graph generation, the graph generation terminates in finite iterations. The generated graph* $G = (V, E)$ *is connected and has a finite number of vertices* $|V| < \infty$ *with* $x_s, x_t \in V$.

The tree generating iteration stops in finite steps.

Convergence Analysis

Theorem

Let $\{\gamma_i\}_{i=1}^m$ be the paths produced by the algorithm with $\{T_i\}_{i=1}^m$ being the stopping time set. If we use the same assumptions in the previous theorem and

$$\sup_i \inf_\epsilon \left\{\epsilon : B(\gamma_i(T_i), \epsilon) \cap \mathcal{O}_c^{T_i} \neq \emptyset\right\} = q < R,$$

holds, then $m < \infty$.

The outer iteration stops in finite steps.

Our algorithm stops in finite steps.

Examples

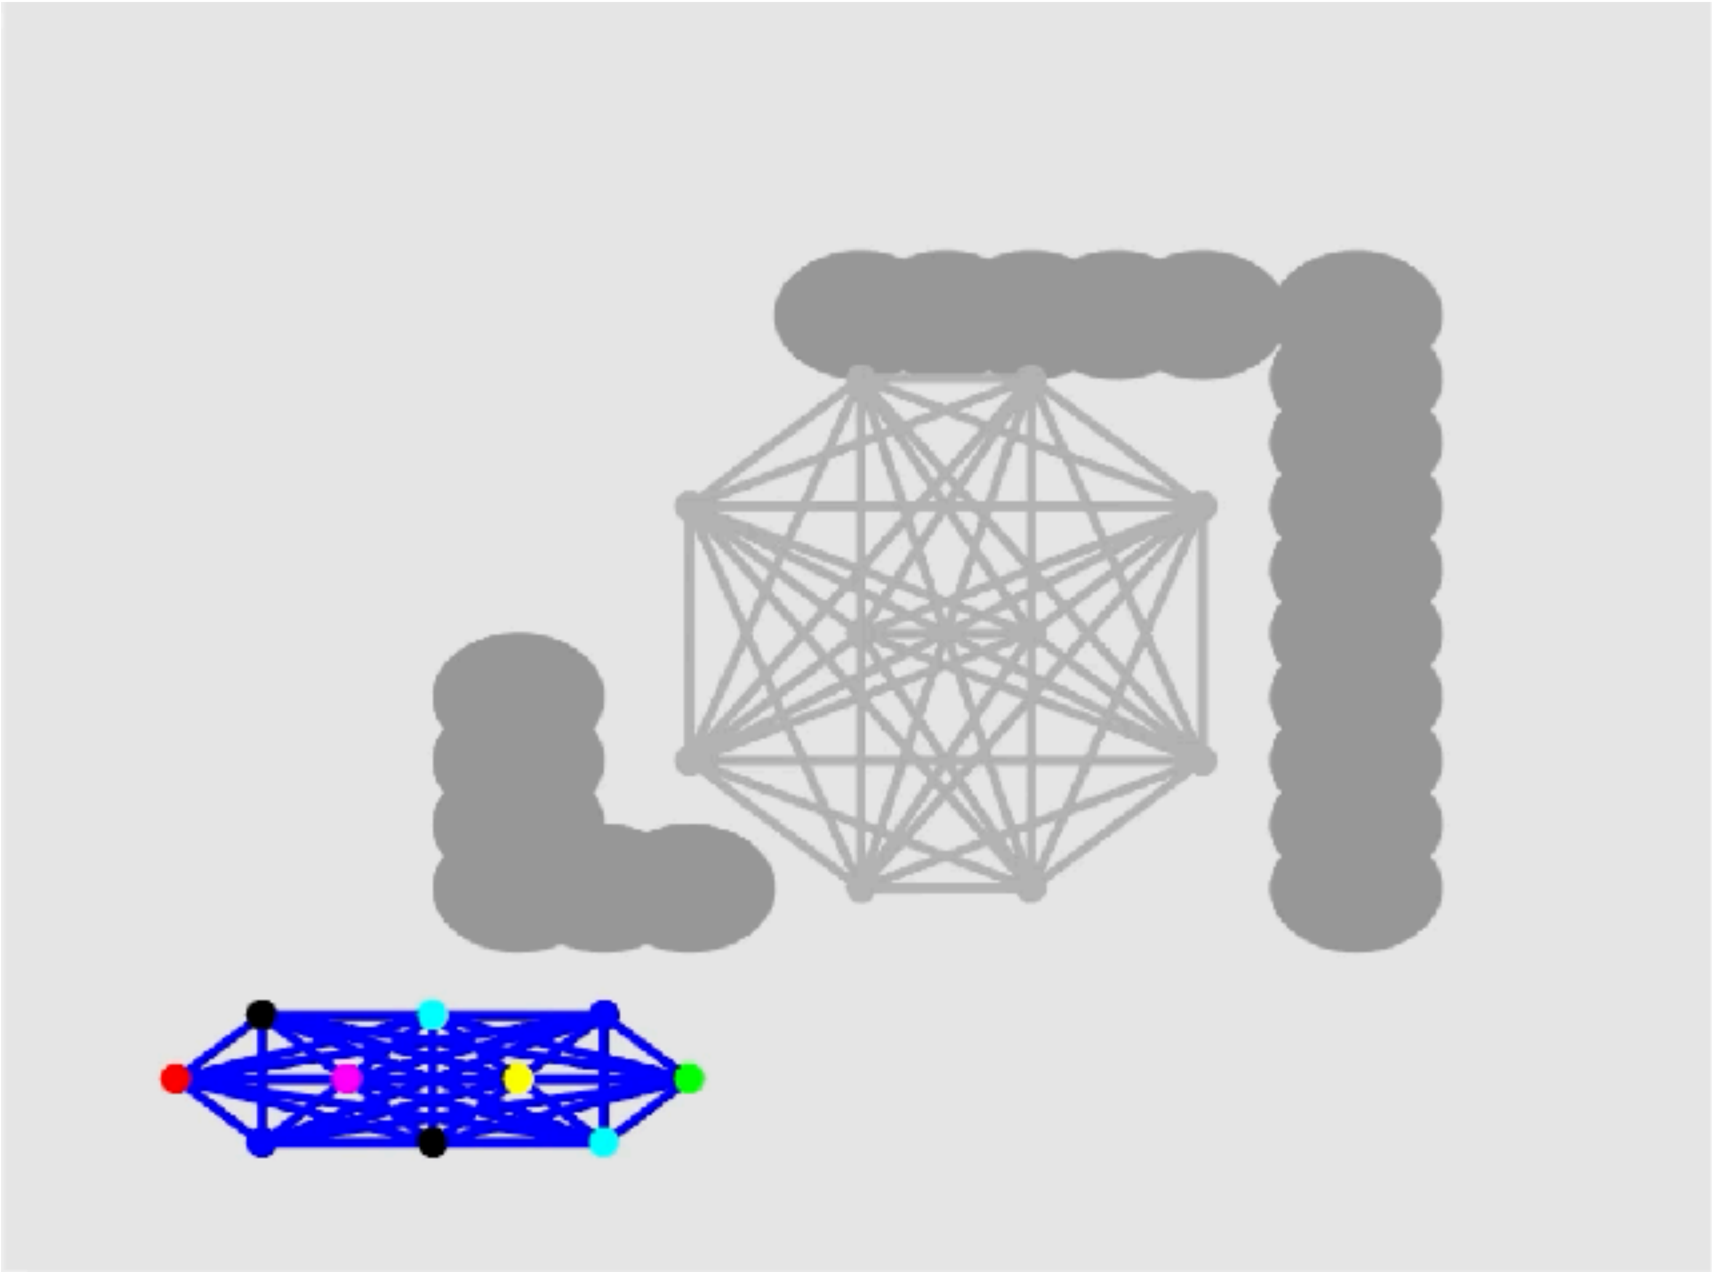

A 10-robot (20 dimensional) example. The entire computation is within 1 minute in Matlab on a laptop.

Examples

A 3-robot example. Most area is not explored.

Summary of the Properties

- The algorithm is deterministic, stops in finite steps,
- Guarantees to find a feasible path if there exists one,
- If the algorithm stops without returning a path, there isn't one that can be identified by the step size.
- The growing rate for the tree is linear, not exponential, w.r.t. the dimension of the configuration space,
- Explores only a limited part of the configuration space.

The method is inspired by optimal transport on trees with intermittent diffusion.

Limited Exploration Region

Theorem

Given any known environment, the generated graph G is bounded by $\mathcal{R}$, produced by evolution of Fokker-Planck equation in the same environment:

$$G \subset \bigcup_{x \in \mathcal{R}} Box(x, l).$$

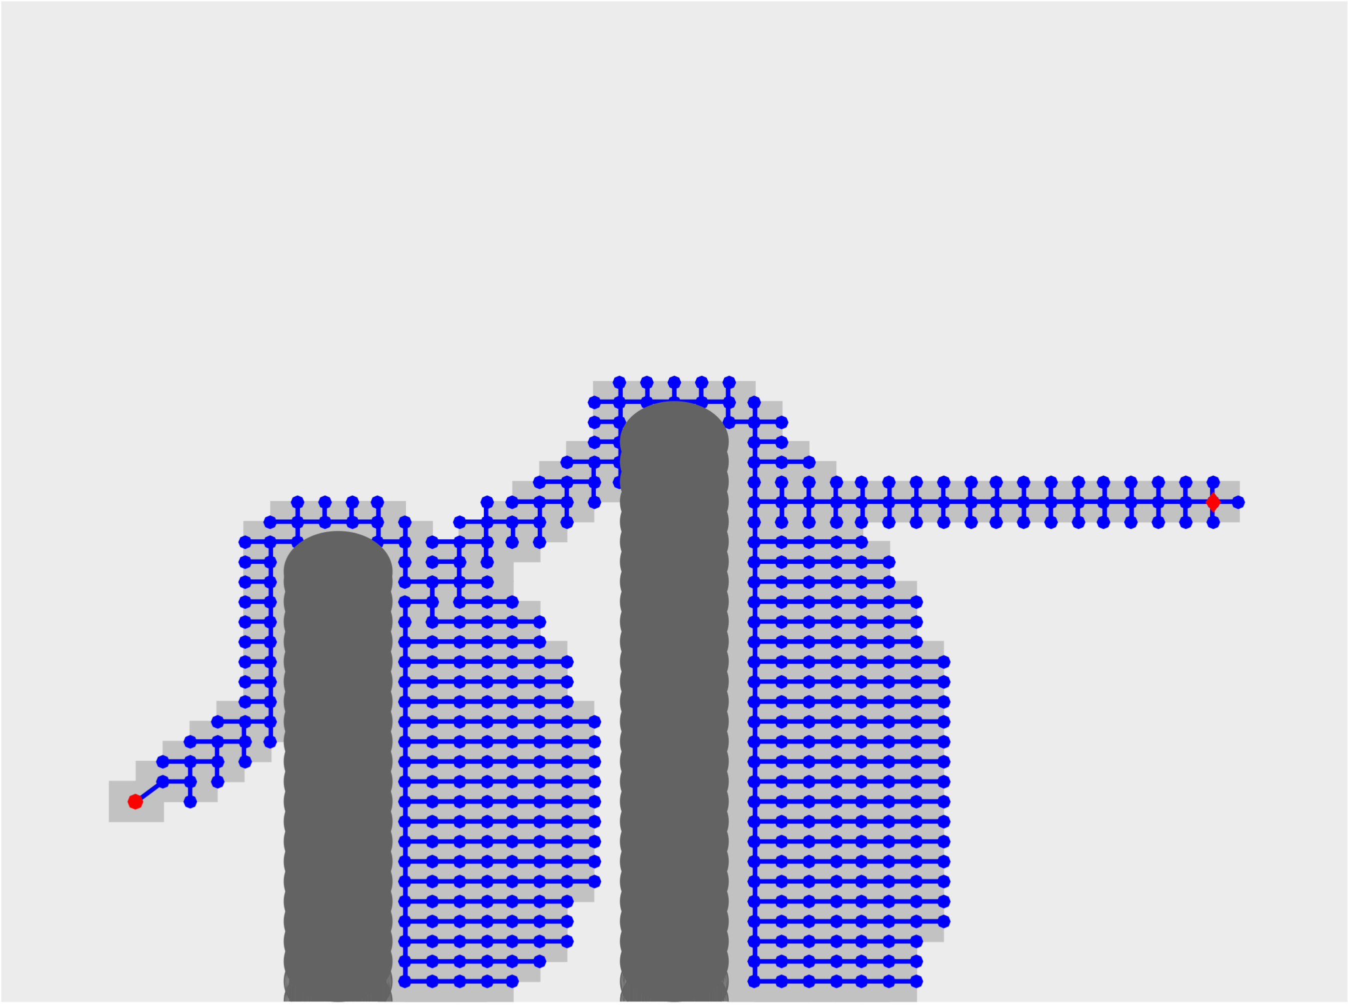

The tree generation contains 2 phases:
projected gradient and diffusion.

Limited Exploration Region

Theorem

Assuming that the robots only stop on node points with assumptions in previous theorems and $R > L$, the complete path γ generated by the algorithm in the unknown environment satisfies

$$\gamma \subset \bigcup_{x \in \mathcal{R}} Box(x, l),$$

where $\mathcal{R}$ is produced with the full knowledge of the environment.

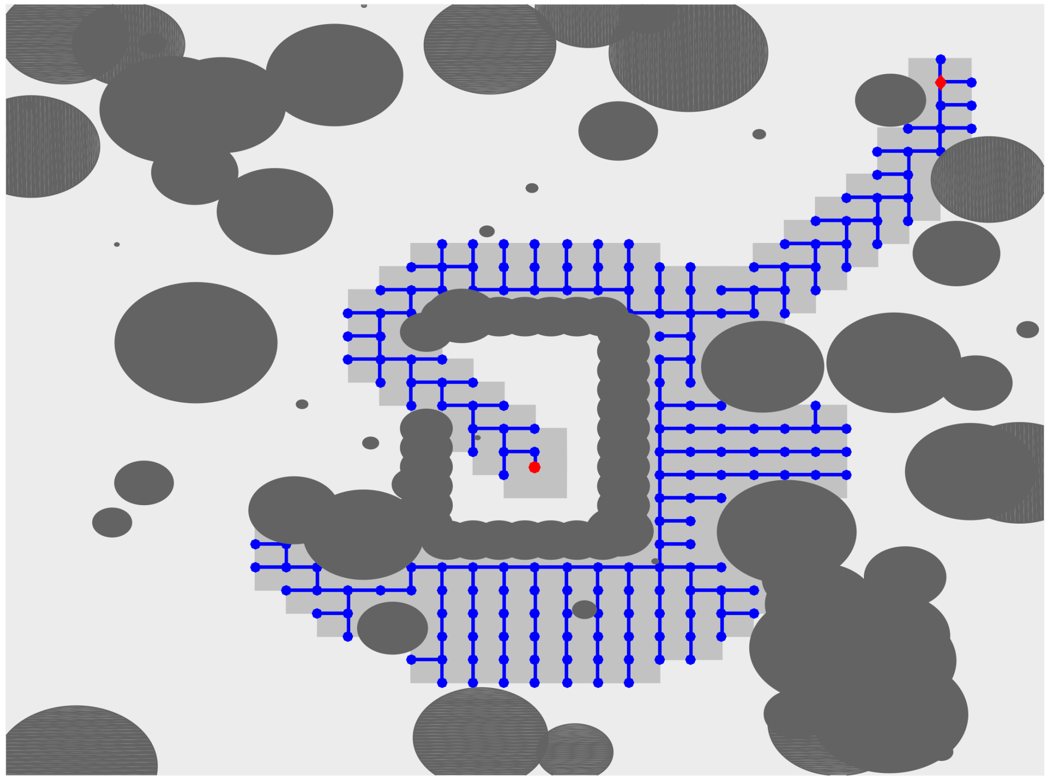

The evolution of Fokker-Planck equation on graph has intermittent diffusion (diffusion coefficient is turned on-and-off).

Optimal Transport

- **Optimal Transport:** Monge (1781), Kantorovich (1942), Otto, Kinderlehrer, Villani, McCann, Carlen, Lott, Strum, Gangbo, Jordan, Evans, Brenier, Benamou, Caffarelli, Figalli, and many many more,

- Related to linear programming, manifold learning, image processing, game theory, …

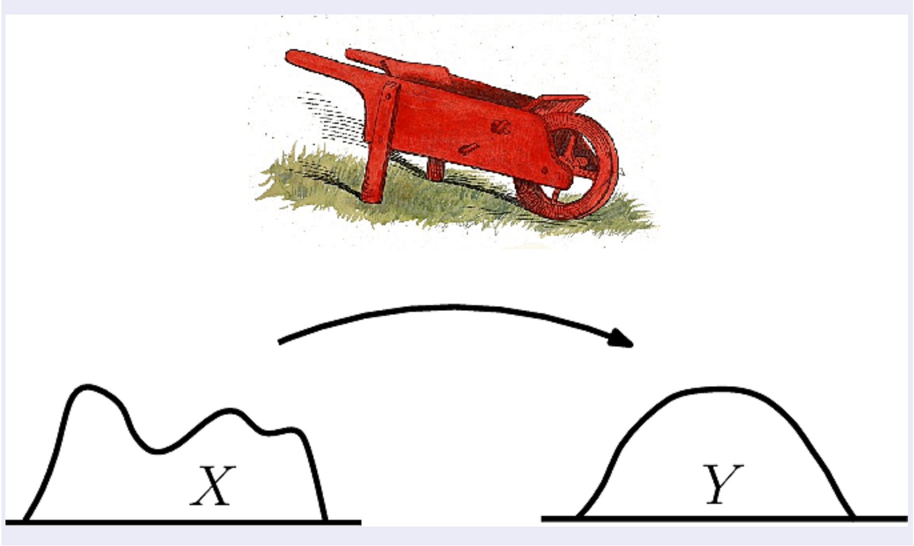

Benamou-Brenier Formula

$$W_2(\rho^0, \rho^1) = \inf_v \left(\int_0^1 \int_{\mathbb{R}^N} v(t,x)^2 \rho(t,x)\, dx dt \right)^{\frac{1}{2}}$$

s.t.

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (v\rho) = 0, \quad \rho(0,x) = \rho^0, \quad \rho(1,x) = \rho^1.$$

2-Wasserstein distance: the minimal cost, with respect to the square of Euclidean distance, to move from ρ^0 to ρ^1.

Fokker-Planck Equations

- Randomly perturbed gradient system:

$$dx = -\nabla\Psi(x)dt + \sqrt{2\beta}dW_t, \quad x \in \mathbb{R}^N$$

- Time evolution of the probability density function, the Fokker-Planck equation:

$$\rho_t(x,t) = \nabla \cdot (\nabla\Psi(x)\rho(x,t)) + \beta\Delta\rho(x,t)$$

- Invariant distribution at steady state -- Gibbs distribution:

$$\rho^*(x) = \frac{1}{K}e^{-\Psi(x)/\beta} \qquad K = \int_{\mathbb{R}^N} e^{-\Psi(x)/\beta}\,\mathrm{d}x$$

Free Energy and Fokker-Planck Equations

- Free energy $$F(\rho) = U(\rho) - \beta S(\rho)$$

- Potential $$U(\rho) = \int_{\mathbb{R}^N} \Psi(x)\rho(x)\mathrm{d}x$$

- Gibbs-Boltzmann Entropy
$$S(\rho) = -\int_{\mathbb{R}^N} \rho(x)\log\rho(x)\mathrm{d}x$$

- Fokker-Planck equation is the gradient flow of the free energy under 2-Wasserstein metric on the manifold of probability space.

- Gibbs distribution is the global attractor of the gradient system.

Optimal Transport

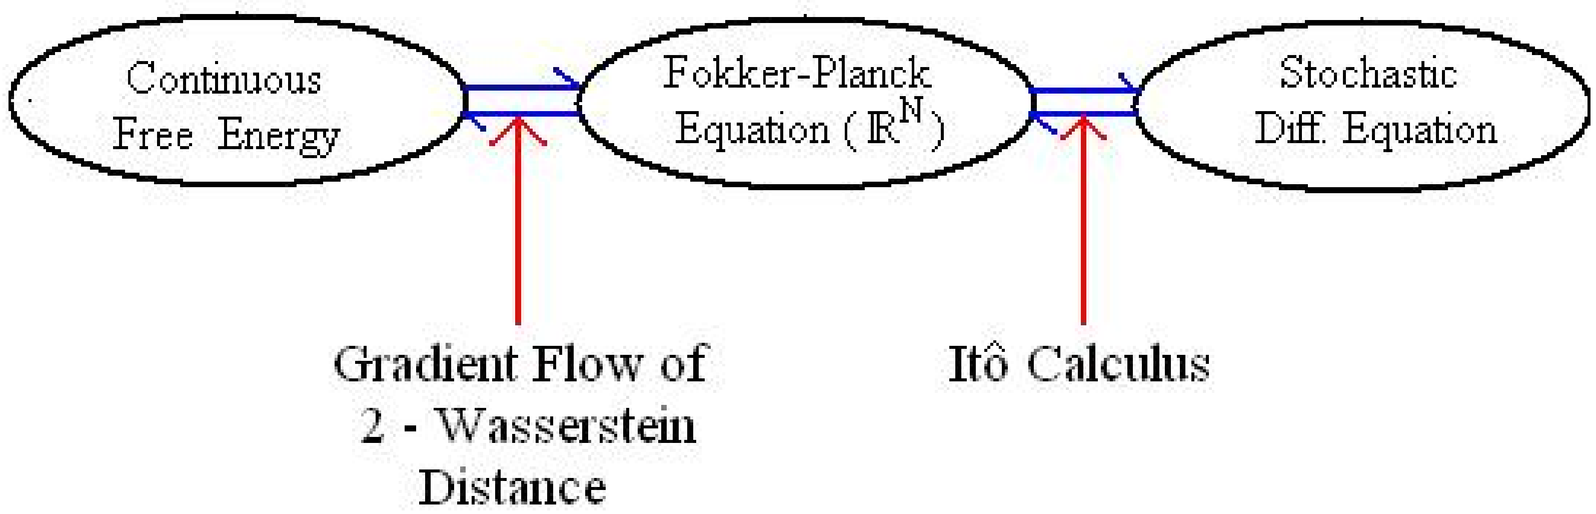

$$F(\rho) = U(\rho) - \beta S(\rho)$$

Free Energy

$$dx = -\nabla \Psi(x)dt + \sqrt{2\beta}dW_t$$

SDE

$$\rho_t = \nabla \cdot (\nabla \Psi \rho) + \beta \Delta \rho = \nabla \cdot (\nabla(\Psi + \beta \log \rho)\rho)$$

Fokker-Planck Equation

Optimal Transport on Finite graphs

- Our Goal: establish optimal transport on graphs with finite vertices.
- Why on graphs: Physical space (number of sites or states) is finite, not necessary from a spatial discretization such as a lattice.
- Applications: game theory, RNA folding, logistic, chemical reactions, machine learning, Markov networks, numerical schemes, ...
- Mathematics: Graph theory, Mass transport, Dynamical systems, Stochastic Processes, PDE's, ...
- Many Recent Developments: Erbar, Mielke, Mass, Gigli, Ollivier, Villani, Tetali, Fathi, Qian, Carlen, …

Basic Set-ups

Graph with finite vertices

$$G = (V, E), \quad V = \{1, \cdots, n\}, \quad \text{E is the edge set.}$$

Probability set

$$\mathcal{P}(G) = \{(\rho_i)_{i=1}^n \mid \sum_{i=1}^n \rho_i = 1,\ \rho_i \geq 0\}.$$

Discrete free energy

$$\mathcal{F}(\rho) = \frac{1}{2}\sum_{i=1}^n \sum_{j=1}^n w_{ij}\rho_i\rho_j + \sum_{i=1}^n v_i\rho_i + \beta\sum_{i=1}^n \rho_i \log \rho_i.$$

Boltzmann-Shannon entropy

Challenges

- Common discretizations of continuous fokker-planck equations often lead to incorrect results,

Theorem: Any given linear discretization of the continuous equation can be written as

$$\frac{d\rho_i}{dt} = \sum_j ((\sum_k e^i_{jk}\Phi_k) + c^i_j)\rho_j.$$

Let

$$A = \{\Phi \in \mathbb{R}^N : \sum_j ((\sum_k e^i_{jk}\Phi_k) + c^i_j)e^{-\frac{\Phi_j}{\beta}} = 0\}.$$

Then A is a zero measure set.

- Graphs are not length spaces and many of the essential techniques cannot be used anymore,

- The notion of random perturbation (white noise) of a Markov process on discrete spaces is not clear.

- Nodes on graphs may have very different neighborhood structures.

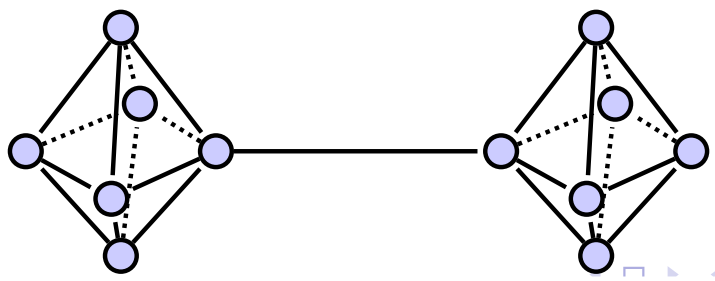

Optimal Transport on Graphs

Discrete 2-Wasserstein distance

For any $\rho^0, \rho^1 \in \mathcal{P}(G)$, define

$$W_{2;\mathcal{F}}(\rho^0, \rho^1) = \inf_v \left(\int_0^1 (v, v)_\rho dt \right)^{\frac{1}{2}}$$

where v and ρ satisfy

$$\frac{d\rho}{dt} + div_G(\rho v) = 0, \quad \rho(0, x) = \rho^0, \rho(1, x) = \rho^1.$$

Vector Operators on Graphs

Vector field on a graph: $v = (v_{ij})_{(i,j)\in E}$, satisfying $v_{ij} = -v_{ji}$

Potential Φ induced vector field : $\nabla_G = (\Phi_i - \Phi_j)_{(i,j)\in E}$

Divergence w. r. t. ρ

$$div_G(\rho v) = -(\sum_{j\in N(i)} v_{ij} g_{ij}^F(\rho))_{i=1}^n$$

Inner product

$$(v,u)_\rho = \sum_{(i,j)\in E} v_{ij} u_{ij} g_{ij}^F(\rho)$$

Here $F_i(\rho) = \frac{\partial}{\partial \rho_i}\mathcal{F}(\rho)$ and

$$g_{ij}^F(\rho) = \begin{cases} \rho_i & \text{if } F_i(\rho) > F_j(\rho),\ j\in N(i); \\ \rho_j & \text{if } F_i(\rho) < F_j(\rho),\ j\in N(i); \\ \frac{\rho_i+\rho_j}{2} & \text{if } F_i(\rho) = F_j(\rho),\ j\in N(i). \end{cases}$$

Discrete Fokker-Planck Equations

Theorem 1

For a finite graph $G = (V, E)$ and a constant $\beta > 0$. The gradient flow of discrete free energy

$$\mathcal{F}(\rho) = \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n} w_{ij}\rho_i\rho_j + \sum_{i=1}^{n} v_i\rho_i + \beta\sum_{i=1}^{n}\rho_i \log \rho_i$$

on the metric space $(\mathcal{P}_o(G), W_{2;\mathcal{F}})$ is

$$\frac{d\rho_i}{dt} = \sum_{j\in N(i)} \rho_j(F_j(\rho) - F_i(\rho))_+ - \sum_{j\in N(i)} \rho_i(F_i(\rho) - F_j(\rho))_+ \qquad (1)$$

for any $i \in V$. Here $F_i(\rho) = \frac{\partial}{\partial \rho_i}\mathcal{F}(\rho)$ and $(\cdot)_+ = \max\{\cdot, 0\}$.

Continuous Fokker-Planck equation

$$\rho_t = \nabla \cdot (\nabla\Psi\rho) + \beta\Delta\rho = \nabla \cdot (\nabla(\Psi + \beta \log \rho)\rho)$$

Fokker-Planck Equation and Exploring Region

In the path planning case, the region is determined by,

$$\frac{\partial \rho_j}{\partial t} = \Big(\sum_{k \in Nb(j)} (F_k(\rho,\sigma) - F_j(\rho,\sigma))_+ \rho_k d_{jk} - \sum_{k \in Nb(j)} (F_j(\rho,\sigma) - F_k(\rho,\sigma))_+ \rho_j d_{jk} \Big) \frac{1}{(\Delta x)^2},$$

where

$$F_i(\rho,\sigma) = p(i) + \sigma \log \rho_i.$$

$p(i)$ is the distance to the target.

$\sigma = 0$ projected gradient,
$\sigma > 0$ diffusion.

Distance value is not used, only the projected gradient direction is used,

Fokker-Planck Equation and Exploring Region

A different view point:

On a generated tree, Fokker-Planck equation selects nodes

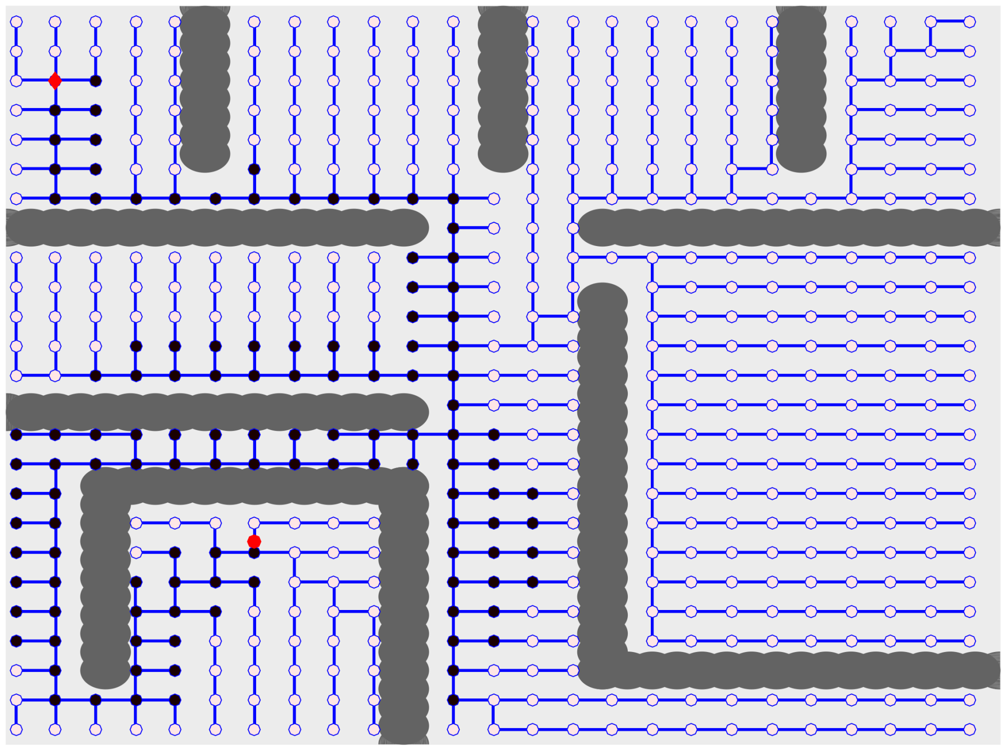

(a) The Evolved Region

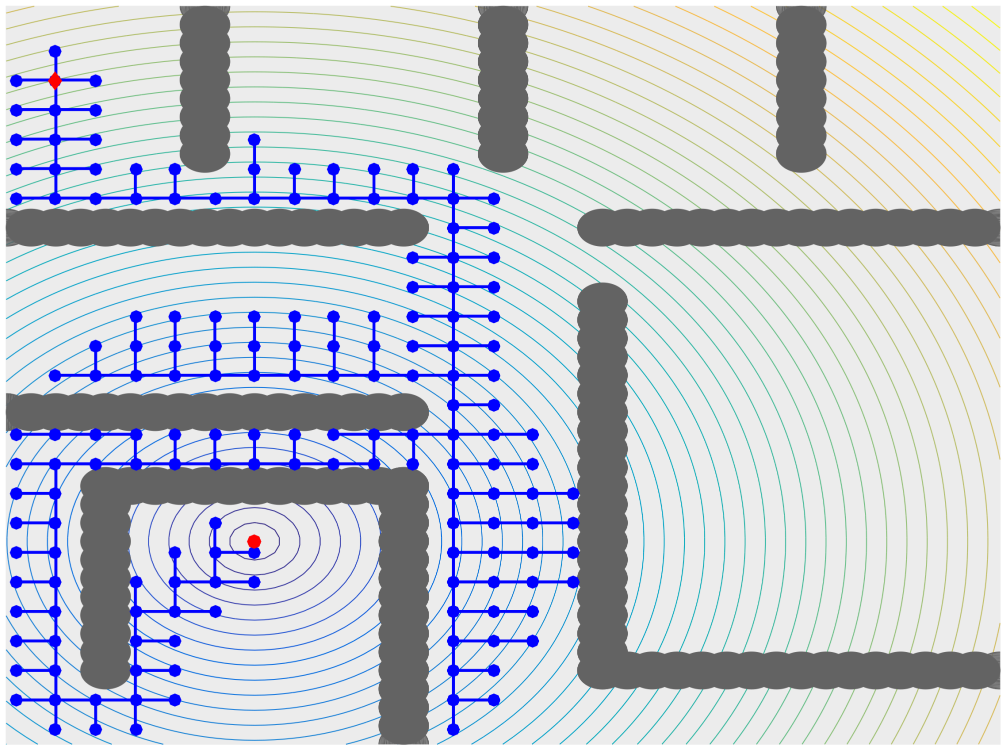

(b) The Generated Graph

General Control Problems

For a complete control system $(\Omega, \mathcal{U}, f)$ that is completely controllable in time T, with $\mathcal{U}$ being a compact set of $\mathbb{R}^n$ and f being a Lipschitz function, we let (Ω, g) be a d dimensional compact Riemannian manifold. There exists a distance function $d(\cdot,\cdot)$ induced by $g(\cdot,\cdot)$. We want to find $u \in \mathcal{U}^{[0,T)}$ such that

$$
\begin{aligned}
&\dot{\gamma}(t) = f(\gamma(t), u(t)),\\
&\gamma(0) = x_0, \gamma(T) = x_f,\\
&\phi(\gamma(t)) \geq 0 \text{ for all } t \in [0, T],\\
&\hat{\psi}(\gamma(t), \gamma, t) \geq 0 \text{ for all } t \in [0, T],
\end{aligned}
$$

where

$$
\hat{\psi}(x, t, \gamma) = \begin{cases} \psi(x) & \text{if } d(x, \gamma(\tau)) \leq R \text{ for some } \tau \leq t \\ 0 & \text{otherwise} \end{cases}
$$

An Example, Path Planning on a Sphere

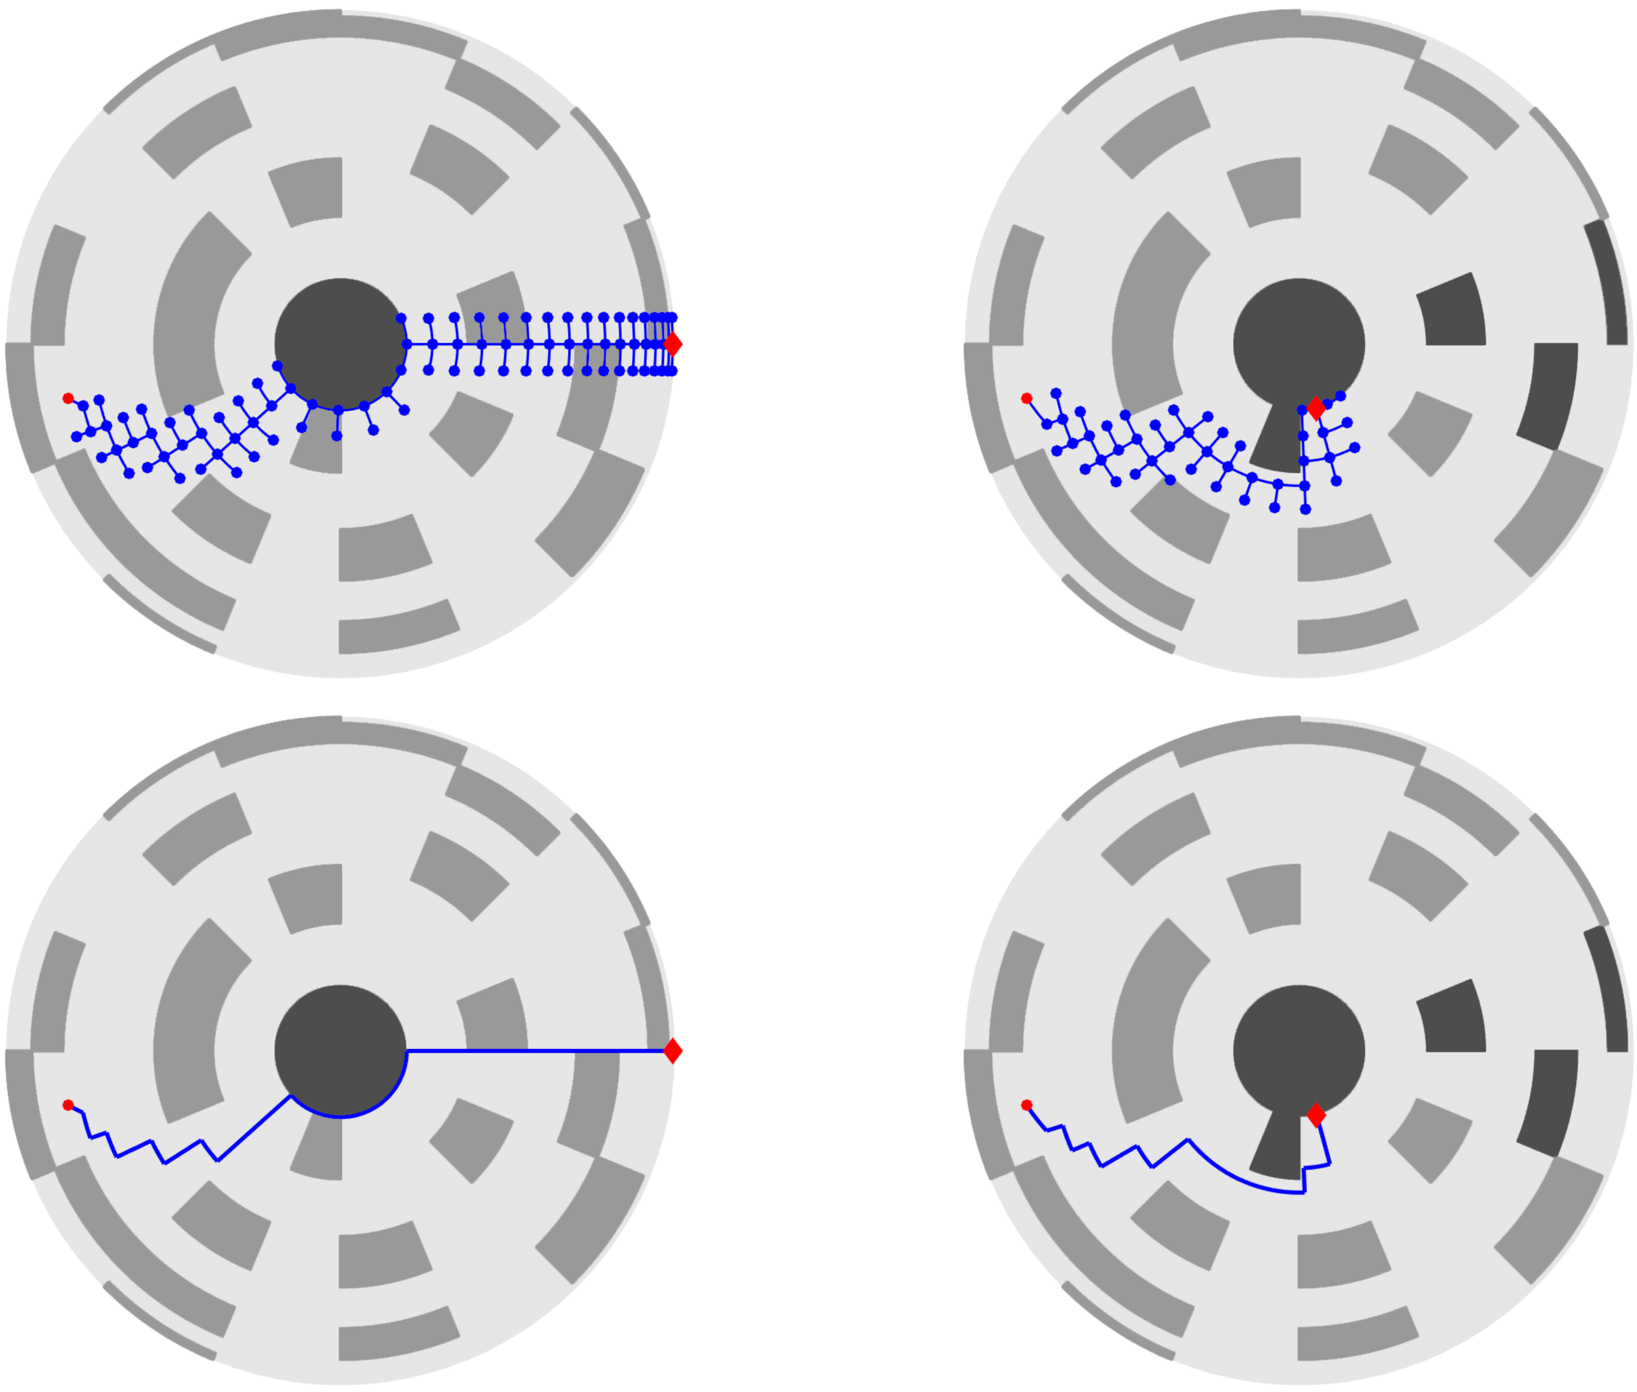

Thank you for your attention!