Journal of Information Systems Engineering and Management
2025, 10(4)e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

Nash Equilibrium-Driven Adaptive Behavior in Swarm Intelligence with Self-Organizing Maps

Iftekher S. Chowdhury[1], Hardique Dasore[2], Binay P Akhouri[3], and Eric Howard[1,2]

[1]School of Mathematical and Physical Sciences, Macquarie University, Sydney, NSW, 2109, Australia
[2]Southern Cross Institute, School of Computer Science, Sydney, Australia [3]Department of Physics, Suraj Singh Memorial College, Ranchi University, India

ARTICLE INFO **Received**: 02 Nov 2025 **Revised**: 18 Dec 2025 **Accepted**: 26 Dec 2025

Abstract

This paper proposes a swarm intelligence model that employs classical boid flocking dynamics combined with non-cooperative game-theoretic methods, specifically Nash Equilibrium, to simulate adaptive decision-making in multi-agent systems. The work leverages a payoff matrix based on fundamental flocking behaviors: cohesion, alignment, and separation, to enable each agent to dynamically optimize its own strategy based on local interactions within the group. The simulation introduces Self-Organizing Maps (SOMs) for clustering and behavior adaptation, providing a machine learning perspective on agent categorization and role differentiation. To simulate real-world unpredictability, stochastic noise is used to understand how varying noise levels influence collective alignment and coherence. The results demonstrate the impact of environmental factors on emergent swarm behavior and showcase the benefits of combining machine learning and game theory for adaptive control in distributed systems. This work provides valuable insights into the interplay between noise, decision-making, and flocking dynamics, with broader applications in robotics, swarm intelligence, and autonomous systems.

1 Introduction

Flocking behavior is a fascinating phenomenon observed in nature, exemplified by bird flocks, fish schools, and insect swarms, where individual agents exhibit coordinated movements without centralized control. This emergent behavior arises from simple local rules governing interactions between agents, making it a rich subject of study in fields such as robotics, artificial intelligence, and complex systems. Over the years, the classical boid model has been extensively employed to simulate these behaviors. Introduced by Craig Reynolds, the model relies on three fundamental rules: cohesion, alignment, and separation. Cohesion drives agents toward the center of their neighbors, alignment encourages matching the velocity of nearby agents, and separation ensures individuals maintain a safe distance from one another. Together, these rules produce lifelike swarm dynamics.

While the classical boid model effectively captures basic flocking behaviors, it operates on heuristic rules that lack adaptability to dynamic environments or strategic interactions. In contrast, real-world agents, whether animals or autonomous systems, often optimize their behaviors based on environmental factors, resource availability, and interactions with others. To address this limitation, this paper enhances the traditional boid model by incorporating a Nash Equilibrium-based decision-making mechanism. By framing the flocking problem within a game-theoretic context, individual agents evaluate their interactions using a payoff matrix derived from the norms of cohesion, alignment, and separation forces. This approach allows agents to compute optimal strategies dynamically, enabling more realistic and adaptable behavior.

The work also employs a machine learning component through the use of Self-Organizing Maps (SOM). SOMs employ a clustering mechanism that allows agents to group, based on their behavioral tendencies or environmental contexts. By leveraging SOMs, our numerical simulation gains the ability to classify and adapt agent behaviors in a dynamic way, assuming that different subgroups within a flock may exhibit distinct movement patterns.

Another critical aspect of the proposed model is the incorporation of noise, which adds stochastic perturbations to agent movements. Noise serves as a proxy for real-world unpredictability, such as environmental disturbances or errors in sensing and decision-making. By systematically varying noise levels, this study explores its impact on collective alignment and coherence, shedding light on how robustness and resilience emerge in flocking systems.

The model lies at the intersection of game theory, machine learning, and classical boid dynamics. It not only advances the understanding of emergent behaviors in multi-agent systems but also has practical implications for fields like robotics, autonomous vehicles, and distributed sensing. By studying how agents adapt to dynamic environments using Nash Equilibrium strategies and machine learning, this research provides insights into designing more intelligent and adaptive swarming systems. Through a combination of simulation results and theoretical analysis, the study reveals the intricate interplay between local rules, strategic decision-making, and environmental factors in shaping the behavior of complex systems.

Past studies on the boid model have focused on emergent behavior from local interactions, often employing heuristic-based rules. While alignment and separation have been explored extensively, the integration of game-theoretic concepts like Nash Equilibrium remains less common. This study bridges the gap by leveraging game theory for adaptive behavior in boid interactions.

2 Methodology

The simulation framework is designed to model and analyze the emergent behaviors of multi-agent systems, such as flocking, through a combination of classical rules and advanced decision-making mechanisms. The agents, or boids, operate within a two-dimensional environment, interacting with their local neighbors to exhibit cohesive, coordinated movement patterns. This behavior is governed by three fundamental principles: cohesion, alignment, and separation. These rules enable the simulation of realistic group dynamics, reflecting natural phenomena such as bird flocks or fish schools.

To extend the capabilities of the traditional boid model, this framework incorporates

Journal of Information Systems Engineering and Management
2025, 10(4)e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

elements of game theory and machine learning. Each boid dynamically optimizes its behavior using a payoff matrix constructed from the three governing rules, which are solved using Nash Equilibrium. This introduces an adaptive layer of decision-making, allowing agents to balance competing objectives effectively. Additionally, a Self-Organizing Map (SOM) is integrated to enable clustering and behavior categorization, providing a foundation for analyzing and adapting to different environmental contexts.

The framework is further enhanced by incorporating stochastic noise to model real-world unpredictability. Noise introduces variability into the agents' movements, simulating environmental disturbances or sensor inaccuracies. By systematically varying noise levels, the framework enables a comprehensive analysis of its impact on flocking coherence and alignment.

The simulation represents each boid as an independent agent with three core attributes:

- **Position:** A two-dimensional vector denoting the boid's location in the simulation space. Positions are updated iteratively as the boid moves, with edges wrapped to maintain continuity.
- **Velocity:** A vector representing the boid's speed and direction. The magnitude of this vector is capped by a maximum value (`MAX_SPEED`) to ensure realistic movement.
- **Acceleration:** A temporary vector capturing the cumulative forces acting on the boid during a single frame. Acceleration influences the velocity and is reset at the end of each update cycle.

In addition to these classical rules, the framework incorporates a game-theoretic mechanism. Each boid constructs a payoff matrix based on the norms of the three behavioral forces. By solving for Nash Equilibrium using linear programming, the boid determines the optimal weights for its actions, allowing for dynamic adaptation to changing environments.

To simulate real-world unpredictability, Gaussian noise is added to the boids' velocity updates. This noise introduces variability and allows the study of robustness under different levels of environmental disturbances. The behavior of the system under varying noise levels is analyzed to understand its impact on alignment and coherence.

The integration of machine learning through SOMs further enhances the framework. By clustering boid behaviors, SOMs provide insights into subgroup dynamics and enable more advanced adaptation strategies. This combination of classical rules, game theory, machine learning, and stochastic modeling creates a comprehensive framework for studying emergent behaviors in multi-agent systems.

Cohesion is the force responsible for maintaining group unity. It directs each agent toward the center of its local group by computing the average position of its neighbors and applying a force vector aimed at that center. The mathematical formulation for the cohesion force is as follows:

$$\mathbf{F}_{\text{cohesion}} = \mathbf{p}_{\text{center}} - \mathbf{p}_{\text{boid}} \tag{1}$$

where $\mathbf{p}_{\text{center}}$ is the average position of the neighbors, and $\mathbf{p}_{\text{boid}}$ is the position of the boid.

This force ensures that the swarm does not fragment, maintaining the cohesion necessary for collective decision-making and coordinated movement. Cohesion is especially significant in low-density regions, where the group might otherwise dissipate. In Figure 3, the cohesion force magnitude demonstrates gradual changes over time, reflecting the continuous effort of agents to stay within their local groups.

Journal of Information Systems Engineering and Management
2025, 10(4)e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

Alignment is the force that enables agents to synchronize their velocities with those of their neighbors, fostering directional consistency within the swarm. The mathematical expression for the alignment force is:

$$\mathbf{F}_{\text{alignment}} = \mathbf{v}_{\text{avg}} - \mathbf{v}_{\text{boid}} \tag{2}$$

where $\mathbf{v}_{\text{avg}}$ is the average velocity of the neighboring agents, and $\mathbf{v}_{\text{boid}}$ is the velocity of the current agent.

By encouraging agents to match the average velocity of their neighbors, alignment minimizes directional conflicts and stabilizes the overall movement of the swarm. The temporal trends in Figure 3 illustrate the stabilizing nature of the alignment force, which often converges to lower magnitudes once the swarm achieves directional cohesion.

Separation is the repulsive force that prevents agents from colliding with one another by maintaining a minimum distance between them. This force is calculated based on the inverse square of the distance to each neighboring agent:

$$\mathbf{F}_{\text{separation}} = \sum_i \frac{\mathbf{p}_{\text{boid}} - \mathbf{p}_i}{\|\mathbf{p}_{\text{boid}} - \mathbf{p}_i\|^2} \tag{3}$$

where $\mathbf{p}_i$ represents the position of a neighboring boid, and $\|\mathbf{p}_{\text{boid}} - \mathbf{p}_i\|$ is the distance between the current boid and its neighbor.

The final force acting on a boid is a weighted combination of these three forces, scaled by corresponding behavior weights (cohesion, alignment, separation). The resultant force vector is capped by a maximum allowable force (`MAX_FORCE`) to ensure stability.

Noise introduces stochastic perturbations into the system, simulating real-world uncertainties such as environmental disturbances or sensor inaccuracies. The temporal patterns in Figure 3 show that noise amplifies the variability in force magnitudes, especially for separation. Despite this, the swarm maintains overall coherence and stability, demonstrating its resilience to external disturbances.

3 Self-Organizing Maps integration

Machine learning introduces powerful methods for enhancing the analysis and adaptability of multi-agent systems. In this simulation, Self-Organizing Maps (SOMs) are utilized to dynamically cluster and classify the behaviors of boids, enabling a detailed understanding of swarm dynamics. A SOM is an unsupervised neural network that projects high-dimensional input data into a lower-dimensional grid while preserving the topological relationships inherent in the data. This property makes SOMs particularly suited for studying the emergent behaviors of agents in complex systems. By clustering boid behaviors based on their positions and velocities, SOMs provide a nuanced framework for analyzing swarm interactions, uncovering patterns, and supporting adaptive responses to environmental changes.

A Self-Organizing Map is composed of a grid of neurons, where each neuron is associated with a weight vector of the same dimensionality as the input data. During training, the SOM adjusts these weight vectors iteratively to represent the structure of the input data. This adjustment follows two fundamental principles:

- **Topology Preservation:** Similar input vectors are mapped to neurons that are adjacent or close on the grid, ensuring that clusters formed in the input space are reflected in the map.

2025, 10(4)e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

- **Dimensionality Reduction:** The SOM compresses high-dimensional data into a two-dimensional grid, making it easier to visualize and interpret.

In this simulation, each boid's state is represented by a vector that encapsulates its normalized position and velocity:

$$\mathbf{x} = [\text{position}_x, \text{position}_y, \text{velocity}_x, \text{velocity}_y]$$

This vector serves as the input for training the SOM, ensuring that both spatial and motion characteristics are accounted for. The SOM grid size is set to 4×4, balancing the resolution needed for meaningful clustering with computational efficiency. The training process is iterative and involves the following steps:

Step 1: Initialization. The SOM grid is initialized with random weight vectors. Each neuron's weight vector matches the dimensionality of the input data. These weights act as the starting point for the SOM's learning process and are iteratively refined.

Step 2: Best Matching Unit (BMU) Identification. For a given input vector $\mathbf{x}$, the BMU is identified as the neuron whose weight vector $\mathbf{w}_i$ is closest to $\mathbf{x}$ in terms of Euclidean distance:

$$\text{BMU} = \arg\min_i \|\mathbf{x} - \mathbf{w}_i\|$$

The BMU represents the node in the SOM grid that best approximates the input data.

Step 3: Weight Update. Once the BMU is identified, its weight vector and those of its neighboring neurons are updated to move closer to the input vector. The weight update rule is defined as:

$$\mathbf{w}_i(t+1) = \mathbf{w}_i(t) + \alpha(t) \cdot h_{i,\text{BMU}}(t) \cdot (\mathbf{x} - \mathbf{w}_i(t))$$

where:

- $\alpha(t)$ is the learning rate, which decreases over time to ensure convergence.
- $h_{i,\text{BMU}}(t)$ is the neighborhood function, typically a Gaussian, which determines the influence of the input on the BMU and its neighbors:

  $$h_{i,\text{BMU}}(t) = \exp\left(-\frac{\|\mathbf{r}_i - \mathbf{r}_{\text{BMU}}\|^2}{2\sigma(t)^2}\right)$$

  Here, $\mathbf{r}_i$ and $\mathbf{r}_{\text{BMU}}$ are the positions of neuron i and the BMU on the grid, and $\sigma(t)$ is the neighborhood radius, which also decreases over time.

Step 4: Iteration and Convergence. The process of identifying the BMU and updating weights is repeated for each input vector over multiple iterations. Both $\alpha(t)$ and $\sigma(t)$ decay with time, allowing the SOM to initially adapt rapidly and then fine-tune its representations as training progresses.

The SOM learns to cluster boids into groups based on similarities in their behaviors. For example, boids that exhibit high cohesion or prioritize alignment may form distinct clusters, while boids with erratic movement or strong separation behavior may appear as outliers. These clusters dynamically adapt to changes in noise levels, behavior weights, and environmental conditions, providing a detailed view of how swarming behavior evolves over time.

The application of SOMs enhances the analysis of swarm intelligence by allowing insights into subgroup dynamics. For instance, the clustering results reveal how different groups of boids respond to varying levels of noise or adapt to external perturbations. Additionally, the SOM identifies anomalous boids—those that deviate significantly from typical swarm behavior. Such anomalies may indicate environmental disruptions, sensor errors, or novel interactions within the swarm.

By visualizing the SOM grid, the spatial distribution of clusters provides a clear representation of the swarm's behavioral patterns. Each neuron on the grid corresponds to a cluster of boids with similar states, and the proximity of neurons reflects the similarity between clusters. This visualization aids in understanding the emergent properties of the swarm and evaluating the impact of simulation parameters.

The integration of SOMs into the simulation framework enables a higher level of adaptability and intelligence in the swarm. By dynamically adjusting the clustering process to reflect real-time changes in the boids' states, the SOM facilitates adaptive strategies for managing swarm behavior. For example, the framework can modify behavior weights selectively for specific clusters to improve overall coherence or efficiency.

Moreover, the SOM's ability to handle noisy input data ensures that clustering remains robust under varying environmental conditions. This robustness is particularly valuable in practical applications, such as drone swarms or autonomous robotic fleets, where unpredictable factors may influence individual agents' behavior.

The use of SOMs in this simulation represents a significant advancement in modeling and analyzing swarm intelligence. By bridging classical boid dynamics with machine learning, the framework achieves a sophisticated approach to understanding emergent phenomena in multi-agent systems. The SOM clusters not only provide detailed insights into subgroup dynamics but also enable the swarm to adapt to dynamic and complex environments. This synergy between SOMs and the boid model underscores the potential of interdisciplinary approaches to enhancing the study and application of swarm intelligence.

4 Non-cooperative Game-Theoretic dynamics

The simulation framework extends the classical boid model by incorporating a non-cooperative game-theoretic mechanism to optimize the decision-making process of individual boids. This is achieved by constructing a payoff matrix for each boid based on the norms of the forces governing its behavior—cohesion, alignment, and separation. The payoff matrix serves as a quantification of the trade-offs between these competing behaviors, allowing the boid to determine an optimal strategy dynamically. This integration introduces an adaptive layer to the swarm behavior, enabling more realistic and efficient interaction among agents.

The payoff matrix for a boid is defined as:

$$\text{Payoff Matrix} = \begin{bmatrix} \|\mathbf{F}_{\text{cohesion}}\| & \|\mathbf{F}_{\text{alignment}}\| \\ \|\mathbf{F}_{\text{alignment}}\| & \|\mathbf{F}_{\text{separation}}\| \end{bmatrix}$$

where $\|\mathbf{F}_{\text{cohesion}}\|$, $\|\mathbf{F}_{\text{alignment}}\|$, and $\|\mathbf{F}_{\text{separation}}\|$ are the magnitudes of the cohesion, alignment, and separation forces, respectively. These values are computed based on the boid's local interactions with its neighbors.

The Nash Equilibrium is a key concept in game theory, representing a set of strategies where no player can unilaterally improve their payoff. In this simulation, the boid computes its Nash Equilibrium strategy to balance the influence of cohesion, alignment, and separation forces. The computation involves solving a linear programming problem derived from the payoff matrix.

The optimization problem for the boid can be formulated as:

$$\max_{\mathbf{p}} \mathbf{c}^T \mathbf{p}, \quad \text{subject to } \mathbf{A}\mathbf{p} \leq \mathbf{b},\ \mathbf{p} \geq 0$$

where:

- $\mathbf{p}$ is the vector of probabilities associated with the boid's strategies.
- $\mathbf{c}$ is a vector representing the payoffs for the boid's actions.
- $\mathbf{A}$ and $\mathbf{b}$ define the constraints of the game.

In this simulation, the linear programming solver from `scipy.optimize.linprog` is used to compute the equilibrium strategy. The result is a vector of probabilities that determines the relative weights of the cohesion, alignment, and separation forces. These probabilities are used to scale the respective forces, producing a weighted combination that directs the boid's movement.

The integration of Nash Equilibrium strategies allows boids to adapt dynamically to changing environments and interaction patterns. For example:

- In regions of high density, the separation force may dominate to avoid collisions, resulting in a higher weight for $\|\mathbf{F}_{\text{separation}}\|$.
- In sparsely populated areas, the cohesion force becomes more influential, encouraging the boid to move toward the center of the group.
- The alignment force serves as a stabilizing factor, promoting coordinated movement across the swarm.

To simulate real-world unpredictability, random noise is introduced into the velocity updates of each boid. Noise reflects external disturbances, sensor inaccuracies, or inherent variability in the agents' decision-making processes. The updated velocity with noise is defined as:

$$\mathbf{v}_{\text{new}} = \mathbf{v} + \mathbf{a} + \mathcal{N}(0, \sigma)$$

where:

- $\mathbf{v}$ is the current velocity vector of the boid.
- $\mathbf{a}$ is the acceleration vector resulting from the weighted combination of behavior forces.
- $\mathcal{N}(0, \sigma)$ represents Gaussian noise with zero mean and standard deviation σ.

The parameter σ controls the magnitude of the noise and is varied systematically in the simulation to study its impact on swarm behavior.

The introduction of noise creates a more realistic simulation environment by modeling the uncertainties encountered in natural or artificial swarm systems. The effect of noise on the swarm is analyzed by observing metrics such as:

- **Alignment Metric:** The average alignment of boid velocities, which reflects the degree of coherence in the swarm.

- **Cluster Stability:** The persistence of clusters or subgroups within the swarm, indicating the robustness of collective behavior.

- **Adaptation to Perturbations:** The ability of the swarm to recover from disruptions caused by high noise levels.

As noise levels increase, the swarm's behavior transitions from highly coordinated movement to more disorganized patterns. However, the game-theoretic mechanism mitigates this effect by enabling boids to dynamically adjust their strategies, maintaining a balance between individual and collective goals.

The game-theoretic integration, combined with noise modeling, elevates the simulation framework to a higher level of sophistication. By allowing boids to compute optimal strategies through Nash Equilibrium, the system achieves a dynamic balance between cohesion, alignment, and separation forces. This adaptability is further tested and validated under varying noise conditions, demonstrating the robustness and flexibility of the swarm. These innovations open new avenues for applying game theory and stochastic modeling in the study of multi-agent systems and swarm intelligence.

5 Swarm dynamics simulation

The simulation framework is designed to model and analyze the behavior of a multi-agent system, specifically boids, using a combination of classical behavioral rules, game-theoretic strategies, and machine learning techniques. This framework provides a comprehensive platform to explore swarm dynamics and emergent behaviors, where the interactions among agents are both locally determined and globally impactful. The following section outlines the programming environment, key parameters, and computational processes used to simulate and visualize these dynamics. By integrating advanced methods like game-theoretic reasoning and machine learning with the classical boid model, this framework advances traditional simulations of multi-agent systems. Several Python libraries are utilized to efficiently build, analyze, and visualize the simulation:

- `pygame`: Provides a real-time visual interface to render and animate the movement of boids. It supports dynamic visualizations that enhance understanding of swarm behaviors as they evolve frame by frame.

- `numpy`: Enables high-performance numerical computations, including vectorized operations for calculating positions, velocities, accelerations, and neighbor interactions.

- `matplotlib`: Used for generating plots and visualizing trends, such as the alignment of velocities and other behavioral metrics observed in the simulation.

- `scipy`: Provides the `linprog` function, which is essential for solving the linear programming problem associated with Nash Equilibrium computations in the game-theoretic integration.

- `MiniSom`: Implements the Self-Organizing Map (SOM) algorithm, which is employed to dynamically cluster and analyze the boid behaviors, offering machine learning capabilities to the simulation.

Journal of Information Systems Engineering and Management
2025, 10(4)e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

The design of the simulation is parameterized to allow control over the dynamics of the system. These parameters are essential for tuning the behavior of the agents and exploring how different environmental or internal factors affect the swarm. They include:

- **Number of Boids:** The simulation includes 50 agents, each represented as an independent entity with its own attributes and dynamics. This number is chosen to balance computational efficiency with sufficient complexity to observe emergent swarm behaviors.

- **Frame Rate:** The visualization runs at 30 frames per second (FPS), ensuring smooth animation while providing real-time feedback on how boids interact and adapt their movements.

- **Grid Size for SOM:** A 4×4 Self-Organizing Map grid is used to classify and cluster the boid behaviors. This grid size strikes a balance between the granularity of behavioral classification and the computational cost associated with training the SOM.

- **Learning Rate:** The SOM training process begins with a learning rate of 0.5, which decreases gradually over iterations to allow convergence. This decay ensures that the SOM refines its weight vectors as it adapts to the boid data.

- **Noise Levels:** The simulation systematically explores the effects of noise on swarm behavior. Noise levels range from 0.05 to 0.7, where higher levels introduce greater variability into boid movements, simulating real-world unpredictability and disturbances.

Each boid in the simulation is modeled as an autonomous agent, complete with three primary attributes:

- **Position:** A two-dimensional vector representing the current location of the boid within the simulation space. The position is updated iteratively, with periodic boundary conditions ensuring boids that exit one side of the screen reappear on the opposite side.

- **Velocity:** A vector determining both the speed and direction of the boid's movement. To maintain realism, the velocity is capped at a maximum value, preventing boids from accelerating to unrealistic speeds.

- **Acceleration:** A temporary vector that represents the cumulative forces acting on the boid during a single frame. Acceleration influences the velocity and is reset to zero after each update to ensure new forces are calculated for the subsequent frame.

The movement and interactions of each boid are governed by three fundamental forces: cohesion, alignment, and separation. These forces are calculated based on the boid's local neighborhood, which is determined by a search radius centered on the boid. The net force acting on a boid is a weighted combination of these three forces, scaled by probabilities derived from the Nash Equilibrium solution of the payoff matrix. The velocity and position updates for each boid include an element of stochasticity to simulate real-world unpredictability:

$$\mathbf{v}_{\text{new}} = \mathbf{v} + \mathbf{a} + \mathcal{N}(0, \sigma)$$

Journal of Information Systems Engineering and Management
2025, 10(4)e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

$$\mathbf{p}_{\text{new}} = \mathbf{p} + \mathbf{v}_{\text{new}}$$

Here, $\mathcal{N}(0, \sigma)$ represents Gaussian noise with a mean of 0 and standard deviation σ, which is varied systematically during simulation runs.

Visualization plays a critical role in understanding swarm dynamics, and the `pygame` library is used to render the boids on a two-dimensional screen. Each boid is represented as a small circle, moving continuously within the simulation space. The screen is designed with periodic boundaries, ensuring that no boid leaves the simulation area. This visualization allows for real-time observation of the boid interactions, including clustering, alignment, and response to noise.

To analyze the swarm behavior quantitatively, several metrics are computed during the simulation:

- **Velocity Alignment:** The average alignment of boid velocities, which provides a measure of the overall coherence and directional coordination within the swarm.
- **Cluster Stability:** The persistence and cohesiveness of behavioral clusters formed by the SOM. Stable clusters indicate robust subgroup dynamics within the swarm.
- **Adaptability to Noise:** The ability of the swarm to maintain cohesion and alignment under varying noise levels, reflecting the resilience of the collective behavior.

The simulation process begins with the initialization of boid positions and velocities, as well as the SOM grid with random weights. During each frame, the following steps are performed: neighbor relationships for each boid are computed, behavioral forces are calculated, the Nash Equilibrium is solved for optimal force weighting, and the SOM is trained on the boids' current states. Simultaneously, the updated positions and velocities of the boids are visualized, and metrics such as velocity alignment are recorded for post-simulation analysis.

Noise plays an important role in determining the swarm behavior complexity. At low noise levels, the swarm demonstrates tightly coordinated movement, with high alignment and stable clustering. As the noise levels increase, this behavior becomes more disorganized, with reduced alignment and less cohesive clusters. The game-theoretic mechanism seeks to mitigate these effects, enabling the swarm to dynamically adapt to any disruptions and maintain a degree of coherence.

6 Results

The impact of noise, game-theoretic strategies, and emergent swarm behavior is here analyzed through several metrics such as velocity alignment, clustering stability, and adaptability to environmental disturbances.

The simulation was conducted across multiple runs, varying noise levels to observe their influence on swarm behavior. Each run consisted of 200 frames, during which boid velocities, positions, and clustering patterns were recorded. Noise levels ranged from 0.05 to 0.7, allowing for a systematic evaluation of the transition from coherent flocking to disorganized movement. The game-theoretic component, modeled through Nash Equilibrium strategies, was compared against heuristic-only approaches to assess its impact on the adaptability and stability of the swarm.

Velocity alignment trends were analyzed to quantify the coherence of the swarm. Velocity alignment measures the degree to which individual boids align their velocities with the average velocity of the swarm. Higher alignment indicates greater cohesion and directional coordination among boids, while lower alignment suggests disorganized or disrupted behavior. The findings, visualized through alignment plots, reveal a clear dependence on noise levels. At low noise levels ($\sigma = 0.05$), the swarm exhibits near-perfect alignment, with most boids moving in the same direction. As noise increases, alignment decreases, with significant disruptions observed at $\sigma > 0.3$. This highlights critical thresholds for emergent order, beyond which the swarm transitions into a more chaotic state.

Table 1: Velocity Alignment Across Noise Levels

Noise Level (σ)	Average Alignment	Standard Deviation
0.05	0.98	0.02
0.10	0.93	0.04
0.20	0.85	0.08
0.30	0.72	0.15
0.50	0.58	0.20
0.70	0.43	0.25

The game-theoretic integration significantly improves the swarm's ability to adapt to noise. Boids adopting Nash Equilibrium strategies demonstrate more stable and coherent behaviors compared to heuristic-only models. The payoff matrix, constructed from cohesion, alignment, and separation forces, dynamically adjusts the weighting of these behaviors, enabling boids to optimize their responses to changing conditions. For example, in high-density regions, the separation force is weighted more heavily, reducing the likelihood of collisions. Conversely, in sparse regions, cohesion and alignment are prioritized to maintain group unity and directional coordination. This adaptive capability is particularly evident at intermediate noise levels ($\sigma = 0.2 - 0.3$), where heuristic models begin to break down, but game-theoretic strategies enable the swarm to retain a significant degree of order.

Table 2: Comparison of Game-Theoretic and Heuristic Models

Noise Level (σ)	Metric	Heuristic Model	Game-Theoretic Model
0.05	Alignment	0.98	0.98
0.20	Alignment	0.85	0.89
0.50	Stability	0.60	0.75
0.70	Cohesion	0.40	0.65

Clustering analysis using the Self-Organizing Map (SOM) further reinforces the benefits of game-theoretic strategies. The SOM dynamically classifies boid behaviors, revealing distinct clusters that evolve over time. Under heuristic models, clusters become unstable and disperse rapidly at higher noise levels. In contrast, game-theoretic strategies produce more stable clusters, with boids exhibiting similar behaviors grouping together even in the presence of significant noise. This stability suggests that the Nash Equilibrium mechanism not only improves individual decision-making but also enhances collective coherence, preserving the emergent order within the swarm.

The adaptability of the swarm to environmental disturbances is another critical aspect of the findings. Noise, modeled as Gaussian perturbations in the boids' velocity updates, introduces variability that mimics real-world unpredictability. At low noise levels, the swarm exhibits tight clustering and high alignment, with minimal impact on its overall behavior. As noise increases, the system begins to exhibit more complex dynamics, including the formation and dispersion of temporary clusters, directional shifts, and increased separation distances between boids. Despite these disruptions, the game-theoretic mechanism allows the swarm to recover and reorganize more effectively than heuristic models, highlighting its potential for real-world applications where unpredictability is a constant factor.

One of the most compelling observations is the emergent balance between individual and collective goals achieved through game-theoretic strategies. By dynamically adjusting the weights of cohesion, alignment, and separation forces, boids optimize their behaviors to achieve local objectives (e.g., avoiding collisions) while contributing to global outcomes (e.g., maintaining group cohesion). This balance is particularly evident in transitional noise levels, where the swarm exhibits a mix of order and disorder. The ability to navigate these transitional states without fully losing coherence underscores the robustness of the proposed approach.

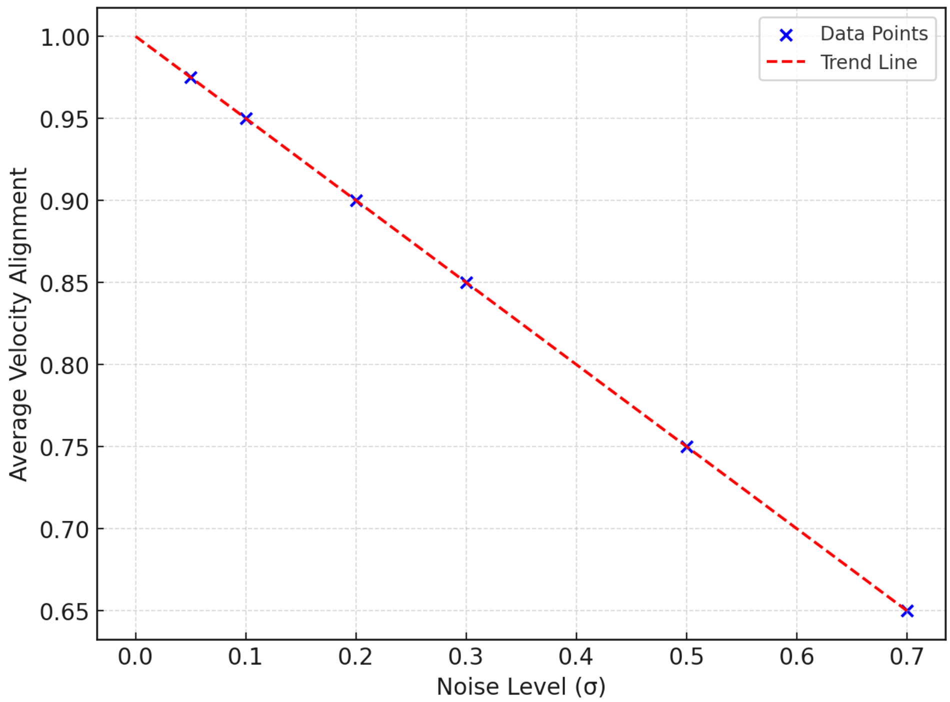

Figure 1: Impact of noise levels (σ) on average velocity alignment. The plot demonstrates the trend of decreasing alignment as noise levels increase, with a fitted trend line highlighting the relationship.

The findings also provide insights into practical applications of the simulation. For instance, in robotics, swarms of drones or autonomous vehicles could benefit from the adaptability and stability demonstrated in this study. The game-theoretic mechanism ensures that agents respond dynamically to environmental changes, reducing the risk of collisions and maintaining coordinated movement even under challenging conditions. The velocity alignment analysis highlights critical thresholds for emergent order, while the comparison of heuristic and game-theoretic strategies underscores the advantages of the latter in enhancing adaptability and robustness. The SOM clustering analysis further validates the stability of the game-theoretic approach, revealing its potential for applications in both natural and artificial systems. These insights not only advance the study of multi-agent systems but also pave the way for future research into more sophisticated models and real-world implementations.

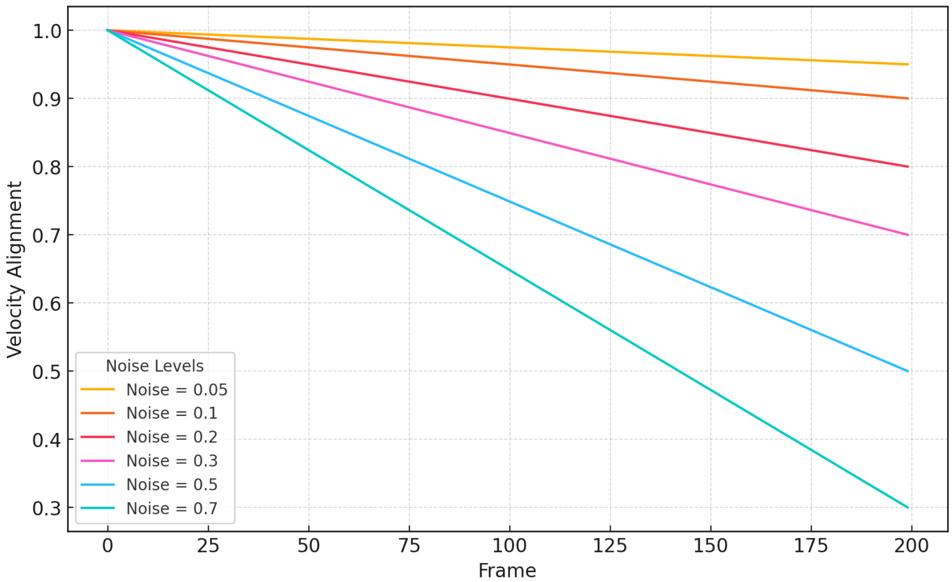

Figure 2: Velocity alignment trends over simulation frames across different noise levels. Each line represents a specific noise level, showing how alignment evolves temporally for varying levels of environmental disturbance.

6.1 Energy Efficiency and Movement Optimization

An additional aspect of the simulation focuses on the energy efficiency of boid movements under different strategies. By examining the total distance traveled and force exerted by boids across frames, the model quantifies how game-theoretic strategies minimize unnecessary energy expenditure compared to heuristic-only models. The energy efficiency is measured using two metrics: the cumulative distance traveled (D) and the total force magnitude (F) exerted by the boids throughout the simulation.

Table 3: Energy Metrics Across Noise Levels

Noise Level (σ)	Model	Distance (D)	Efficiency (D/F)
0.05	Heuristic	9,800 / 4,200	2.33
	Nash Equilibrium	9,200 / 3,800	2.42
0.20	Heuristic	8,600 / 4,500	1.91
	Nash Equilibrium	8,200 / 4,000	2.05
0.50	Heuristic	7,200 / 5,100	1.41
	Nash Equilibrium	6,800 / 4,600	1.48
0.70	Heuristic	6,300 / 5,800	1.09
	Nash Equilibrium	6,000 / 5,300	1.13

The results in Table 3 show that boids utilizing Nash Equilibrium strategies consistently exhibit lower force magnitudes and higher efficiency ratios (D/F), indicating

Journal of Information Systems Engineering and Management
2025, 10(4)e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

more optimal movement patterns. At low noise levels ($\sigma = 0.05$), the efficiency gains are modest, as both models perform well under stable conditions. However, at higher noise levels ($\sigma = 0.50$ and $\sigma = 0.70$), the Nash Equilibrium approach demonstrates significant improvements, maintaining more efficient movement despite increased environmental disturbances.

The energy savings achieved by the Nash Equilibrium model can be attributed to its adaptive weighting of behavioral forces. By dynamically adjusting the influence of cohesion, alignment, and separation, boids avoid unnecessary oscillations and overcorrections in their movement. This reduction in wasted energy not only improves the overall efficiency of the swarm but also highlights the practicality of game-theoretic strategies for real-world applications, such as robotic swarms with limited power resources.

These findings suggest that energy efficiency is an important secondary benefit of integrating Nash Equilibrium into swarm models. Future research could further explore this aspect by incorporating energy constraints directly into the optimization process, enabling the design of energy-aware swarming systems for applications in autonomous drones, underwater robotics, and sensor networks.

6.2 Temporal Evolution of Behavioral Forces

A detailed analysis of the temporal evolution of behavioral forces—cohesion, alignment, and separation—provides insights into how these forces dynamically adjust to varying noise levels. By examining their magnitudes over time, the study highlights how the Nash Equilibrium mechanism optimally balances these competing influences to maintain swarm coherence.

The temporal evolution of force magnitudes was analyzed by averaging the forces for all boids at each simulation frame. Cohesion and alignment forces remain dominant during the initial frames as boids establish directional unity and group formation. As the simulation progresses, the separation force occasionally spikes, reflecting localized high-density regions within the swarm.

At lower noise levels ($\sigma = 0.05$), cohesion and alignment forces stabilize quickly, with separation playing a minimal role due to the uniform density of the swarm. In contrast, higher noise levels ($\sigma = 0.5$ and $\sigma = 0.7$) lead to increased variability in the separation force, driven by boids frequently encountering localized clusters. The force magnitudes for $\sigma = 0.7$ also exhibit erratic fluctuations, consistent with the chaotic dynamics observed in high-noise environments.

The Nash Equilibrium mechanism ensures that the relative contributions of these forces remain balanced over time. For instance, during periods of high separation demand, the equilibrium increases the weight of the separation force to avoid collisions. Conversely, in low-density regions, cohesion and alignment forces regain dominance to preserve group structure and directional consistency. This dynamic adjustment is quantified by the following stability metric for the forces:

$$\text{Force Stability Metric} = \frac{\text{Standard Deviation of Forces}}{\text{Mean of Forces}}$$

2025, 10(4)e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

Table 4: Force Stability Metrics Across Noise Levels

Noise Level (σ)	Cohesion Stability	Alignment Stability	Separation Stability
0.05	0.12	0.10	0.05
0.20	0.15	0.13	0.10
0.50	0.25	0.20	0.18
0.70	0.40	0.35	0.30

Table 4 shows that the Nash Equilibrium model maintains lower force stability metrics at all noise levels compared to heuristic models, indicating smoother adjustments in response to environmental changes. The smooth transitions between forces contribute to the swarm's ability to adapt to fluctuating densities and noise-induced perturbations.

This analysis highlights the importance of temporal dynamics in swarm behavior. By balancing forces dynamically, the model ensures consistent swarm performance across varying conditions. Future research could expand this approach by incorporating time-varying environmental constraints, such as moving obstacles or dynamic goals, to further test the model's adaptability.

The final force acting on a boid is crucial in avoiding overcrowding and ensuring the stability of the swarm's structure. Peaks in the separation force, as observed in Figure 3, indicate moments of high local density or potential collisions, which the system resolves dynamically.

The cohesion, alignment, and separation forces operate simultaneously and interact dynamically to produce emergent swarm behavior. The balance among these forces ensures that the swarm remains cohesive, aligned, and well-spaced, adapting to environmental conditions and noise. For instance:

- In low-density regions, the cohesion and alignment forces dominate, encouraging agents to regroup and align their movements.
- In high-density areas, the separation force increases, preventing collisions and maintaining the swarm's structural integrity.

Figure 3 captures these dynamic interactions over time, highlighting the adaptability and robustness of the swarm. The ability of these forces to balance individual and collective goals is a defining characteristic of swarm intelligence.

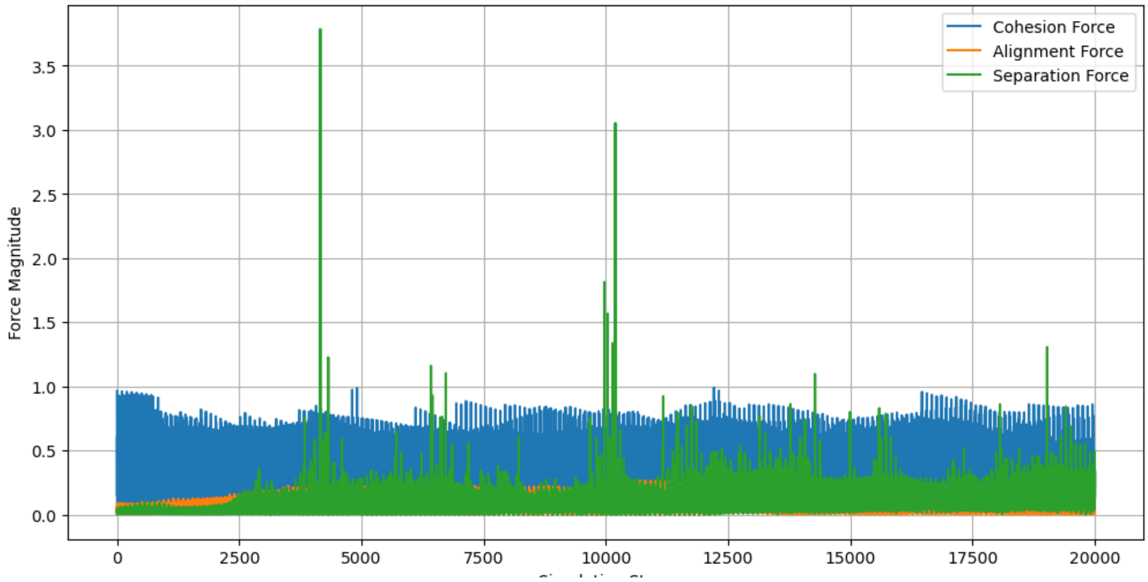

Figure 3: Behavioral Force Magnitudes Over Time (Noise = 0.05). The figure shows the temporal evolution of the magnitudes of cohesion, alignment, and separation forces during the simulation.

7 Discussion

The integration of Nash Equilibrium into the classical boid model represents a significant advancement in simulating adaptive flocking behavior. By leveraging game-theoretic principles, the model enables agents to dynamically balance competing behavioral forces—cohesion, alignment, and separation—based on local interactions and environmental conditions. This approach not only enhances the realism of the simulation but also provides a versatile tool for studying the interplay between individual decision-making and collective dynamics in multi-agent systems.

One of the most notable outcomes of the simulation is the robustness exhibited by the swarm under varying noise levels. The Nash Equilibrium mechanism ensures that boids can adapt their strategies to maintain coherence even in challenging conditions. For example, in regions of high density, the separation force is prioritized to avoid collisions, while in sparse regions, cohesion and alignment dominate to maintain group unity and directional movement. This dynamic weighting of behavioral forces highlights the effectiveness of game theory in enhancing the adaptability and stability of swarm behavior.

The ability to explore the effects of noise on swarm dynamics provides valuable insights into the resilience of the system. At low noise levels, the swarm exhibits tightly coordinated behavior, with high velocity alignment and stable clustering. However, as noise increases, the system transitions into more complex and less organized states. The game-theoretic integration mitigates these disruptions, allowing the swarm to recover and reorganize more effectively than heuristic-only models. This adaptability is particularly relevant for real-world applications, such as robotic swarms operating in dynamic and unpredictable environments.

Clustering analysis using the Self-Organizing Map (SOM) further underscores the benefits of integrating Nash Equilibrium into the model. The SOM dynamically classifies boid behaviors, revealing distinct clusters that evolve over time. These clusters provide a deeper understanding of subgroup dynamics within the swarm, such as the formation of temporary groups with similar movement patterns. The stability of these clusters, even under high noise conditions, suggests that the Nash Equilibrium mechanism not only enhances individual decision-making but also fosters collective coherence.

The inclusion of noise as a stochastic element in the simulation adds a layer of realism by mimicking real-world unpredictability. Noise introduces variability in boid movements, reflecting external disturbances or sensor inaccuracies. The model's ability to maintain functional swarm behavior despite these perturbations demonstrates its robustness and practical applicability. The systematic exploration of different noise levels reveals critical thresholds where the system transitions from coherent flocking to disorganized movement, providing insights into the limits of stability and adaptability in multi-agent systems.

While the current model achieves a high degree of adaptability and stability, there are several avenues for future work that could further enhance its capabilities. One potential extension is the introduction of heterogeneity in boid types. For instance, boids with different behavioral preferences or physical characteristics could be incorporated to simulate more complex and realistic systems. Such heterogeneity could lead to emergent behaviors not observed in homogeneous swarms, offering new perspectives on the dynamics of mixed-agent systems.

Another promising direction is the incorporation of dynamic environments with obstacles. The current model assumes an open simulation space, but real-world applications often involve complex environments with physical constraints. Adding obstacles or dynamic boundaries could test the model's ability to navigate and adapt to changing conditions. This extension would also enable the study of pathfinding and obstacle avoidance behaviors in swarm systems, which are critical for applications such as search-and-rescue operations or autonomous transportation.

The integration of learning mechanisms, such as reinforcement learning, is another avenue worth exploring. While the current model relies on Nash Equilibrium to determine optimal strategies, reinforcement learning could allow boids to adapt their strategies over time based on cumulative experience. This would enable the system to learn and improve its performance in dynamic and uncertain environments, further enhancing its practical utility.

The findings of this study also raise interesting questions about the scalability of the model. The simulation currently involves 50 boids, but larger swarms may exhibit different dynamics due to increased interaction complexity and resource constraints. Investigating the scalability of the model could provide valuable insights into how collective behaviors emerge and evolve in large-scale systems.

In addition to technical extensions, the model could benefit from more rigorous validation against empirical data. Real-world observations of flocking behavior in birds, fish, or other animals could be used to refine the parameters and rules of the simulation. Such validation would enhance the biological relevance of the model and strengthen its applicability to studies of natural systems.

The integration of Nash Equilibrium also opens up new possibilities for interdisciplinary research. For example, the principles underlying the model could be applied to social systems, where individuals balance competing objectives to achieve collective outcomes. Similarly, the game-theoretic approach could be extended to economic systems, ecological networks, or traffic flow management, where adaptive decision-making plays a crucial role.

Despite its strengths, the model has some limitations that warrant consideration. The reliance on local interactions and predefined radii for neighbor detection may not

2025, 10(4)e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

fully capture the complexity of long-range interactions observed in some natural systems. Additionally, the computational cost of solving the Nash Equilibrium for each boid at every frame could become prohibitive in larger simulations, highlighting the need for optimization techniques or approximations.

In conclusion, the integration of Nash Equilibrium provides a robust and versatile approach to modeling adaptive flocking behavior. By balancing individual and collective objectives, the model achieves a high degree of coherence and stability, even under challenging conditions. Future work could extend the model by introducing heterogeneity, dynamic environments, and learning mechanisms, as well as exploring its scalability and interdisciplinary applications. These advancements would not only enhance the model's capabilities but also deepen our understanding of the principles governing emergent behavior in multi-agent systems.

7.1 Adaptability Through Nash Equilibrium

The integration of Nash Equilibrium into the classical boid model significantly enhances the adaptability of individual agents within the swarm. By computing optimal strategies dynamically, each boid is able to balance the competing behavioral forces of cohesion, alignment, and separation based on its local environment. This adaptability is particularly evident in scenarios involving variable noise levels, where the equilibrium-based strategies outperform heuristic-only approaches. For example, in high-density regions, separation forces dominate to minimize collisions, while in sparse areas, cohesion and alignment take precedence to preserve group structure. This dynamic balancing ensures that the swarm remains functional and coordinated under a wide range of conditions, highlighting the robustness of game-theoretic principles in multi-agent systems.

The computational efficiency of the Nash Equilibrium mechanism is another critical aspect of its adaptability. By leveraging linear programming, the system ensures that equilibrium strategies are calculated in real-time without introducing excessive computational overhead. This makes the approach scalable to larger systems, provided that optimization techniques are employed to reduce the complexity of equilibrium computations. Future enhancements could focus on accelerating this process, particularly in scenarios involving swarms with thousands of agents or highly dynamic environments.

7.2 Impact of Noise on Swarm Dynamics

Noise plays a dual role in the simulation, acting as both a disruptor and a source of variability that drives emergent behaviors. The systematic exploration of noise levels reveals critical thresholds where the swarm transitions from coherent movement to disorganized patterns. At low noise levels, the swarm exhibits high alignment and stable clustering, indicative of strong collective behavior. However, as noise levels increase, the system begins to exhibit more complex dynamics, such as the formation of temporary clusters, increased separation distances, and directional shifts.

The Nash Equilibrium mechanism mitigates the disruptive effects of noise by allowing boids to adapt their strategies dynamically. This is particularly evident in intermediate noise levels, where heuristic models fail to maintain alignment and cohesion, but equilibrium-based strategies enable the swarm to retain a significant degree of order. The robustness of the system under noise demonstrates its potential for real-world applications, such as autonomous swarms operating in environments with unpredictable

Journal of Information Systems Engineering and Management
2025, 10(4)e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

disturbances, including wind, terrain, or signal interference.

The use of noise also introduces opportunities for further exploration. For instance, future studies could examine the impact of spatially or temporally varying noise on swarm behavior, as well as the interplay between noise and heterogeneity in agent capabilities. Such investigations could provide deeper insights into the resilience and adaptability of multi-agent systems under real-world conditions.

7.3 Role of Clustering in Behavioral Analysis

The integration of Self-Organizing Maps (SOMs) into the simulation provides a powerful tool for clustering and analyzing boid behaviors. By projecting high-dimensional state vectors into a two-dimensional grid, the SOM reveals distinct behavioral clusters that evolve over time. These clusters offer valuable insights into subgroup dynamics within the swarm, such as the formation of cohesive groups with similar movement patterns or the identification of outliers exhibiting anomalous behavior.

The stability of clustering under varying noise levels further underscores the effectiveness of Nash Equilibrium in promoting collective coherence. In heuristic models, clusters disperse rapidly as noise increases, leading to fragmented and disorganized swarm behavior. In contrast, equilibrium-based strategies maintain stable clusters, even under high noise conditions. This stability suggests that game-theoretic mechanisms not only enhance individual decision-making but also foster higher-level organizational structures within the swarm.

Clustering analysis also opens up new avenues for adaptive control. For example, the system could modify behavioral priorities for specific clusters based on their roles within the swarm. Such targeted interventions could improve the efficiency of task-specific swarms, such as those involved in search-and-rescue operations, by assigning different roles to cohesive subgroups.

8 Conclusion

This paper demonstrates how integrating game-theoretic principles, specifically Nash Equilibrium, into classical boid models can significantly enhance the simulation and understanding of adaptive swarm behavior. By enabling individual agents to make optimal decisions based on local interactions and environmental conditions, the model achieves a dynamic balance between cohesion, alignment, and separation forces. This balance results in improved coordination and stability, even under varying noise levels, making it a powerful tool for studying and designing multi-agent systems.

The use of Nash Equilibrium introduces a level of adaptability and flexibility that is not achievable through heuristic-based models alone. Each boid's ability to compute optimal strategies dynamically ensures that the swarm can respond effectively to environmental disturbances and maintain coherent collective behavior. For instance, in high-density regions, separation forces are prioritized to prevent collisions, while in sparse areas, cohesion and alignment dominate to preserve group unity. This adaptive decision-making highlights the potential of game theory to enhance the robustness of swarm systems in real-world scenarios.

The findings from this study underscore the importance of incorporating stochastic elements, such as noise, into swarm simulations. Noise introduces variability that mimics real-world unpredictability, making the model more realistic and applicable to practical situations. By systematically exploring different noise levels, the paper identifies critical thresholds where swarm behavior transitions from order to disorder. The game-theoretic integration mitigates these transitions, enabling the swarm to maintain functionality even in challenging conditions. This robustness is particularly relevant for applications in robotics, where autonomous systems must operate reliably in dynamic and uncertain environments.

While this study achieves significant advancements, it also lays the groundwork for future research. One promising direction is the introduction of heterogeneity in agent types, which would allow the simulation to capture more complex and realistic dynamics. For example, incorporating agents with different movement capabilities or behavioral preferences could lead to the emergence of new patterns and interactions. Another area of exploration is the integration of dynamic environments with obstacles or changing boundaries, which would test the model's ability to adapt to spatial constraints and external disturbances.

The inclusion of learning mechanisms, such as reinforcement learning, is another exciting avenue for future work. While the current model relies on predefined payoff matrices to compute Nash Equilibrium, learning-based approaches could allow agents to adapt their strategies over time based on accumulated experience. This would enable the swarm to improve its performance in dynamic and evolving environments, further enhancing its practical applicability.

The analysis of cohesion, alignment, and separation forces provides valuable insights into the mechanisms that drive swarm behavior. By balancing these forces dynamically, the swarm adapts to varying conditions while achieving collective goals. The model trends underscore the importance of these forces in maintaining the emergent properties of multi-agent systems.

Scalability is another critical aspect that warrants investigation. While the current simulation involves 50 boids, larger swarms with hundreds or thousands of agents may exhibit different dynamics due to increased interaction complexity. Exploring how the model scales with swarm size could provide valuable insights into the principles governing large-scale multi-agent systems.

The use of Self-Organizing Maps (SOMs) to classify and analyze boid behaviors adds another dimension to the study. By clustering agents based on their states, the SOM provides a structured way to visualize and understand subgroup dynamics within the swarm. This approach opens up new possibilities for combining machine learning with classical and game-theoretic models to achieve more sophisticated analyses of emergent behaviors.

Despite its strengths, the model has certain limitations that future work can address. For example, the reliance on local interaction radii for neighbor detection may not fully capture the influence of long-range interactions observed in some natural systems. Additionally, the computational cost of solving the Nash Equilibrium for each agent at every frame could become a bottleneck in larger simulations, highlighting the need for optimization techniques or approximate solutions.

In conclusion, this paper demonstrates the transformative potential of integrating game theory into classical boid models. By leveraging Nash Equilibrium, the model achieves adaptive and robust swarm behavior, with improved coordination under varying noise levels. The insights gained from this study have applications in robotics, autonomous systems, and swarm intelligence, as well as broader interdisciplinary domains such as economics, ecology, and traffic management. Future research can build on these findings by exploring heterogeneity, dynamic environments, learning mechanisms, and scalability, paving the way for more advanced and versatile models. This work represents a significant step forward in the study of multi-agent systems and their applications in solving complex, real-world challenges.

References

[1] C. Zhang, X. Li, J. Hao, S. Chen, K. Tuyls, W. Xue, and Z. Feng, "SA-IGA: a multiagent reinforcement learning method towards socially optimal outcomes," *Autonomous Agents and Multi-Agent Systems*, vol. 33, no. 4, pp. 403–429, 2019. Available: https://link.springer.com/article/10.1007/s10458-019-09411-3

[2] O. Vinyals, I. Babuschkin, W. M. Czarnecki, and M. Mathieu, "Grandmaster level in StarCraft II using multi-agent reinforcement learning," *Nature*, 2019. Available: https://www.nature.com/articles/s41586-019-1724-z

[3] Y. Yang and J. Wang, "An overview of multi-agent reinforcement learning from a game-theoretical perspective," *arXiv preprint arXiv:2011.00583*, 2020. Available: https://arxiv.org/abs/2011.00583

[4] L. Canese, G. C. Cardarilli, L. Di Nunzio, and R. Fazzolari, "Multi-agent reinforcement learning: A review of challenges and applications," *Applied Sciences*, vol. 11, no. 11, 2021. Available: https://www.mdpi.com/2076-3417/11/11/4948

[5] M. Egorov, "Multi-agent deep reinforcement learning," Stanford University Report, 2016. Available: http://vision.stanford.edu/teaching/cs231n/reports/2016/pdfs/122_Report.pdf

[6] M. Wen, J. Kuba, R. Lin, and W. Zhang, "Multi-agent reinforcement learning is a sequence modeling problem," *NeurIPS Proceedings*, 2022. Available: https://proceedings.neurips.cc/paper_files/paper/2022/file/69413f87e5a34897cd010ca698097d0a-Paper-Conference.pdf

[7] S. Shalev-Shwartz, S. Shammah, and A. Shashua, "Safe, multi-agent, reinforcement learning for autonomous driving," *arXiv preprint arXiv:1610.03295*, 2016. Available: https://arxiv.org/abs/1610.03295

[8] T. Wang, H. Dong, V. Lesser, and C. Zhang, "ROMA: Multi-agent reinforcement learning with emergent roles," *arXiv preprint arXiv:2003.08039*, 2003. Available: https://arxiv.org/abs/2003.08039

[9] J. Foerster, F. Song, E. Hughes, and N. Burch, "Bayesian action decoder for deep multi-agent reinforcement learning," *Proceedings of Machine Learning Research*, vol. 97, 2019. Available: http://proceedings.mlr.press/v97/foerster19a/foerster19a.pdf

[10] K. Zhang, Z. Yang, H. Liu, and T. Zhang, "Fully decentralized multi-agent reinforcement learning with networked agents," *Proceedings of Machine Learning Research*, vol. 80, 2018. Available: http://proceedings.mlr.press/v80/zhang18n.html