

SIGMAS: Second-Order Interaction-based Grouping for Overlapping Multi-Agent Swarms

Minah Lee

The University of Texas at Dallas
Richardson, TX, USA
minah.lee@utdallas.edu

Saibal Mukhopadhyay

Georgia Institute of Technology
Atlanta, GA, USA
saibal.mukhopadhyay@ece.gatech.edu

ABSTRACT

Swarming systems, such as drone fleets and robotic teams, exhibit complex dynamics driven by both individual behaviors and emergent group-level interactions. Unlike traditional multi-agent domains such as pedestrian crowds or traffic systems, swarms typically consist of a *few large groups* with *inherent and persistent* memberships, making group identification essential for understanding fine-grained behavior. We introduce the novel task of **group prediction in overlapping multi-agent swarms**, where latent group structures must be inferred directly from agent trajectories without ground-truth supervision. To address this challenge, we propose **SIGMAS** (Second-order Interaction-based Grouping for Multi-Agent Swarms), a self-supervised framework that goes beyond direct pairwise interactions and model *second-order interaction* across agents. By capturing how similarly agents interact with others, SIGMAS enables robust group inference and adaptively balances individual and collective dynamics through a learnable gating mechanism for joint reasoning. Experiments across diverse synthetic swarm scenarios demonstrate that SIGMAS accurately recovers latent group structures and remains robust under simultaneously overlapping swarm dynamics, establishing both a new benchmark task and a principled modeling framework for swarm understanding.

KEYWORDS

Swarm Dynamics; Multi-Agent Systems; Group Identification; Self-Supervised Learning; Trajectory Prediction

ACM Reference Format:

Minah Lee and Saibal Mukhopadhyay. 2026. SIGMAS: Second-Order Interaction-based Grouping for Overlapping Multi-Agent Swarms. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/MZXX4786>

1 INTRODUCTION

Swarming systems, including drone fleets, robotic teams, and autonomous vehicles, are key enablers of scalable and robust autonomous operations [2, 4, 21, 24, 28]. Through simple local interactions, these systems exhibit emergent group-level behaviors that support coordination, collision avoidance, and efficient resource utilization in dynamic and uncertain environments [13, 19].

Compared to traditional multi-agent scenarios such as pedestrian motion, sports analytics, or traffic navigation [16, 25, 27], swarming systems differ in two fundamental aspects. First, traditional settings typically involve many small groups with sparse interactions, whereas swarms typically consist of a *few large groups* with dense intra-group coordination and complex inter-group dynamics, as illustrated in Figure 1(a). Second, the notion of a “group” also differs fundamentally. In traditional domains, groups are often defined by transient interactions arising from spatial proximity, while in swarming systems, group membership is an *intrinsic and persistent* property of each agent. Agents from different swarms may even occupy the same spatial region at the same time, yet maintain distinct coordination patterns.

A key challenge in multi-agent swarms is understanding the latent group structure that governs agent coordination. This is critical for robust behavior modeling and fine-grained interpretation of intra-swarm dynamics. To this end, we introduce the novel task of **group prediction in overlapping multi-agent swarms**, where the goal is to infer latent group memberships directly from agent trajectories without any ground-truth supervision. This self-supervised formulation is motivated by real-world swarm scenarios, where group affiliations are inherent but unobservable and may evolve over time. Since manually annotated group labels are rarely available in practice, we aim to learn group structure directly from motion dynamics without relying on explicit labels.

However, most existing trajectory prediction models are not well suited to this regime for two main reasons. First, many are designed for low-density scenes with sparse interactions, such as the ETH/UCY datasets [12, 17], which typically include only 1–6 agents within each observation and prediction window. Second, and more importantly, these models primarily rely on spatial proximity or pairwise motion similarity to model direct influence between agents, here referred to as *first-order attention* mechanisms (Figure 1(b)). While such mechanisms can be effective for short-term prediction, they are often insufficient to capture true group affiliations, especially in scenarios where multiple swarms overlap and occupy the same spatial region simultaneously, or where members of the same swarm are spatially dispersed. These limitations highlight the need for a specialized approach capable of understanding latent group structure in dense and dynamic swarm settings.

Therefore, we propose modeling social behavior through **second-order interactions**, where second-order interactions refer to the similarity between agents’ first-order attention patterns. In other words, instead of asking *whether* two agents directly interact, we examine *how similarly* they interact with others, as illustrated in Figure 1(c). While first-order attention captures localized dependencies and can be effective for short-term trajectory forecasting,



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/MZXX4786>

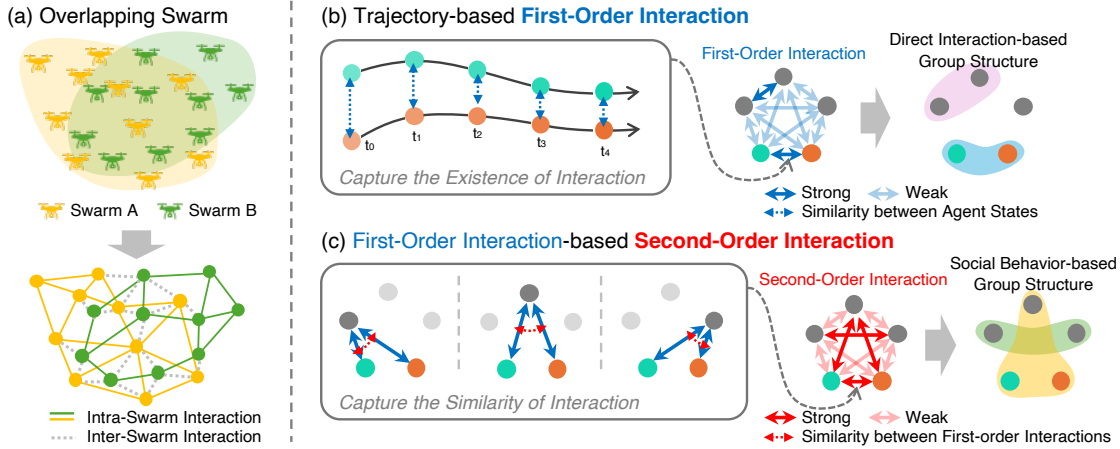


Figure 1: (a) Overlapping swarms with distinct coordination patterns, making group identification difficult based on proximity alone. (b) *First-order interactions* capture direct motion correlations between agents. (c) *Second-order interactions* compare agents’ first-order attention patterns to measure how similarly they interact with others. This captures *social behavior similarity* and enables robust group identification, even in the absence of direct interactions.

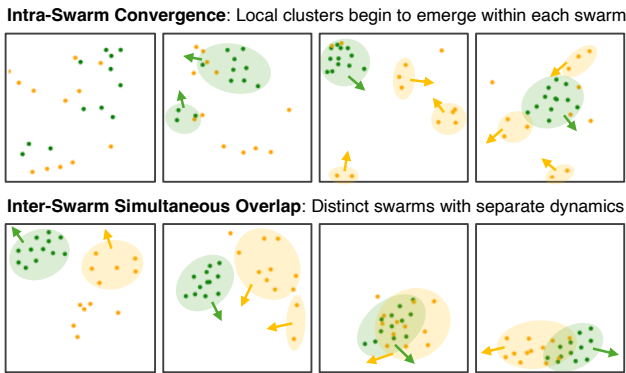


Figure 2: Illustration of emergent swarm behaviors. **Top:** Intra-swarm convergence, where agents within the same swarm gradually align and form cohesive clusters. **Bottom:** Inter-swarm overlap, where distinct swarms simultaneously share spatial regions while maintaining separate coordination dynamics.

it often conflates transient spatial proximity with true group structure. In contrast, second-order interactions compare each agent’s full interaction profile, revealing shared behavioral patterns and uncovering latent affiliations. This approach is particularly useful in dense, overlapping, or partially converged swarm formations, where direct interactions are no longer a reliable signal of group membership.

To address these challenges, we introduce **SIGMAS (Second-order Interaction-based Grouping for Multi-Agent Swarms)**, a self-supervised framework for latent group inference in swarming systems. SIGMAS learns directly from agent trajectories without requiring manual annotations. SIGMAS features a novel swarm-level encoder that computes *second-order interactions* by comparing agents’ first-order attention patterns. This enables SIGMAS to model interaction similarity and recover cohesive group structures even when agents are spatially entangled or dynamically dispersed. In addition, SIGMAS integrates a dynamic balancing mechanism to fuse individual and group-level dynamics, and applies spectral clustering for self-supervised group inference. These components jointly allow SIGMAS to capture both localized motion and global coordination patterns, enhancing group understanding in complex

multi-agent environments. This paper makes the following unique contributions:

- We formulate a novel task of **group prediction in overlapping multi-agent swarms**, where latent group memberships must be inferred from dense agent trajectories without ground-truth supervision. To the best of our knowledge, this is the first work to explicitly address this problem setting.
- We present **SIGMAS**, a self-supervised framework that models group-level dynamics via second-order interaction modeling and adaptively fuses them with agent-level features through a learnable gating mechanism, enabling joint reasoning over local and global behaviors.
- We empirically demonstrate that SIGMAS accurately recovers latent group structures and maintains robust performance across diverse and overlapping swarm scenarios.

2 BACKGROUND

2.1 Swarm Behavior and Interaction Modeling

Craig W. Reynolds [20] introduced three canonical rules for flocking behavior in agent-based systems: *alignment*, *cohesion*, and *separation*. Each agent adjusts its velocity based on its local neighbors, leading to emergent collective behaviors such as coordinated motion and dynamic formation. We simulate these behaviors using the *agentpy* framework [5], which provides fine-grained control over agent-level parameters and swarm configurations. In our setup, agents within the same swarm share identical alignment, cohesion, and separation strengths, reflecting the assumption that *group membership dictates underlying motion dynamics*. For agents in different swarms, only the separation rule applies to avoid collision. Figure 2 illustrates two key swarm phenomena that emerge from these local interaction rules:

Intra-Swarm Convergence (Top Row). Agents within the same swarm gradually align their velocities, forming cohesive motion patterns over time. As interactions accumulate, local clusters emerge and evolve into tightly coordinated formations. The rate of convergence depends on intra-group interaction strength: higher alignment and cohesion lead to faster convergence and denser formations (green agents), while lower values yield slower and more diffuse structures (yellow agents).

Inter-Swarm Overlap (Bottom Row). Distinct swarms may occupy overlapping spatial regions simultaneously while maintaining separate coordination dynamics. Depending on the relative strengths of intra- and inter-swarm interactions, agents may either preserve their original affiliations or partially mix across swarm boundaries.

These examples illustrate the challenge of inferring latent group structures in swarming systems. Spatial proximity alone is often misleading, as agents from different swarms may overlap while following distinct coordination patterns. This motivates the design of SIGMAS, which infers group membership by modeling *second-order interaction*, offering a more robust foundation for group prediction in complex, overlapping swarm environments.

2.2 Trajectory-based Interaction Modeling in Multi-Agent Systems

Trajectory prediction has long been a core problem in multi-agent systems, with applications ranging from pedestrian modeling to human-robot interaction and autonomous navigation [6, 22, 29]. A key challenge lies in capturing complex inter-agent dependencies, that influence motion, coordination, and intent understanding. Early approaches relied on handcrafted social rules [1] or pooled neighborhood encodings [10], while recent works adopt graph-based and attention-based neural architectures for more expressive interaction modeling.

Most existing works model inter-agent interactions through *first-order mechanisms*, where interactions are captured directly at the node or edge level. Node-centric approaches, such as SocialGAN [6], AgentFormer [29], and Leapfrog [14], focus on how individual agents respond to their neighbors based on proximity, attention, or goal-conditioned latent representations. Edge-centric models, such as Neural Relational Inference (NRI) [9], infer latent pairwise relations between agents by predicting the presence and type of edges that best explain observed dynamics. These methods are effective in low-density environments with moderate agent counts and clearly separated groups, such as pedestrian crowds or physical particle simulations.

To address the limitations of purely pairwise modeling, several methods incorporate group-level cues as auxiliary signals to improve prediction accuracy. For example, GroupNet [26] optimizes a social affinity matrix to identify latent groups, while SocialCircle [25] enforces temporal group consistency to enhance trajectory forecasting. Other methods leverage group priors or apply graph partitioning to better structure the interaction space [3, 8, 11]. However, these methods do not explicitly model group emergence as a central objective.

In contrast, our work targets a fundamentally different regime: inferring latent group structure in dense, overlapping swarms. In such settings, group identity is not a transient outcome of spatial proximity but an intrinsic, persistent property of each agent, emerging from how similarly they interact with others. This requires new modeling methods that go beyond direct pairwise influence. To this end, we introduce *second-order interactions*, which measure the similarity between agents’ first-order attention patterns. This mechanism enables robust group inference in densely interacting, dynamically evolving swarms where direct interactions may be

Table 1: Comparison of key capabilities for group prediction in overlapping multi-agent swarms. ‘Traj. Pred.’ denotes trajectory prediction and ‘Inter. Type’ refers to the modeled interaction type. Δ indicates that group modeling is used internally in the latent space to support trajectory prediction, but is not evaluated as a standalone prediction objective.

Models	Method	Traj. Pred.	Inter. Type	Group Modeling	Group Inference
SocialGAN [6]	LSTM	✓	1st	✗	✗
Trajectron++ [22]	LSTM	✓	1st	✗	✗
AgentFormer [29]	Transformer	✓	1st	✗	✗
Leapfrog [14]	Diffusion	✓	1st	✗	✗
GroupNet [26]	GNN	✓	1st	Δ	✗
SocialCircle [25]	Transformer	✓	1st	Δ	✗
SIGMAS (Ours)	Transformer	✓	1st & 2nd	✓	✓

ambiguous or unreliable. As summarized in Table 1, SIGMAS is the only framework that simultaneously supports trajectory prediction, second-order interaction reasoning, and self-supervised group inference, all of which are essential for understanding complex group behavior in multi-agent swarms.

3 PROBLEM STATEMENT

Since group labels and the number of groups are unavailable during training, we formulate overlapping swarm identification as a **joint task of trajectory prediction and latent group inference** in a multi-agent swarm. Intuitively, given only the past H -step trajectories of agents (as shown in Figure 1(a)), the model must jointly predict their future motion over T time steps and identify the underlying group structure, i.e., $\mathcal{G} = \{\{yellows\}, \{greens\}\}$.

To formalize this task, let N agents be observed up to time $t \leq 0$. This joint state is denoted by $\mathbf{X}^t = (x_1^t, x_2^t, \dots, x_N^t)$, where $x_n^t \in \mathbb{R}^d$ is the position of agent n at time t . The observation history up to time 0 is $\mathbf{X} = (\mathbf{X}^{-H}, \dots, \mathbf{X}^0)$, with temporal horizon H . For future time steps $t > 0$, the corresponding joint state is $\mathbf{Y}^t = (y_1^t, y_2^t, \dots, y_N^t)$, where $y_n^t \in \mathbb{R}^d$ denotes the future position of agent n . The full future trajectories over T time steps is $\mathbf{Y} = (\mathbf{Y}^1, \dots, \mathbf{Y}^T)$.

In addition to predicting \mathbf{Y} , the goal is to infer latent group memberships for all agents. Let $\mathcal{G} = \{g_1, \dots, g_C\}$ denote the set of groups, where the number of groups C is unknown during training, and each $g_c \subseteq \{a_1, \dots, a_N\}$ represents the subset of agents assigned to group c .

The overall objective is to jointly estimate both the future trajectories \mathbf{Y} and the latent group structure \mathcal{G} from past observations \mathbf{X} in a self-supervised manner: $p_\theta(\mathbf{Y}, \mathcal{G} \mid \mathbf{X})$, where θ denotes the model parameters. This formulation enables joint modeling of motion dynamics and latent group organization without requiring ground-truth group labels during training.

4 SIGMAS

We propose **SIGMAS (Second-order Interaction-based Grouping for Multi-Agent Swarms)**, a self-supervised framework that learns swarm-aware trajectory embeddings by jointly modeling individual motion and collective coordination. As illustrated in Figure 3, SIGMAS extracts agent-level and swarm-level trajectory embedding separately, and then fuses them to create a unified, swarm-aware representation for each agent. This fused embedding

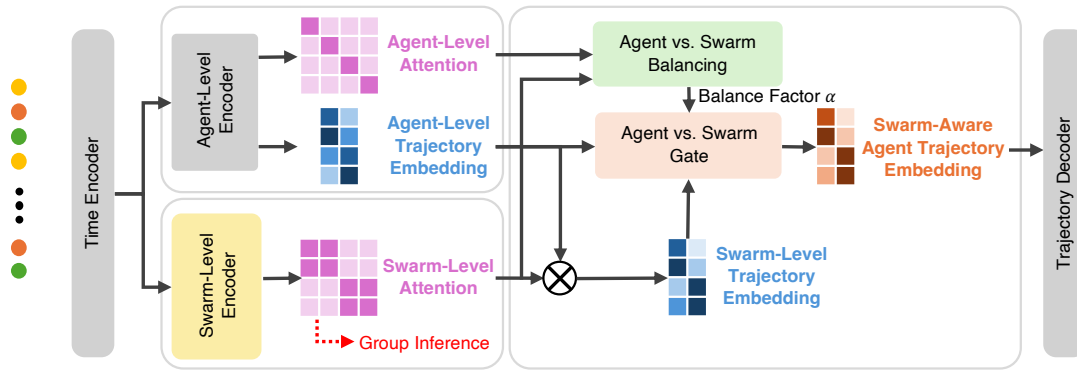


Figure 3: SIGMAS architecture. Agent-level and swarm-level features are extracted via separate encoders, and fused using balancing and gating to form swarm-aware embeddings for trajectory prediction and group inference.

is used to predict future trajectories. In addition, the swarm-level attention matrix from the Swarm-Level Encoder is used to infer latent group structures during testing. The framework consists of three main components: (1) an **Agent-Level Encoder** that models first-order interactions and individual motion patterns, (2) a **Swarm-Level Encoder** that captures second-order interaction similarity across agents, and (3) a **Balancing & Gating Module** that adaptively fuses individual and swarm-level representations.

4.1 Agent-Level Encoder

The Agent-Level Encoder captures localized dynamics and direct pairwise interactions among agents. Given temporally encoded trajectories, it produces two types of representations:

- **Agent-Level Attention:** captures pairwise behavioral correlations (*first-order interactions*) by computing attention scores based on trajectory similarity.
- **Agent-Level Trajectory Embedding:** represents each agent’s individual motion context.

We adopt the encoder architecture from AgentFormer [29], which employs an agent-aware spatiotemporal attention mechanism. This allows each agent to attend differently to its neighbors depending on motion context, enabling flexible modeling of time-varying inter-agent dependencies. The resulting attention matrix reflects first-order interaction patterns, as illustrated in Figure 1(b).

4.2 Swarm-Level Encoder

Swarm-Level Encoder captures *social behavior similarity* among agents by modeling second-order interactions. It is motivated by the observation that agents within the same swarm tend to interact with others in similar ways, even without direct influence. The architecture is illustrated in Figure 4.

We first compute *first-order interaction features* that capture direct correlations between agents based on their past trajectories. Given the trajectory embeddings $\mathbf{Q}, \mathbf{K} \in \mathbb{R}^{N \times d}$, we compute the first-order attention matrix as

$$\mathbf{A}^{(1)} = \text{Softmax} \left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}} \right)$$

where $A_{ij}^{(1)}$ represents how strongly agent i attends to agent j . These attention scores reflect directional influence and localized behavioral dependencies among agents. For example, in Figure 4(a), the red agent attends strongly to the purple agent, while the green agent distributes its attention more broadly across the population.

To understand group-level structure, we compute *second-order interactions* by comparing first-order attention patterns across agents:

$$\mathbf{A}^{(2)} = \mathbf{A}^{(1)} \mathbf{A}^{(1)T}$$

Each entry $A_{ij}^{(2)}$ measures the similarity between agents i and j based on how similarly they attend to others, thus quantifying their *social behavior similarity*. If two agents have similar attention distribution, they likely exhibit aligned social behaviors and therefore belong to the same swarm. In Figure 4(b), the green and blue agents both attend strongly to the red and purple agents, indicating similar social behavior and suggesting shared group affiliation.

The resulting second-order interaction matrix encodes pairwise similarity in social behavior. During training, we regularize this matrix to be *symmetric* and *positive semi-definite* (Section 5.2), enforcing mutual consistency in social similarity. That is, if agent i views agent j as socially similar, the converse must also hold. During testing, $\mathbf{A}^{(2)}$ is passed to the clustering module (Section 5.1), which infers soft group assignments and, for instance, successfully groups the red, yellow, green, and blue agents into the same swarm.

4.3 Agent vs. Swarm Balancing and Gating

As shown in Figure 3, the agent-level encoder produces two outputs: a first-order attention map $\mathbf{A}^{(1)} \in \mathbb{R}^{HN \times HN}$, which captures pairwise interactions across H observed time steps, and an agent-level trajectory embedding $\mathbf{f}_n^{\text{agent}} \in \mathbb{R}^{HN \times d}$. The swarm-level encoder generates a second-order attention matrix $\mathbf{A}^{(2)} \in \mathbb{R}^{HN \times HN}$ that encodes second-order social similarity between agents. Using this matrix, we derive the swarm-level trajectory embedding via attention-weighted aggregation: $\mathbf{f}_n^{\text{group}} = \mathbf{A}^{(2)} \mathbf{f}_n^{\text{agent}} \in \mathbb{R}^{HN \times d}$.

The key challenge lies in how to combine $\mathbf{f}_n^{\text{agent}}$ and $\mathbf{f}_n^{\text{group}}$ into a unified representation that balances individual motion cues with group-level coordination. The optimal balance can vary depending on the swarm state, for example, the level of intra-group convergence, overlap across groups, or directional alignment, making static fusion suboptimal.

To address this, we introduce an adaptive gating mechanism that fuses the two embeddings in a context-aware manner. A soft gating coefficient $\alpha_n \in [0, 1]$ is computed for each agent using a multi-layer perceptron (MLP) over its agent-level embedding $\mathbf{f}_n^{\text{agent}}$, swarm-level embedding $\mathbf{f}_n^{\text{group}}$, and normalized group-assignment

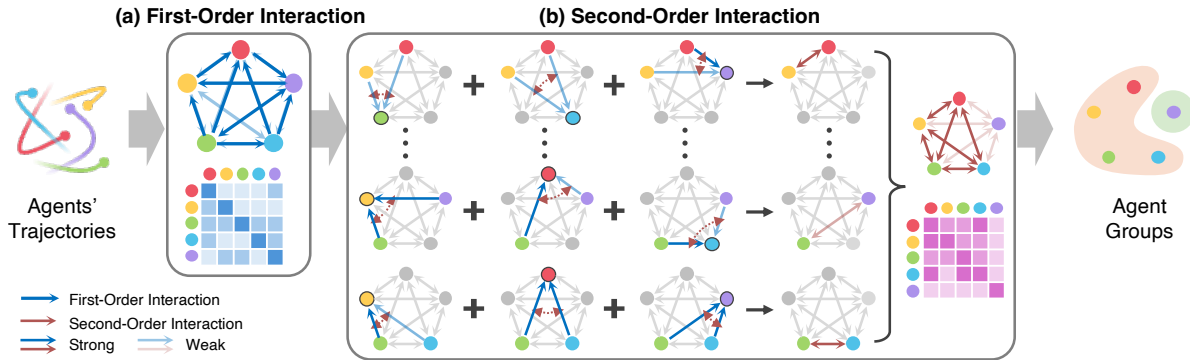


Figure 4: Overview of second-order interaction modeling in the Swarm-Level Encoder. (a) First-order interactions encode directional influence between agents. (b) Second-order interactions compare first-order attention patterns across agents, revealing social behavior similarity. The resulting matrix highlights agents with similar interaction profiles, enabling group identification via clustering. In the attention matrix, rows correspond to attending agents and columns to agents being attended to. Darker cells indicate stronger outgoing attention.




Figure 5: SIGMAS-based CVAE prediction framework. Dotted paths are used during training only.

entropy \mathcal{H}_n :

$$\alpha_n = \text{MLP} \left([f_n^{\text{agent}} \parallel f_n^{\text{group}} \parallel \mathcal{H}_n] \right),$$

where $\mathcal{H}_n = -\sum_{j=1}^{HN} A_{nj}^{(2)} \log A_{nj}^{(2)} / \log(HN)$ is the normalized entropy of the n -th row of $A^{(2)}$. This entropy reflects the level of convergence within the swarm. High entropy indicates broad attention across peers (indicating a cohesive swarm), while low entropy suggests focused attention on a few agents (indicating pre-converged or fragmented structures).

The final trajectory embedding is computed as:

$$f_n^{\text{final}} = (1 - \alpha_n) \cdot f_n^{\text{agent}} + \alpha_n \cdot f_n^{\text{group}}.$$

This dynamic fusion allows the model to adjust the contribution of individual versus group-level cues based on the agent’s social context.

5 SIGMAS-BASED PREDICTION FRAMEWORK

We integrate SIGMAS into a conditional variational autoencoder (CVAE)-based multi-agent trajectory prediction framework following [26] (Figure 5). During the encoder stage, both past and future trajectories are used to estimate the posterior distribution parameters (μ_q, σ_q) . SIGMAS is applied in this stage to produce swarm-aware trajectory embeddings that capture both individual and group-level behavior via first- and second-order interactions. In the decoder stage, a series of residual blocks take the past trajectory and a sampled latent embedding as input to generate both the predicted future trajectory and a reconstruction of the past. This residual design helps preserve input information while refining future predictions.

5.1 Group Inference via Spectral Clustering.

The second-order attention matrix $A^{(2)} \in \mathbb{R}^{N \times N}$ encodes meaningful group-relevant information by capturing similarities in social behavior. At early stages of swarm formation, when agents from the same group are still spatially dispersed, $A^{(2)}$ may reflect multiple subgroup structures corresponding to partially converged clusters. In this case, a post-processing step during inference is necessary to merge these subgroups based on their dynamic similarity and given the number of groups k .

To infer group membership from $A^{(2)}$, we treat it as a soft adjacency matrix and apply spectral clustering using the normalized graph Laplacian [15]. Specifically, we compute:

$$L = I - D^{-1/2} A^{(2)} D^{-1/2},$$

where D is the degree matrix with diagonal entries $D_{ii} = \sum_j A_{ij}^{(2)}$. We then extract the k smallest eigenvectors of L to obtain low-dimensional agent embeddings that capture the global interaction structure. k -means clustering is then applied to these embeddings to produce the predicted group assignments.

This procedure is fully self-supervised and relies on the geometry of the learned attention matrix to uncover cohesive swarm structures, even under spatial overlap or incomplete convergence, without requiring ground-truth labels.

5.2 Training Objectives

The training objective of SIGMAS consists of three components: a trajectory prediction loss, a swarm-level attention regularization loss, and a balancing factor regularization loss. These are jointly optimized to ensure accurate forecasting while promoting structured group-aware representations.

Trajectory Prediction Loss. We adopt a composite loss [26] that includes: (1) an evidence lower bound (ELBO) loss from the CVAE formulation, (2) a reconstruction loss on the past trajectory, and (3) a diversity-promoting variety loss following Social-GAN:

$$\begin{aligned} \mathcal{L}_{\text{elbo}} &= \|\widehat{X}^+ - X^+\|_2^2 + \text{KL} \left(\mathcal{N}(\mu_q, \text{Diag}(\sigma_q^2)) \parallel \mathcal{N}(0, \lambda I) \right), \\ \mathcal{L}_{\text{rec}} &= \|\widehat{X} - X\|_2^2, \\ \mathcal{L}_{\text{variety}} &= \min_k \|\widehat{X}^{+(k)} - X^+\|_2^2, \end{aligned}$$

$$\mathcal{L}_{\text{traj}} = \mathcal{L}_{\text{elbo}} + \mathcal{L}_{\text{rec}} + \mathcal{L}_{\text{variety}},$$

where $\widehat{\mathbf{X}}$ is the reconstructed past trajectory, $\widehat{\mathbf{X}}^+$ is the predicted future trajectory, \mathbf{X}^+ is the ground-truth future trajectory, and $\widehat{\mathbf{X}}^{+(k)}$ represents the k -th stochastic sample. $\|\cdot\|_2$ denotes the ℓ_2 norm, and $\text{KL}(\cdot\|\cdot)$ denotes the Kullback–Leibler divergence. This loss encourages the model to generate accurate, diverse predictions while preserving past dynamics and maintaining a regularized latent representation.

Swarm-Level Attention Regularization. To ensure that the second-order interaction matrix $\mathbf{A}^{(2)}$ captures meaningful group structure, we introduce a regularization loss on the swarm-level similarity matrix $\mathbf{S} \in \mathbb{R}^{N \times N}$, obtained by averaging $\mathbf{A}^{(2)}$ across the observation horizon. We enforce two key properties to make \mathbf{S} a valid affinity kernel:

- **Symmetry Loss.** Social similarity should be mutual: if agent i is behaviorally similar to agent j , then the reverse should also hold. We encourage symmetry by minimizing:

$$\mathcal{L}_{\text{sym}} = \|\mathbf{S} - \mathbf{S}^\top\|_F^2,$$

which penalizes asymmetric affinity scores and promotes bidirectional consistency across agent pairs.

- **Positive Semi-Definite (PSD) Loss.** To ensure \mathbf{S} is a valid inner-product kernel for spectral clustering, we encourage it to be positive semi-definite. Specifically, we penalize negative eigenvalues:

$$\mathcal{L}_{\text{psd}} = \sum_{i=1}^N [\max(0, -\lambda_i)]^2,$$

where λ_i is the i -th eigenvalue of \mathbf{S} . This regularization ensures geometric coherence in the learned affinity space.

The total swarm-level regularization is defined as:

$$\mathcal{L}_{\text{swarm}} = \mathcal{L}_{\text{sym}} + \mathcal{L}_{\text{psd}}.$$

These constraints guide $\mathbf{A}^{(2)}$ to form a symmetric, geometrically valid similarity matrix that supports interpretable and robust group identification during inference.

Balancing Factor Regularization. To enable adaptive fusion between agent-level and swarm-level dynamics, the model learns a gating factor $\alpha_n \in [0, 1]$ for each agent n , as introduced in Section 4.3. Ideally, α_n should consider both the degree of swarm convergence as well as agent’s local motion.

To guide this behavior, we introduce a regularization loss that aligns the learned α_n with a target value based on the entropy of the swarm-level attention. Let \mathcal{H}_n denote the normalized entropy of the n -th row of the second-order attention matrix $\mathbf{A}^{(2)}$. We define the target balancing factor as:

$$\tilde{\alpha}_n = 1 - \mathcal{H}_n.$$

Low entropy indicates that the agent strongly attends to a small subset of peers, suggesting sub-convergence and a stronger need to rely on swarm-level information. We encourage alignment via the loss:

$$\mathcal{L}_{\text{balance}} = \frac{1}{N} \sum_{n=1}^N (\alpha_n - \tilde{\alpha}_n)^2.$$

This term promotes dynamic adjustment of the agent’s attention balance based on the current level of social organization, improving adaptability across diverse swarm formations.

Total Loss. The total objective is a weighted sum of the three terms:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{traj}} + \mathcal{L}_{\text{swarm}} + \mathcal{L}_{\text{balance}}.$$

6 SIMULATION RESULTS

6.1 Experimental Setup

Datasets. Due to the lack of public datasets for large, overlapping multi-agent swarms, we simulate swarm dynamics using a flocking-based agent simulator (AgentPy [5]) based on Reynolds’ rules [20]. To constrain agent motion, we apply a reflective boundary condition and keep agents bouncing back at the simulation edge. The dataset contains 120 sequences in total: 100 for training, 10 for validation, and 10 for testing. Each sequence consists of 24 agents divided into two groups of 12 agents each, moving over 200 time steps. Agents are initialized with random positions and velocities and gradually converge into coherent swarms through local interactions.

Evaluation Metrics. Let $\mathbf{z} = (z_1, \dots, z_N)$ denote the ground-truth group labels and $\hat{\mathbf{z}} = (\hat{z}_1, \dots, \hat{z}_N)$ the predicted assignments. We evaluate clustering performance using three standard metrics: (1) *Adjusted Rand Index (ARI)* [7]: ARI measures pairwise agreement between predicted and true groupings, correcting for chance. Let n_{ij} be the number of agents in true cluster i assigned to predicted cluster j , $a_i = \sum_j n_{ij}$, $b_j = \sum_i n_{ij}$, and $N = \sum_{i,j} n_{ij}$. Then:

$$\text{ARI} = \frac{\sum_{i,j} \binom{n_{ij}}{2} - \frac{\sum_i \binom{a_i}{2} \sum_j \binom{b_j}{2}}{\binom{N}{2}}}{\frac{1}{2} \left[\sum_i \binom{a_i}{2} + \sum_j \binom{b_j}{2} \right] - \frac{\sum_i \binom{a_i}{2} \sum_j \binom{b_j}{2}}{\binom{N}{2}}}.$$

ARI ranges from -1 (anti-correlation) to 1 (perfect match), with 0 indicating random labeling.

(2) *Normalized Mutual Information (NMI)* [23]: NMI measures the mutual information between ground-truth and predicted clusters, normalized by their entropies. Let p_{ij} , p_i , and p_j denote the joint and marginal probabilities from the contingency table:

$$I(C; \hat{C}) = \sum_{i,j} p_{ij} \log \left(\frac{p_{ij}}{p_i p_j} \right)$$

$$H(C) = - \sum_i p_i \log p_i, \quad H(\hat{C}) = - \sum_j p_j \log p_j$$

The normalized score is:

$$\text{NMI}(C, \hat{C}) = \frac{2I(C; \hat{C})}{H(C) + H(\hat{C})}.$$

which lies in $[0, 1]$, with higher values indicating stronger agreement.

(3) *F-score* [18]: The F-score combines pairwise precision and recall. For all agent pairs, define TP (true positives), FP (false positives), and FN (false negatives). Then:

$$\text{Precision} = \frac{TP}{TP + FP}, \quad \text{Recall} = \frac{TP}{TP + FN}, \quad F = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}.$$

F-score balances both false groupings and missed groupings, with $F \in [0, 1]$.



Figure 6: Comparison of first-order and second-order attention for group inference. (a) Agent trajectories show two overlapping swarms. (b) Agent-level attention $A^{(1)}$ reflects direct interaction relevance for trajectory prediction. (c) Swarm-level attention $A^{(2)}$ captures second-order interaction similarity across agents, enabling accurate latent group discovery via spectral clustering.

Table 2: Swarm configurations for different scenarios.

Swarm Config	Cohesion	Separation	Alignment	Speed
Swarm A	[0.005, 0.007]	[0.1, 0.2]	[0.3, 0.1]	[1.0, 1.2]
Swarm B	[0.005, 0.010]	[0.1, 0.3]	[0.1, 0.3]	[1.0, 1.3]
Swarm C	[0.005, 0.010]	[0.1, 0.4]	[0.1, 0.3]	[1.0, 1.5]

6.2 Validation on Group Inference

We validate the group inference capability of SIGMAS in a simulated environment with two swarming groups exhibiting overlapping trajectories and distinct internal dynamics. While both groups follow Reynolds’ flocking rules, they differ in cohesion, alignment, and movement speed, resulting in heterogeneous intra-swarm behavior. No explicit inter-group interaction is applied, creating scenarios where agents from different swarms spatially interleave. Swarm A in Table 2 is used.

Latent Clustering. This experiment demonstrates how the proposed swarm-level encoder, based on second-order interactions, enables accurate group inference in partially converged, overlapping swarm scenarios. Figure 6(a) visualizes trajectories from the two swarms, where each swarm forms local sub-groups that have not yet fully converged. Effective group prediction in such scenarios requires identifying fine-grained sub-structures while preserving global swarm affiliation.

Figure 6(b) presents the first-order attention matrix $A^{(1)}$ from the agent-level encoder (Section 4.1), which captures direct behavioral correlations. While this matrix is useful for trajectory prediction, it primarily reflect transient spatial dependencies. For example, agents {13, 18, 19} from the yellow swarm attend to nearby green agents to avoid collisions. Such interactions are helpful for motion forecasting but do not reliably indicate shared group membership, making $A^{(1)}$ less effective for clustering.

In contrast, Figure 6(c) presents the second-order attention matrix $A^{(2)}$ (Section 4.2), which compares first-order attention patterns across agents to capture social behavior similarity. This matrix

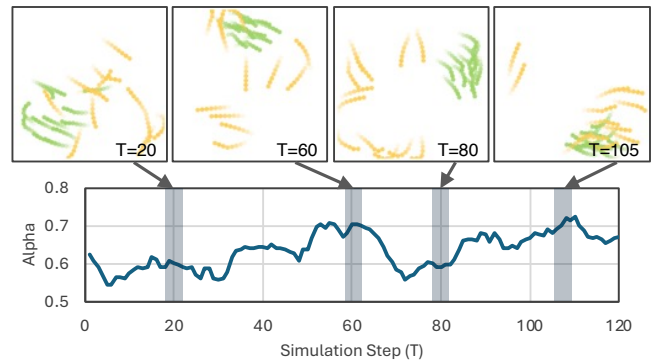


Figure 7: Adaptive balancing factor α over time in response to swarm dynamics. The model increases α during inter-swarm overlap and decreases it when individual motion cues are more predictive. $\alpha = 0$ corresponds to relying solely on agent-level trajectory embeddings, while $\alpha = 1$ reflects full reliance on swarm-level embeddings.

exhibits a clear block-diagonal structure with low inter-swarm similarity, effectively separating the green and yellow swarms. It also uncovers four latent sub-groups (S1–S4). Spectral clustering analyzes the eigenspace of the graph Laplacian and merges S1 and S2 into one swarm, and S3 and S4 into another, based on the known number of groups.

This example demonstrates that second-order interaction modeling captures both intra-swarm coordination and inter-swarm separation, enabling robust group inference even under partial convergence.

Individual vs Swarm Balancing. Since SIGMAS is trained with trajectory supervision, the utility of latent group information depends on the swarm’s configuration. It is therefore crucial to adaptively balance agent-level and swarm-level embeddings to capture both individual motion and collective behavior. To achieve this, SIGMAS learns a dynamic gating factor $\alpha \in [0, 1]$ that adjusts the contribution of swarm-level information relative to individual dynamics.

Figure 7 shows how α evolves over time in a multi-swarm simulation, along with corresponding snapshots at selected timesteps. Early in the sequence (e.g., $T = 20$), agents are loosely coordinated, so trajectory prediction relies more on individual behavior, resulting in low α . As intra-swarm cohesion increases, agents begin to exhibit more structured group motion. However, when swarms remain well-separated (e.g., $T = 80$), group identity becomes trivial, and individual cues again dominate, lowering α .

In contrast, during inter-swarm overlap (e.g., $T = 60, T = 105$), spatial proximity between agents from different swarms introduces ambiguity. In these cases, swarm-level context becomes essential for resolving latent affiliations, leading to higher α values. This adaptive gating mechanism enables SIGMAS to dynamically shift its reliance based on the swarm interactions.

Comparison with Prior Work. We evaluate group prediction accuracy by comparing SIGMAS against relevant baselines. As group prediction in swarming systems is a novel task, few existing models are directly applicable. Furthermore, many multi-agent

Table 3: Comparison of group prediction performance across baselines.

	AgentFormer [29]	SIGMAS_IG	SIGMAS
ARI \uparrow	0.3155	0.3306	0.4045
NMI \uparrow	0.3739	0.3860	0.4086
F-score \uparrow	0.6863	0.6887	0.6968

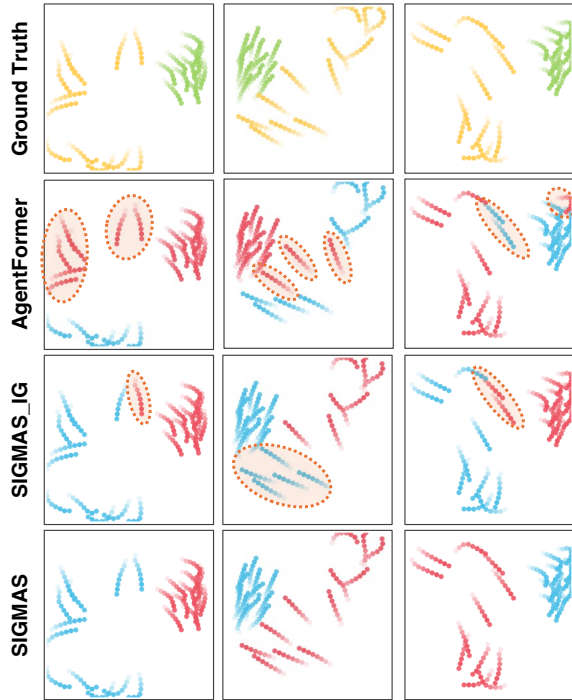


Figure 8: Qualitative results of group prediction. Red dotted ellipses indicate incorrect group assignments. Group labels are permutation-invariant, so different colors indicate distinct predicted clusters. Results videos are in Supplementary.

trajectory models are not designed to handle large-scale swarms. We consider the following baselines:

- **AgentFormer [29]:** A transformer-based trajectory model that uses self-attention to capture pairwise interactions. For group prediction, we apply spectral clustering to its final attention matrix.
- **SIGMAS_IG:** A variant of our model that uses both encoders during training but applies spectral clustering to the agent-level attention matrix $A^{(1)}$ instead of the second-order matrix.

Our full model, SIGMAS, performs group inference using $A^{(2)}$ from the swarm-level encoder, which captures second-order social behavior similarity.

Figure 8 presents qualitative results. AgentFormer and SIGMAS_IG frequently produce incorrect groupings in cases with spatial overlap or partial convergence, often grouping agents with similar motion direction rather than true swarm affiliation (see red ellipses). In contrast, SIGMAS successfully recovers the underlying group structure by leveraging second-order interactions, enabling robust inference even under ambiguous conditions. These observations are reflected




Figure 9: Group prediction performance across three heterogeneous swarm configurations (Swarm A, B, C) with increasing intra-swarm disparity. Shaded regions indicate standard deviation across multiple runs. Results videos are in Supplementary.

in Table 3, where SIGMAS achieves better performance across all clustering metrics.

6.3 Case study: Intra-Swarm Heterogeneity

To evaluate the robustness of SIGMAS under varying swarm dynamics, we conduct experiments across three heterogeneous swarm settings (Swarm A, B, C in Table 2). In each case, the two swarms differ in cohesion, separation, alignment, and speed, simulating different levels of swarm heterogeneity. A single SIGMAS model trained on Swarm A is evaluated on all three configurations to assess its generalizability.

Figure 9 shows the Adjusted Rand Index (ARI) over time for all three scenarios. As agents gradually converge into coherent swarms, ARI improves steadily. Importantly, SIGMAS achieves consistent performance across all settings despite the growing heterogeneity in swarm dynamics. This indicates that SIGMAS reliably infers group structure even under diverse intra-swarm dynamics. Similar trends are observed for NMI and F-score, confirming the model’s generalizability to a wide range of heterogeneous swarm behaviors.

7 CONCLUSIONS

Swarming systems exhibit complex, emergent behaviors driven by dense local interactions. Unlike traditional multi-agent settings, multiple swarms often operate in overlapping regions while maintaining persistent group identities, making latent group structure inference both critical and underexplored. In this paper, we introduce the novel task of **group prediction in overlapping multi-agent swarms**, where the objective is to recover group memberships directly from trajectory data without ground-truth supervision. This problem setting requires modeling beyond local motion interactions to capture global coordination behavior inherent in swarm dynamics

To address this challenge, we propose **SIGMAS**, a self-supervised framework that jointly captures individual and collective behavior through a novel *second-order interaction* mechanism. By measuring how similarly agents attend to others, SIGMAS quantifies social behavior similarity and supports robust group inference even in spatially entangled or partially converged swarm formations. An adaptive gating module fuses agent-level and swarm-level embeddings based on social context. Experiments demonstrate that SIGMAS outperforms prior baselines and generalizes well across diverse heterogeneous swarm scenarios.

ACKNOWLEDGMENTS

This work is supported by the Defense Advanced Research Projects Agency (DARPA) under Grant Numbers GR00019153. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Department of Defence, DARPA, or the U.S. Government.

REFERENCES

- [1] Alexandre Alahi, Kratharth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. 2016. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 961–971.
- [2] Faris A Almalki and Marios C Angelides. 2024. A synthesis of machine learning and internet of things in developing autonomous fleets of heterogeneous unmanned aerial vehicles for enhancing the regenerative farming cycle. *Computing* 106, 12 (2024), 4167–4192.
- [3] Inhwan Bae, Jin-Hwi Park, and Hae-Gon Jeon. 2022. Learning pedestrian group representations for multi-modal trajectory prediction. In *European Conference on Computer Vision*. Springer, 270–289.
- [4] Jayson Boubin, Codi Burley, Peida Han, Bowen Li, Barry Porter, and Christopher Stewart. 2021. Programming and deployment of autonomous swarms using multi-agent reinforcement learning. *arXiv preprint arXiv:2105.10605* (2021).
- [5] Joël Foramitti. 2021. AgentPy: A package for agent-based modeling in Python. *Journal of Open Source Software* 6, 62 (2021), 3065.
- [6] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. 2018. Social gan: Socially acceptable trajectories with generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2255–2264.
- [7] Lawrence Hubert and Phipps Arabie. 1985. Comparing partitions. *Journal of classification* 2, 1 (1985), 193–218.
- [8] Sungjune Kim, Hyung-gun Chi, Hyerin Lim, Karthik Ramani, Jinkyu Kim, and Sangpil Kim. 2024. Higher-order Relational Reasoning for Pedestrian Trajectory Prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15251–15260.
- [9] Thomas Kipf, Ethan Fetaya, Kuan-Chieh Wang, Max Welling, and Richard Zemel. 2018. Neural relational inference for interacting systems. In *International conference on machine learning*. Pmlr, 2688–2697.
- [10] Namhoon Lee, Wongun Choi, Paul Vernaza, Christopher B Choy, Philip HS Torr, and Manmohan Chandraker. 2017. Desire: Distant future prediction in dynamic scenes with interacting agents. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 336–345.
- [11] Seongju Lee, Junseok Lee, Yeonguk Yu, Taeri Kim, and Kyoobin Lee. 2024. Mart: Multiscale relational transformer networks for multi-agent trajectory prediction. In *European Conference on Computer Vision*. Springer, 89–107.
- [12] Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. 2007. Crowds by example. In *Computer graphics forum*, Vol. 26. Wiley Online Library, 655–664.
- [13] Somaiyeh MahmoudZadeh, Amirmehdi Yazdani, Atabak Elmi, Amin Abbasi, and Pooria Ghanooni. 2022. Exploiting a fleet of UAVs for monitoring and data acquisition of a distributed sensor network. *Neural Computing and Applications* (2022), 1–14.
- [14] Weibo Mao, Chenxin Xu, Qi Zhu, Siheng Chen, and Yanfeng Wang. 2023. Leapfrog Diffusion Model for Stochastic Trajectory Prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 5517–5526.
- [15] Andrew Ng, Michael Jordan, and Yair Weiss. 2001. On spectral clustering: Analysis and an algorithm. *Advances in neural information processing systems* 14 (2001).
- [16] Pedro Reyes Norambuena, Juan Bekios-Calfa, and Javier Martínez Torres. 2023. Study of machine learning techniques for pedestrian dynamics simulation models. In *2023 IEEE 13th International Conference on Pattern Recognition Systems (ICPRS)*. 1–8. <https://doi.org/10.1109/ICPRS58416.2023.10179045>
- [17] Stefano Pellegrini, Andreas Ess, Konrad Schindler, and Luc Van Gool. 2009. You'll never walk alone: Modeling social behavior for multi-target tracking. In *2009 IEEE 12th international conference on computer vision*. IEEE, 261–268.
- [18] David MW Powers. 2020. Evaluation: from precision, recall and F-measure to ROC, informedness, markedness and correlation. *arXiv preprint arXiv:2010.16061* (2020).
- [19] Alberto Pretto, Stéphanie Aravecchia, Wolfram Burgard, Nived Chebrolo, Christian Dornhege, Tillmann Falck, Freya Fleckenstein, Alessandra Fontenla, Marco Imperoli, Raghav Khanna, Frank Liebisch, Philipp Lottes, Andres Milioto, Daniele Nardi, Sandro Nardi, Johannes Pfeifer, Marija Popović, Ciro Potena, Cédric Pradalier, Elisa Rothacker-Feder, Inkyu Sa, Alexander Schaefer, Roland Siegwart, Cyrill Stachniss, Achim Walter, Wera Winterhalter, Xiaolong Wu, and Juan Nieto. 2021. Building an Aerial-Ground Robotics System for Precision Farming: An Adaptable Solution. *IEEE Robotics & Automation Magazine* 28, 3 (2021), 29–49. <https://doi.org/10.1109/MRA.2020.3012492>
- [20] Craig W. Reynolds. 1987. Flocks, Herds and Schools: A Distributed Behavioral Model. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*. ACM, New York, NY, USA.
- [21] Angela Ribeiro and Jesus Conesa-Muñoz. 2021. Multi-robot systems for precision agriculture. *Innovation in Agricultural Robotics for Precision Agriculture: A Roadmap for Integrating Robots in Precision Agriculture* (2021), 151–175.
- [22] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. 2020. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *European Conference on Computer Vision*. Springer, 683–700.
- [23] Alexander Strehl and Joydeep Ghosh. 2002. Cluster ensembles—a knowledge reuse framework for combining multiple partitions. *Journal of machine learning research* 3, Dec (2002), 583–617.
- [24] Emilian Vlasceanu, Dan Popescu, and Loretta Ichim. 2019. Aerial Robotic Team for Complex Monitoring in Precision Agriculture. In *2019 15th International Conference on Distributed Computing in Sensor Systems (DCOSS)*. 167–169.
- [25] Conghao Wong, Beihao Xia, Ziqian Zou, Yulong Wang, and Xinge You. 2024. SocialCircle: Learning the Angle-based Social Interaction Representation for Pedestrian Trajectory Prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 19005–19015.
- [26] Chenxin Xu, Maosen Li, Zhenyang Ni, Ya Zhang, and Siheng Chen. 2022. GroupNet: Multiscale Hypergraph Neural Networks for Trajectory Prediction with Relational Reasoning. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [27] Chenxin Xu, Robby T. Tan, Yuhong Tan, Siheng Chen, Xinchao Wang, and Yanfeng Wang. 2023. Auxiliary Tasks Benefit 3D Skeleton-based Human Motion Prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 9509–9520.
- [28] Ömür Yildirim, Klaus Diepold, and Revna Acar Vural. 2019. Decision Process of Autonomous Drones for Environmental Monitoring. In *2019 IEEE International Symposium on Innovations in Intelligent Systems and Applications (INISTA)*. 1–6. <https://doi.org/10.1109/INISTA.2019.8778341>
- [29] Ye Yuan, Xinshuo Weng, Yanglan Ou, and Kris M Kitani. 2021. Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting. In *Proceedings of the IEEE/CVF international conference on computer vision*. IEEE, Piscataway, NJ, USA, 9813–9823.