Decoding Complex UAV Swarm Maneuvers under Dynamical Leadership

Christos N. Mavridis, Nilesh Suriyarachchi and John S. Baras
Department of Electrical and Computer Engineering and the Institute for Systems Research, University of Maryland, College Park, USA.

Abstract

The advent and proliferation of the use of Unmanned Aerial Vehicles (UAV's) poses new challenges to swarm robotics and air defense systems. In this work we consider the problem of defending against adversarial attacks from UAV swarms performing complex maneuvers, driven by multiple, dynamically changing, leaders. We rely on short-time observations of the trajectories of the UAVs and develop a leader detection scheme based on the notion of Granger causality. We proceed with the estimation of the swarm's coordination laws, modeled by a generalized Cucker-Smale model with non-local repulsive potential functions and dynamically changing leaders, through an appropriately defined iterative optimization algorithm. Similar problems exist in communication and computer networks, as well as social networks over the Internet. Thus, the methodology and algorithms proposed can be applied to many types of network swarms including detection of influential malevolent "sources" of attacks and "miss-information". The proposed algorithms are robust to missing data and noise.

Introduction and Problem Definition

While modern high-precision targeting anti-air defenses are capable of taking down a single UAV, when it comes to a large swarm of UAV's attacking simultaneously, these defences can be rendered ineffective. These problems are even more challenging when in UAV swarms, a few units are managed by humans (we refer to them as "leaders", while most units "follow", leading to very effective management of large swarms. When the role of leaders can be dynamically re-assigned the monitoring and defense against such swarms becomes even more difficult.

The first question that needs to be addressed in creating a defense against a hostile UAV swarm is understanding the control (coordination) and communication laws governing how the drones move and interact with each other. We view the interconnected problems of modeling and learning the interaction laws of a swarm as one problem that can be analyzed in the microscopic scale as a port-Hamiltonian networked system. We extend existing simulation models, such as the Boids and the Cucker-Smale models, to incorporate interaction, communication and dynamics terms that can capture realistic complex swarm maneuvers and develop corresponding simulation models in the macroscopic domain [8, 10]. Consider an interacting system of N particles and the leader sets $\mathcal{L}(i),\ \ 1 \leq i \leq N$ of cardinality $|\mathcal{L}(i)| = 1$ assigned to each particle representing the index of the leader particle that it is following. We introduce

I. a scalable simulation algorithm, based on the Boids model [11], that can capture interaction laws and communication protocols of complex swarm maneuvers, including (a) velocity alignment, (b) spatial cohesion, (c) collision avoidance, and (d) response to dynamically changing leaders:

$$\begin{cases} \dot{x}_i &= v_i \\ \dot{v}_i &= -c\nabla U_c(x) - a\nabla U_a(x,v) + s\nabla U_s(x) \end{cases} \tag{1}$$

II. a large-scale learning algorithm, based on the generalized Cucker-Smale model [1] and automatic differentiation, designed to work on state-of-the-art deep learning platforms that can identify the interaction laws (a)-(d) by observing particle trajectories of position and velocity:

$$\begin{cases} \dot{x}_i &= v_i \\ \dot{v}_i &= \frac{K}{N}\sum_{j=1}^{N} \psi_{ij}(x(t), v(t)) \end{cases} \tag{2}$$

where

$$\psi_{ij}(x) = \begin{cases} -\nabla U(\|x_i - x_j\|), & j \notin \mathcal{L}(i), j \neq i \\ G(\|x_i - x_j\|)(v_j(t) - v_i(t)) - \nabla U(\|x_i - x_j\|), & j \in \mathcal{L}(i) \end{cases} \tag{3}$$

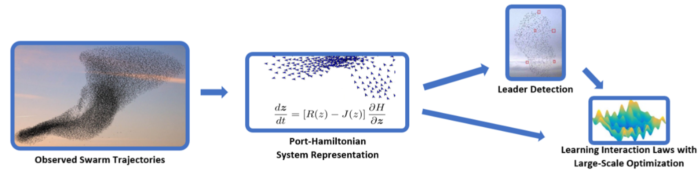

Figure 1: Reconstructing complex swarm dynamics. The agents' trajectories are observed and used to detect leaders and identify a port-Hamiltonian networked system modeling their interaction rules [8].

It can be shown that the Cucker-Smale model with leadership is equivalent to a fully connected N-dimensional network of generalized mass-spring-dampers with appropriately defined Hamiltonian functions, that can be written in an input-state-output port-Hamiltonian form [5, 9]:

$$\dot{\mathbf{z}} = [\mathbf{J}(\mathbf{z}) - \mathbf{R}(\mathbf{z})]\frac{\partial \mathbf{H}(\mathbf{z})}{\partial \mathbf{z}} + \mathbf{g}(\mathbf{z})\mathbf{u}, \tag{4}$$

where $z = (q, p)$, with $q, p \in \mathbb{R}^{\frac{N(N-1)}{2}}$ being the vectors of relative distances and momenta between each pair of particles, and the quantities $J = -J^{\mathrm{T}}$, H, R and $g(z)$ are appropriately defined, and u is an external control input that affects only the leader particles and is responsible for their trajectories.

Methodology

For leader detection, we adopt a majority vote criterion, where each particle i votes for the particle j to be the leader, according to a measure related to the observed trajectories of the particles. In particular, we each vote wit the existence of a causal effect of the trajectory of particle j, Y, on the trajectory of particle j, X. The causal relationship is determined by a *hypothesis test* designed according to Granger's definition of causality [2] which is based on two principles:

i. The cause happens prior to its effect
ii. The cause has unique information about the future values of its effect.

Given these assumptions, we say that a time series Y Granger-causes X if the past values of Y provide statistically significant information about the future values of X.

The first question one needs to answer when dealing with leader detection, however, is the number of leaders that the algorithm is trying to find. We view this problem as a clustering problem given a window of position and velocity observations of the particles, since it is reasonable to assume that particle trajectories will be more 'similar' to each other if they are following the same leader. However, the number of the clusters is not known a priori, which makes standard clustering algorithms inadequate for this application [7]. Instead, we need an unsupervised learning algorithm that progressively estimates the number of clusters by adding new clusters only when some measure of distortion is high enough to support this decision. In this regard, we use the Online Deterministic Annealing algorithm [6] as a fitting clustering algorithm for estimating the number of leaders.

For the learning task we model the networked system of interacting agents as a port-Hamiltonian system (4). We make use of the position

Extended abstract submitted to the Maryland Robotics Center (MRC) Research Symposium 2021. Research partially supported by the Defense Advanced Research Projects Agency (DARPA) under Agreement No. HR00111990027, by ONR grant N00014- 17-1-2622, and by a grant from Northrop Grumman Corporation. Authors' e-mails: {mavridis, nileshs, baras}@umd.edu.

and velocity trajectories of the particles to recover the resistive terms $R(z)$ and the Hamiltonian $H(z)$, which is equivalent to recovering the interaction functions $\psi_{ij}(x,v)$ of a general Cucker-Smale model (2). The components of the interaction model (resistive element and the spring Hamiltonian) are modeled as neural-networks with one hidden layer, and the following optimization problem with a mean square error (MSE) loss function is formulated

$$\min_w \frac{1}{n}\sum_{i=1}^{n} \|\dot{\mathbf{z}}(t_i) - \hat{\dot{\mathbf{z}}}(t_i;w)\|^2 \tag{5}$$

$$\text{s.t.} \quad \dot{\mathbf{z}}(t_i) = [\mathbf{J}(\mathbf{z}(t_i)) - \mathbf{R}(\mathbf{z}(t_i))]\frac{\partial \mathbf{H}(\mathbf{z}(t_i))}{\partial \mathbf{z}} + \mathbf{g}(\mathbf{z})\mathbf{u} \tag{6}$$

$$\hat{\dot{\mathbf{z}}}(\mathbf{t_i};\mathbf{w}) = \left[\mathbf{J}(\mathbf{z}(\mathbf{t_i})) - \hat{\mathbf{R}}(\mathbf{z}(\mathbf{t_i};\mathbf{w}))\right]\frac{\partial \hat{\mathbf{H}}(\mathbf{z}(\mathbf{t_i};\mathbf{w}))}{\partial \mathbf{z}} + \mathbf{g}(\mathbf{z})\mathbf{u}, \tag{7}$$

where n is the number of time samples, $w = \{W^{[0]}, b^{[0]}, W^{[1]}, b^{[1]}\}$ is the set of optimization variables, and $(\hat{\cdot})$ represents quantities estimated by the neural networks. We approach the solution w^* of (5) with respect to $V_p(\theta) := \sum_{\tau=t_0}^{t_f} \|\dot{z}^*(\tau) - \dot{z}(\tau)\|^2$ with an iterative gradient descent method

$$\theta^{n+1} = \theta^n - \alpha_n(\nabla_\theta V_p(\theta^n)),\ n = 0,1,2,\ldots \tag{8}$$

where the iteration maps $\alpha_n : \mathbb{R}^2 \to \mathbb{R}^2$, $n \geq 0$ are defined in accordance with the Adam method of moments for stochastic optimization [3], and the computation of the gradient vectors is implemented using automatic differentiation [4].

Experimental Results and Future Directions

We showcase the proposed algorithm in the complex swarm movements shown in Fig. 2 and 3, where the trajectories of the particles are generated by the Cucker-Smale and extended Boids models with one leader, respectively. The decoding results are detailed in [8] and show successful leader detection and reconstruction of the coordination laws of the UAV swarm.

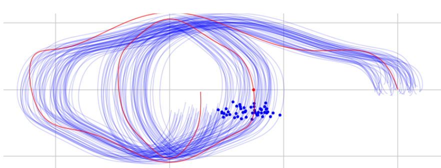

Figure 2: An example of 2D particle trajectories of a swarm following the dynamics of a Cucker-Smale model with one leader.

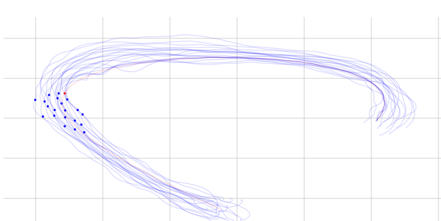

Figure 3: An example of 2D particle trajectories of a swarm following the dynamics of an extended Boids model with one leader.

In the case of two leaders, we consider the complex swarm movement shown in the Fig. 4. In order to apply our port-Hamiltonian based learning algorithm, we first estimate the sets $\mathcal{L}(i)$, $1 \leq i \leq N$ with the leader detection described above. The number of leaders was successfully detected, as well as the leaders themselves, even though there is an intersection between the trajectories. The interaction function is reconstructed with a mean squared error of $MSE = 0.193657$ [8].

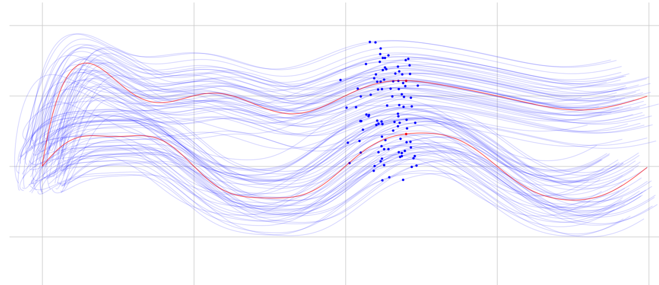

Figure 4: An example of 2D particle trajectories of a swarm following the dynamics of a Cucker-Smale model with two leaders.

While the key focus of this work is related to the defense against hostile UAV swarms, similar problems are found in many other types of large networked systems, including communication and computer networks, sensor networks, networked cyber-physical systems, biological systems, and social networks over the Internet. In such systems there are corresponding notions of leaders, such as initiators of a malicious attack, coordinators of malevolent behavior, initiators of a biological cell-malfunction, or influential sources of miss-information or untrustworthiness. In all these problems fast identification of the leaders and the associated follower groups (or influence groups) is essential for defending and correcting such malevolent actions and functions. Thus the applicability of the ideas and methods proposed in this work is very broad, with the appropriate modeling and semantic changes for the various domains. Important directions of our current and future research include extensions of the framework and algorithms to these broader domains, as well as the utilization of game theoretic methods for their analysis (non-cooperating, cooperating and mean-field games).

References

[1] Felipe Cucker and Steve Smale. Emergent behavior in flocks. *IEEE Transactions on automatic control*, 52(5):852–862, 2007.

[2] C. W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, 1969.

[3] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[4] D. Maclaurin, D. Duvenaud, M. Johnson, and J. Townsend. Autograd, 2018.

[5] Ion Matei, Christos Mavridis, John S. Baras, and Maksym Zhenirovskyy. Inferring particle interaction physical models and their dynamical properties. In *2019 IEEE Conference on Decision and Control (CDC)*, pages 4615–4621. IEEE, 2019.

[6] Christos Mavridis and John Baras. Online deterministic annealing for classification and clustering. *arXiv preprint arXiv:2102.05836*, 2021.

[7] Christos N Mavridis and John S Baras. Convergence of stochastic vector quantization and learning vector quantization with bregman divergences. In *21rst IFAC World Congress*. IFAC, 2020.

[8] Christos N Mavridis, Nilesh Suriyarachchi, and John S Baras. Detection of dynamically changing leaders incomplex swarms from observed dynamic data. In *2020 Conference on Decision and Game Theory for Security (GameSec)*, 2020.

[9] Christos N Mavridis, Amoolya Tirumalai, and John S Baras. Learning interaction dynamics from particle trajectories and density evolution. In *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, 2020.

[10] Christos N Mavridis, Amoolya Tirumalai, John S Baras, and Ion Matei. Semi-linear poisson-mediated flocking in a cucker-smale model. In *24th International Symposium on Mathematical Theory of Networks and Systems (MTNS)*. IFAC, 2021.

[11] C.W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *ACM SIGGRAPH computer graphics*, volume 21, pages 25–34. ACM, 1987.