

Multi-Agent Shape Control with Optimal Transport

Alex Tong Lin, Stanley J. Osher

June 30, 2022

Abstract

We introduce a method called MASCOT (Multi-Agent Shape Control with Optimal Transport) to compute optimal control solutions of agents with shape/formation/density constraints. For example, we might want to apply shape constraints on the agents – perhaps we desire the agents to hold a particular shape along the path, or we want agents to spread out in order to minimize collisions. We might also want a proportion of agents to move to one destination, while the other agents move to another, and to do this in the optimal way, i.e. the source-destination assignments should be optimal. In order to achieve this, we utilize the Earth Mover’s Distance from Optimal Transport to distribute the agents into their proper positions so that certain shapes can be satisfied. This cost is both introduced in the terminal cost and in the running cost of the optimal control problem.

1 Introduction

Optimal control seeks to find the best policy for an agent that optimizes a certain criterion. This general formulation allows optimal control theory to be applied in numerous areas such as robotics, finance, aeronautics, and many other fields. Inherently, optimal control optimizes the control of a single agent, but in recent years, extending optimal control problems to the realm of multi-agents has been a popular trend. Indeed, there are numerous cases where we want to model not just a single agent, but many, e.g. a fleet of drones.

Here we introduce MASCOT: Multi-Agent Shape Control with Optimal Transport, a method to compute solutions to multi-agent optimal control problems that involve shape, formation, or density constraints among the agents. These constraints can be formulated in the running cost of the agents, or as a terminal cost, or even both.

We first introduce the reader to optimal control and its multi-agent version. We then review the idea of optimal transport and Earth Mover’s Distance. Finally, we demonstrate the method on some examples.

2 Related Works

In terms of multi-agent shape/formation constraints, as noted in [29], some methods rely on the use of leaders [5, 6, 15, 28, 10] to help agents make formations. Others are more behavior-based [2, 3, 20]. And there are some that rely on a virtual-structure approach [16]. Methods based on consensus of agents has also been considered in [13, 26, 29, 22]. There is also work based on potential functions [30, 9]. There are also many algorithms for solving multi-agent optimal control problems based on different assumptions on how the multi-agents interact: [23, 25, 8, 17].

In terms of using the assignment problem, [12] used various assignment algorithms to match agents with destinations so as to improve surveillance. In [21], they experiment with solving the assignment problem in order make agent formations, but do not place it into an optimal control framework. Recently, in [11] they apply discrete optimal transport in a capability-aware fashion to assign agents to moving targets. A survey of multi-agent formation control is provided in [24].

Further, related to the field of multi-agent optimal control are the field of Mean Field Control and Mean Field Games [19, 27, 14].

3 Background

We provide background on optimal control, multi-agent optimal control, and optimal transport.

3.1 Optimal Control

Optimal control seeks to find a control law that best optimizes a cost or payoff criterion. Here we will stick with convention of minimizing a cost.

Given an initial point $x \in \mathbb{R}^n$ and an initial time $t \in [0, T]$, the system will follow the dynamics:

$$\begin{cases} \dot{\mathbf{x}}(s) = \mathbf{f}(\mathbf{x}(s), \mathbf{u}(s)), & t < s < T \\ \mathbf{x}(t) = x \end{cases}$$

where $x, \mathbf{x}(s) \in \mathbb{R}^n$ for all $s \in (0, T)$, $\mathbf{f}: (\mathbb{R}^n \times U) \rightarrow \mathbb{R}^n$, and $U \subseteq \mathbb{R}^m$. We call \mathbf{x} the *state*, and \mathbf{u} the *control*. Then we want to minimize the functional $J_{x,t}: \mathcal{U} \rightarrow \mathbb{R}$ where,

$$J_{x,t}[\mathbf{u}] \stackrel{\text{def}}{=} g(\mathbf{x}(T)) + \int_t^T L(\mathbf{x}(s), \mathbf{u}(s)) ds$$

and where $\mathcal{U} \stackrel{\text{def}}{=} \{\mathbf{u}: \mathbf{u}: (0, T) \rightarrow U\}$ is called the *admissible control set*, $g: \mathbb{R}^n \rightarrow \mathbb{R}$ is the *terminal cost* and $L: (\mathbb{R}^n \times U) \rightarrow \mathbb{R}$ is the *running cost*. So optimal control seeks to find $\mathbf{u} \in \mathcal{U}$ that minimizes $J_{x,t}$.

3.2 Multi-Agent Optimal Control

There are many formulations of multi-agent optimal control, based on whether the control is centralized or decentralized, or whether the agents communicate or not, and many other factors. In this work, we consider a simple extension of optimal control with a centralized controller and a finite number of agents. Later, the control can perhaps be decentralized with an imitation learning algorithm as demonstrated in [18].

In this work, we consider the following multi-agent control problem where the dynamics modeling the N agents are:

$$\begin{cases} \dot{\mathbf{x}}^{(i)}(s) = \mathbf{f}(\mathbf{x}^{(i)}(s), \mathbf{u}^{(i)}(s)), & t < s < T, \quad 1 \leq i \leq N \\ \mathbf{x}^{(i)}(t) = \mathbf{x}_0^{(i)} \end{cases}$$

and they want to collectively minimize the following cost functional:

$$J[\{\mathbf{u}^{(i)}\}_{i=1}^N] \stackrel{\text{def}}{=} g(\{\mathbf{x}^{(i)}(T)\}_{i=1}^N) + \int_t^T L(\{\mathbf{x}^{(i)}(s)\}_{i=1}^N, \{\mathbf{u}^{(i)}(s)\}_{i=1}^N) ds.$$

We note that technically, if we stack the states $\{\mathbf{x}^{(i)}\}$ into one concatenated vector, and we stack the controls $\{\mathbf{u}^{(i)}\}$ also into one concatenated vector, then this can be viewed as a single-agent control problem. This would model the realistic situation where a fleet of drones are being controlled by a centralized controller, for example.

3.3 Earth Mover's Distance and Optimal Transport

The problem of Optimal Transport seeks to find a transportation plan between two probability distributions that is optimal. The cost of the plan also provides a distance metric between the probability distributions, called *Earth Mover's Distance* (also called the *Wasserstein distance*).

To aid explanation and because this is the most relevant case to us, we restrict ourselves to discrete distributions: Suppose we are given two sets of points: $\{\mathbf{x}^{(i)}\}_{i=1}^N$ and $\{\mathbf{z}^{(i)}\}_{i=1}^M$, and we

weight them with distributions $\mathbf{a} = (a_1, \dots, a_N) \in \mathbb{R}^N$ and $\mathbf{b} = (b_1, \dots, b_M) \in \mathbb{R}^M$. So $\mathbf{x}^{(i)}$ has weight a_i for example. Then we can define a cost matrix,

$$C = (c_{ij}) \in \mathbb{R}^{N \times M}, \quad c_{ij} = \|a_i - b_j\|^2, 1 \leq i \leq N, 1 \leq j \leq M.$$

Then optimal transport seeks to find a transportation plan $\pi = (\pi_{ij}) \in \mathbb{R}^{N \times M}$ that minimizes the following,

$$\text{EMD}_C(\mathbf{a}, \mathbf{b}) = \min_{\pi} \left\{ \sum_{ij} \pi_{ij} c_{ij} : \sum_j \pi_{ij} = a_i, \text{ for all } i, \sum_i \pi_{ij} = b_j, \text{ for all } j \right\}$$

and value of the minimum is called the *Earth Mover's Distance*. The argmin is the transportation plan π .

In the special case where $N = M$ and \mathbf{a} and \mathbf{b} are the uniform distribution (i.e. $a_i = b_i = 1/N$ for all $1 \leq i \leq N$), then we can just set $\mathbf{a} = (1, \dots, 1)$ and $\mathbf{b} = (1, \dots, 1)$. Then our problem simply becomes an assignment problem. Although, optimal transport is actually much more general and we take advantage of this: we can demand distributional preferences. For example, we may want a proportion of agents to cover one destination and the rest to cover another destination. In this case, we may have N agents, but only 2 destination points, which significantly reduces the computation cost of the cost matrix from $O(N^2)$ (using the assignment algorithm), to $O(2N) = O(N)$ (using optimal transport).

Related to the computation of the Earth Mover's Distance, is the computation of a regularized Earth Mover's Distance utilizing Sinkhorn iteration [4] and variants therein, which are able to achieve faster computational times, sometimes even achieving near-linear time complexity [1].

4 Methods

In this section, we provide details on how to implement our method. We first explain the case of controlling the multi-agent shape/density at terminal time. Then we explain how to implement the method for the running cost.

4.1 Shape control at terminal time

As noted in Section 3.2, the agents' dynamics are,

$$\begin{cases} \dot{\mathbf{x}}^{(i)}(s) = \mathbf{f}(\mathbf{x}^{(i)}(s), \mathbf{u}^{(i)}(s)), & t < s < T, \quad 1 \leq i \leq N \\ \mathbf{x}^{(i)}(t) = \mathbf{x}^{(i)} \end{cases}$$

and they want to minimize the cost-functional,

$$J[\{\mathbf{u}^{(i)}\}_{i=1}^N] = g(\{\mathbf{x}^{(i)}(T)\}_{i=1}^N) + \int_t^T L(\{\mathbf{x}^{(i)}(s)\}_{i=1}^N, \{\mathbf{u}^{(i)}(s)\}_{i=1}^N) ds.$$

Suppose we want the agents to satisfy the shape/density of the reference points $\{\mathbf{z}^{(j)}\}_{j=1}^M$. We set the cost matrix to be,

$$C = (c_{ij}), \quad c_{ij} = d(\mathbf{x}^{(i)}(T), \mathbf{z}^{(j)})$$

where d is the cost between $\mathbf{x}^{(i)}$ and $\mathbf{z}^{(j)}$. For example, d can be the squared Euclidean distance. Then we choose a distribution on the $\{\mathbf{x}^{(i)}\}_{i=1}^N$, which we call $\mathbf{a} = (a_1, \dots, a_N)$. We also choose a distribution on the $\{\mathbf{z}^{(j)}\}_{j=1}^M$, which we call $\mathbf{b} = (b_1, \dots, b_M)$. Then we can add to the terminal cost,

$$g_{\text{shape}}(\mathbf{x}^{(1)}(T), \dots, \mathbf{x}^{(N)}(T)) = \text{EMD}_C(\mathbf{a}, \mathbf{b})$$

where EMD_C is the Earth Mover's Distance as presented in Section 3.3, with cost matrix C . The Earth Mover's Distance and the transportation plan can be computed using standard linear

programming, or one can use Sinkhorn iteration [4] to approximate the distance and transportation plan.

Intuitively, the transportation plan π weights the cost matrix C , so that agents can then use this information to find optimal assignments to the proper reference points.

4.2 Shape control in the running cost

In order to apply shape/density constraints along the trajectory of the agents, we first subtract the mean from the agents and the reference points:

$$\mathbf{y}^{(i)}(t) = \mathbf{x}^{(i)}(t) - \bar{\mathbf{x}}(t), \quad \text{and} \quad \mathbf{w}^{(j)}(t) = \mathbf{z}^{(j)} - \bar{\mathbf{z}},$$

where $\bar{\mathbf{x}}(t) = \frac{1}{N} \sum_{i=1}^N \mathbf{x}^{(i)}(t)$, and $\bar{\mathbf{z}}(t) = \frac{1}{M} \sum_{j=1}^M \mathbf{z}^{(j)}(t)$ are the mean. Then letting \mathbf{a} and \mathbf{b} be the weight distribution of the \mathbf{y} and \mathbf{w} , as before, then shape constraints can be applied with,

$$L_{\text{shape}}(\{\mathbf{x}^{(i)}(t)\}_{i=1}^N) = \text{EMD}_C(\mathbf{a}, \mathbf{b}),$$

where the cost matrix C is as before.

Subtracting the mean allows the agents to become agnostic to the actual position of the reference points, and now the agents can focus on the relative positioning amongst each other in order to satisfy the shape constraint. This also allows the optimal control problem itself to find optimal paths for the agents, because otherwise the actual positions of the reference points would interfere. One can even try to normalize the agents and reference points using bounding boxes, so agent formations are agnostic to the actual diameter of the reference points.

4.3 Classical Direct Shooting for Optimal Control

Our approach to inducing shape constraints on the agents is generally agnostic to the method used to compute optimal control problems. Thus in our numerical demonstration, we employ a very simple and straightforward method to compute optimal control solutions – the direct shooting method.

We first discretize the time domain:

$$0 = t_0 < t_1 < \dots < t_S = T.$$

For convenience, we let the discretization be uniform with mesh size Δt . Then denoting $\mathbf{x}_s^{(i)} \stackrel{\text{def}}{=} \mathbf{x}^{(i)}(t_s)$ and similarly with $\mathbf{u}_s^{(i)}$, if we use forward Euler approximations for the agents' dynamics, then we have,

$$\begin{cases} \mathbf{x}_{s+1}^{(i)} = \mathbf{x}_s^{(i)} + \Delta t \mathbf{f}(\mathbf{x}_s^{(i)}, \mathbf{u}_s^{(i)}), & t < s < T, \quad 1 \leq i \leq N \\ \mathbf{x}_0^{(i)} = \mathbf{x}^{(i)}(0) \end{cases} \quad (1)$$

Our cost function then becomes,

$$J[\{\mathbf{u}_s^{(i)}\}_{i=1, s=1}^{i=N, s=S}] = g(\{\mathbf{x}_S^{(i)}\}_{i=1}^N) + \Delta t \sum_{s=0}^{S-1} L(\{\mathbf{x}_s^{(i)}\}_{i=1}^N, \{\mathbf{u}_s^{(i)}\}_{i=1}^N).$$

Then the shooting method starts out by making an initial guess for $\{\mathbf{u}_s^{(i)}\}_{i=1, s=1}^{i=N, s=S}$, and then we integrate according to the agent dynamics (1). We then update each $\mathbf{u}_s^{(i)}$ by gradient descent on the cost function,

$$\mathbf{u}_s^{(i)} \leftarrow \mathbf{u}_s^{(i)} - \alpha \nabla_{\mathbf{u}_s^{(i)}} J$$

where α is the gradient descent step-size. We repeat this iteration until some convergence criterion has reached, e.g. when the controls $\{\mathbf{u}_s^{(i)}\}_{i=1, s=1}^{i=N, s=S}$ cease to stop changing very much between iterations.

5 Numerical Demonstrations

Here we provide numerical demonstrations of our method. We first introduce the particular agent dynamics we will use for all our experiments – the 2D double integrator. Then using the direct shooting method as mentioned in Section 4.3 we examine the cases of having shape constraints in the terminal constraint, then in the running cost. Afterwards, we examine the interplay between shape constraints and obstacle avoidance and congestion/collision minimization. To compute the Earth Mover’s Distance, we utilize the Python Optimal Transport (POT) toolbox [7].

5.1 Agent dynamics and the general cost function

In these numerical demonstrations, for each agent we will use the 2D double integrator:

$$\begin{cases} \ddot{x} = \alpha \\ \ddot{y} = \beta \\ x(0) = x_0, y(0) = y_0, \dot{x}(0) = 0, \dot{y}(0) = 0 \end{cases}$$

Of course we then turn this into the first-order system by letting $x_1 = x$, $x_2 = y$, $x_3 = \dot{x}$, $x_4 = \dot{y}$, so if $\mathbf{x} = (x_1, x_2, x_3, x_4)^\top$, and we let $\mathbf{u} = (\alpha, \beta)^\top$, so finally we get,

$$\begin{cases} \dot{\mathbf{x}}(t) = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \mathbf{x}(t) + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \mathbf{u}(t) \\ \mathbf{x}(0) = (x_0, y_0, 0, 0) \end{cases}$$

Our cost functional will generally have the form,

$$\begin{aligned} J[\{\mathbf{u}^{(i)}\}] &= g_{\text{shape}}(\{\mathbf{x}^{(i)}(T)\}_{i=1}^N) + \frac{1}{N} \sum_{j=1}^N \frac{1}{2} \|(x_3(T), x_4(T))\|^2 \\ &+ \int_0^T \sum_{j=1}^N \frac{1}{2} \|\mathbf{u}^{(j)}(s)\|^2 + L_{\text{shape}}(\{\mathbf{x}^{(i)}(t)\}_{i=1}^N) ds \end{aligned}$$

So in all cases, we want the agents to have near zero velocity at terminal time, and we are at least minimizing the squared norm of the control in the running cost.

5.2 Shape constraints at terminal time

In order to employ shape constraints at terminal time, we have reference points $\{\mathbf{z}^{(j)}\}_{j=1}^M$ and letting \mathbf{a} be the weights for the agents and \mathbf{b} be the weights for the reference points, then we have

$$g_{\text{shape}}(\{\mathbf{x}^{(i)}(T)\}_{i=1}^N) = \text{EMD}_C(\mathbf{a}, \mathbf{b})$$

for a given cost matrix $C = (c_{ij})$. Here, we just let the cost matrix be the squared Euclidean distance between agents and reference points, i.e. $c_{ij} = \|\mathbf{x}^{(i)} - \mathbf{z}^{(j)}\|^2$.

In Figure 1, we show that the agents are able to move to a “circle within a circle” shape/distribution. The agents at the starting time and terminal times are opaque and colored in blue and orange respectively. Intermediate times are transparent.

We can also enforce proportionality constraints on the agents. In Figure 2, we have the agents move to the destinations $(2, 1.5)$ and $(2, -1.5)$, but enforce that $2/5$ of the agents go to the upper destination, and the rest $3/5$ go to the lower destination. We note that in this case, the number of reference points $\{\mathbf{z}^{(j)}\}$ is 2, and we merely let $\mathbf{b} = (\frac{2}{5}, \frac{3}{5})$. This saves on computational cost as compared to the full assignment method as noted in Section 3.3.

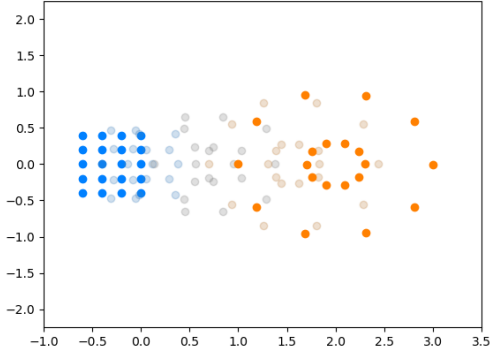


Figure 1 – Agents start from a rectangular shape/distribution and move to a “circle in a circle” shape/distribution. The agents at the starting and ending times are opaque and colored in blue and orange, whereas agents at intermediate times are plotted to be transparent.

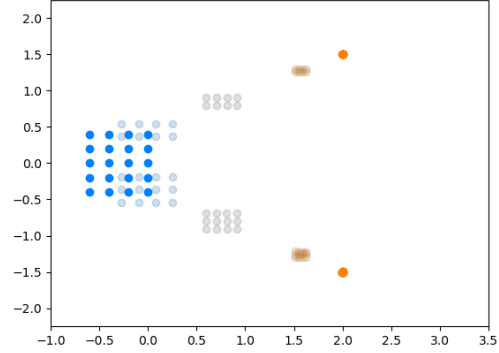


Figure 2 – Agents start from a rectangular shape/distribution and their destination is both $(2, 1.5)$ and $(2, -1.5)$. We also enforce that they divide themselves proportionally in a $2/5 - 3/5$ proportion.

5.3 Shape constraints in the running cost

We can also enforce the agents to adhere to shape constraints along the path for the agents. In the running cost, we add the term,

$$L_{\text{shape}}(\{\mathbf{x}^{(i)}(t)\}_{i=1}^N) = \text{EMD}_C(\mathbf{a}, \mathbf{b}).$$

In Figure 3 we make the agents maintain a “Flying V” shape along the path. In this case, in order to maintain the shape even at terminal time, we also apply the same shape constraint at terminal time, and we loosen the constraint so that it is merely the agent average that should reach the destination of $(2, 0)$.

We can also enforce proportionality constraints in the running cost as well. In Figure 4, we enforce the agents to perform a pincer maneuver, but to also split in a $(\frac{2}{5}, \frac{3}{5})$ proportion. Also, as in Figure 2, the number of reference points is 2 again, saving on computational costs.

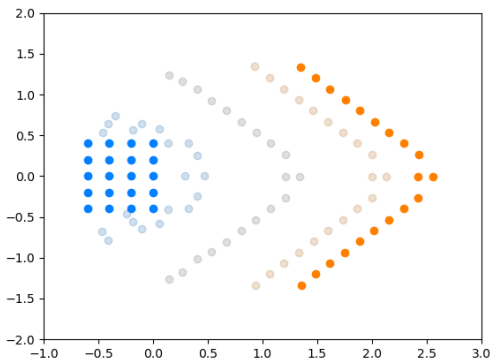


Figure 3 – Agents forming a “Flying V” along the path. This is a demonstration of applying shape constraints in the running cost.

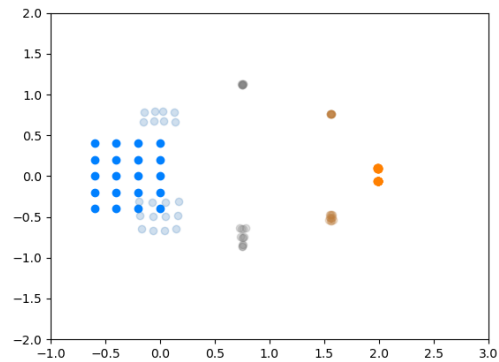


Figure 4 – Agents performing a pincer maneuver but also maintaining a $(\frac{2}{5}, \frac{3}{5})$ split along the path.

5.4 Congestion and Obstacles

Here we demonstrate that we can still enforce shape constraints while avoiding collisions and obstacles.

In Figure 5, we perform the same pincer maneuver as in Figure 4, but now we apply a congestion penalty. This is enforced by using the Gaussian kernel in the running cost:

$$\text{congestion penalty} = \sum_{ij} \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{x}^{(i)}(t) - \mathbf{x}^{(j)}(t)\|^2\right)$$

where we chose $\sigma = 0.15$. We see that the agents do indeed avoid colliding.

And in Figure 6, the agents form the same “Flying V” as in Figure 3, but now with congestion/collision penalty, and avoiding an obstacle. We note that compared to the case without an obstacle, the “tip” of the “Flying V” differ. As the agents move towards their destination, we can see a dynamic reassigning of roles – initially the “tip” of “Flying V” is the middle agent in the right-furthest column, as can be seen in Figure 3. But with an obstacle, the agents dynamically re-assign themselves, so now the role of the “tip” is a different agent.

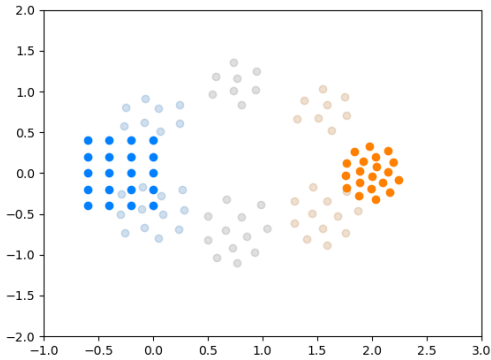


Figure 5 – Agents performing a pincer maneuver but now with a congestion/collision penalty, but also maintaining a $(\frac{2}{5}, \frac{3}{5})$ split.

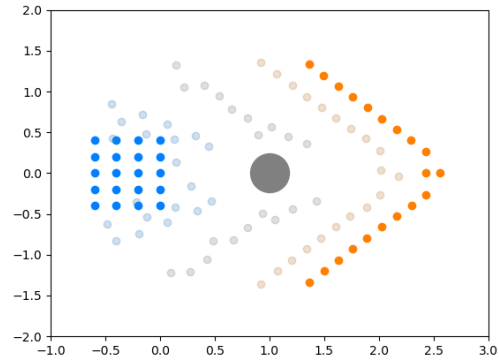


Figure 6 – Agents forming a “Flying V” formation, but now with a congestion/collision penalty, and avoiding an obstacle.

6 Conclusion

In this work, we demonstrate that with MASCOT using Optimal Transport and the Earth Mover’s Distance, we can apply shape constraints to enforce shape/formation/density constraints on agents. We provide numerical examples and demonstrate that dynamic re-assigning of roles can take place.

7 Acknowledgments

Alex Tong Lin and Stanley J. Osher were supported by the grants: AFOSR MURI FA9550-18-1-0502, and ONR grants: N00014-20-1-2093, N00014-20-1-2787, and also by DOD-AF-AIR FORCE RESEARCH LABORATORY (AFRL) Sponsor Award Number: 21-EPA-RQ-50:002, and Alex Tong Lin was further supported by Air Fare AWARD 16-EPA-RQ-09.

References

- [1] Jason Altschuler, Jonathan Niles-Weed, and Philippe Rigollet. Near-linear time approximation algorithms for optimal transport via sinkhorn iteration. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [2] Ronald C Arkin, Ronald C Arkin, et al. *Behavior-based robotics*. MIT press, 1998.
- [3] Tucker Balch and Ronald C Arkin. Behavior-based formation control for multirobot teams. *IEEE transactions on robotics and automation*, 14(6):926–939, 1998.
- [4] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.
- [5] Magnus Egerstedt and Xiaoming Hu. Formation constrained multi-agent control. *IEEE transactions on robotics and automation*, 17(6):947–951, 2001.
- [6] Magnus Egerstedt, Xiaoming Hu, and Alexander Stotsky. Control of mobile platforms using a virtual vehicle approach. *IEEE transactions on automatic control*, 46(11):1777–1782, 2001.
- [7] Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z. Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, Léo Gautheron, Nathalie T.H. Gayraud, Hicham Janati, Alain Rakotomamonjy, Ievgen Redko, Antoine Rolet, Antony Schutz, Vivien Seguy, Danica J. Sutherland, Romain Tavenard, Alexander Tong, and Titouan Vayer. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021.
- [8] Greg Foderaro, Silvia Ferrari, and Thomas A Wettergren. Distributed optimal control for multi-agent trajectory optimization. *Automatica*, 50(1):149–154, 2014.
- [9] Kristian Hengster-Movrić, Stjepan Bogdan, and Ivica Draganjac. Multi-agent formation control based on bell-shaped potential functions. *Journal of Intelligent and Robotic Systems*, 58(2):165–189, 2010.
- [10] Meng Ji, A. Muhammad, and M. Egerstedt. Leader-based multi-agent coordination: controllability and optimal control. In *2006 American Control Conference*, pages 6 pp.–, 2006.
- [11] Koray Gabriel Kachar and Alex Arkady Gorodetsky. Dynamic multiagent assignment via discrete optimal transport. *IEEE Transactions on Control of Network Systems*, 9(1):151–162, 2022.
- [12] Kwan S Kwok, Brian J Driessen, Cynthia A Phillips, and Craig A Tovey. Analyzing the multiple-target-multiple-agent scenario using optimal assignment algorithms. *Journal of Intelligent and Robotic Systems*, 35(1):111–122, September 2002.
- [13] Gerardo Lafferriere, Alan Williams, J Caughman, and JJP Veerman. Decentralized control of vehicle formations. *Systems & control letters*, 54(9):899–910, 2005.
- [14] Mathieu Laurière, Sarah Perrin, Matthieu Geist, and Olivier Pietquin. Learning mean field games: A survey. *arXiv preprint arXiv:2205.12944*, 2022.
- [15] Naomi Ehrich Leonard and Edward Fiorelli. Virtual leaders, artificial potentials and coordinated control of groups. In *Proceedings of the 40th IEEE Conference on Decision and Control (Cat. No. 01CH37228)*, volume 3, pages 2968–2973. IEEE, 2001.
- [16] M Anthony Lewis and Kar-Han Tan. High precision formation control of mobile robots using virtual structures. *Autonomous robots*, 4(4):387–403, 1997.

-
- [17] Alex Tong Lin, Yat Tin Chow, and Stanley J Osher. A splitting method for overcoming the curse of dimensionality in hamilton–jacobi equations arising from nonlinear optimal control and differential games with applications to trajectory generation. *Communications in Mathematical Sciences*, 16(7), 2018.
- [18] Alex Tong Lin, Mark DeBord, Katia Estabridis, Gary Hewer, Guido Montufar, and Stanley Osher. Decentralized multi-agents by imitation of a centralized controller. In *Mathematical and Scientific Machine Learning*, pages 619–651. PMLR, 2022.
- [19] Alex Tong Lin, Samy Wu Fung, Wuchen Li, Levon Nurbekyan, and Stanley J. Osher. Alternating the population and control neural networks to solve high-dimensional stochastic mean-field games. *Proceedings of the National Academy of Sciences*, 118(31):e2024713118, 2021.
- [20] Vladimir J. Lumelsky and KR Harinarayan. Decentralized motion planning for multiple mobile robots: The cocktail party model. *Autonomous Robots*, 4(1):121–135, 1997.
- [21] Edward A Macdonald. *Multi-robot assignment and formation control*. PhD thesis, Georgia Institute of Technology, 2011.
- [22] Arindam Mondal, Laxmidhar Behera, Soumya Ranjan Sahoo, and Anupam Shukla. A novel multi-agent formation control law with collision avoidance. *IEEE/CAA Journal of Automatica Sinica*, 4(3):558–568, 2017.
- [23] Kristian Hengster Movric and Frank L Lewis. Cooperative optimal control for multi-agent systems on directed graph topologies. *IEEE Transactions on Automatic Control*, 59(3):769–774, 2013.
- [24] Kwang-Kyo Oh, Myoung-Chul Park, and Hyo-Sung Ahn. A survey of multi-agent formation control. *Automatica*, 53:424–440, 2015.
- [25] Derek Onken, Levon Nurbekyan, Xingjian Li, Samy Wu Fung, Stanley Osher, and Lars Ruthotto. A neural network approach applied to multi-agent optimal control. In *2021 European Control Conference (ECC)*, pages 1036–1041. IEEE, 2021.
- [26] Wei Ren. Consensus based formation control strategies for multi-vehicle systems. In *2006 American Control Conference*, pages 6–pp. IEEE, 2006.
- [27] Lars Ruthotto, Stanley J. Osher, Wuchen Li, Levon Nurbekyan, and Samy Wu Fung. A machine learning framework for solving high-dimensional mean field game and mean field control problems. *Proceedings of the National Academy of Sciences*, 117(17):9183–9193, 2020.
- [28] J Shao, G Xie, and L Wang. Leader-following formation control of multiple mobile vehicles. *IET Control Theory & Applications*, 1(2):545–552, 2007.
- [29] Feng Xiao, Long Wang, Jie Chen, and Yanping Gao. Finite-time formation control for multi-agent systems. *Automatica*, 45(11):2605–2611, 2009.
- [30] Michael M Zavlanos and George J Pappas. Sensor-based dynamic assignment in distributed motion planning. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, pages 3333–3338. IEEE, 2007.