Unraveling Hidden Interactions in Complex Systems with Deep Learning

Seungwoong Ha

Korea Advanced Institute of Science and Technology

Hawoong Jeong (✉ hjeong@kaist.edu)

Korea Advanced Institute of Science and Technology <https://orcid.org/0000-0002-2491-8620>

Article

Keywords: micro-dynamics, AgentNet, long intrigued researchers, conventional data-driven approaches

Posted Date: December 21st, 2020

DOI: <https://doi.org/10.21203/rs.3.rs-125594/v1>

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

[Read Full License](#)

Unraveling hidden interactions in complex systems with deep learning

Seungwoong Ha¹, Hawoong Jeong^{1,2}

¹*Department of Physics, Korea Advanced Institute of Science and Technology, Daejeon 34141, Korea*

²*Institute for the BioCentury, Korea Advanced Institute of Science and Technology, Daejeon 34141, Korea*

Rich phenomena from complex systems have long intrigued researchers, and yet modeling system micro-dynamics and inferring the forms of interaction remain challenging for conventional data-driven approaches, being generally established by human scientists. In this study, we propose AgentNet, a model-free data-driven framework consisting of deep neural networks to reveal and analyze the hidden interactions in complex systems from observed data alone. AgentNet utilizes a graph attention network with novel variable-wise attention to model the interaction between individual agents, and employs various encoders and decoders that can be selectively applied to any desired system. Our model successfully captured a wide variety of simulated complex systems, namely cellular automata (discrete), the Vicsek model (continuous), and active Ornstein–Uhlenbeck particles (non-Markovian) in which, notably, AgentNet’s visualized attention values coincided with the true variable-wise interaction strengths and exhibited collective behavior that was absent in the training data. A demonstration with empirical data from a flock of birds showed that AgentNet could identify

hidden interaction ranges exhibited by real birds, which cannot be detected by conventional velocity correlation analysis. We expect our framework to open a novel path to investigating complex systems and to provide insight into general process-driven modeling.

1 Main

Complex systems are collections of interactive agents that exhibit non-trivial collective behavior. They have gathered a significant amount of research interest in the last several decades in a wide variety of academic fields from spin systems to human societies. In particular, the domain of physics mainly focuses on investigating the micro-level processes that govern emergent behavior in complex systems and modeling them mathematically. The Vicsek model¹ is a representative example of such approaches, which attempted to explain collective behaviors of active matter like a bird flock with minimal microscopic description. Unfortunately, due to the intrinsic complexity of these systems, extracting hidden micro-dynamics from the observed data of an unknown complex system is virtually infeasible in most cases. Although conventional process-driven modeling is intelligible and provides the conceptual framework, its application to complex systems, to date, still strongly relies on human intuition with various prior assumptions.

To overcome these obstacles, data-driven modeling (DDM), a methodology that finds a relationship between state variables or their time evolution from observed data, has emerged as a powerful tool for system analysis alongside the emergence of machine learning and large-scale data. In previous literature^{2–11}, DDM was employed to discover hidden parameters or dynamics from data in an automated manner. Particularly, active matter modeling greatly relies on DDM by

first designing a model with intuition from observed data and then performing parameter fitting to match the data ^{12–16}, although many of them suffer from sparse, noisy, or discontinuous observation data.

Among various DDM techniques, deep neural networks (DNNs) have recently shown phenomenal performance in pattern recognition and function approximation. One specialized DNN variant for graph-structured data is the graph neural network (GNN) ¹⁷, which models dependencies between linked agents on a graph and has enabled remarkable progress in graph analysis. Similar to ¹⁸, one may depict a complex system as a dynamically changing graph in which each vertex is an agent, with links between agents indicating interactions. In this approach, the problem of modeling the micro-dynamics of single agents becomes equivalent to properly inferring the effect from other agents on a graph and estimating the state transition of each agent at the next time step. Several attempts have been separately made to employ GNNs in the prediction and analysis of specific complex systems and physical models ^{19–27}, but these approaches are mostly limited to the verification of a single system or a small number of agents, and more significantly, it remains difficult to interpret the characteristics of the interaction due to the neural network’s notorious black-box nature. Recently, graph attention networks (GATs) ²⁸ and its applications ^{29,30} showed a path to interpretable GNN by assigning attention to important neighbors, but this attention value cannot be directly interpreted with a physical meaning. For instance, in a multi-dimensional system, the interaction strength cannot be a scalar value since each directional state variable possesses its own interaction range and strength.

Inspired by these recent attempts, we introduce AgentNet, a generalized neural network framework based on GATs with a novel attention scheme to model various complex systems in an physically interpretable manner. AgentNet approximates the transition function of the states of individual agents by training the neural network to predict the future state variables. Due to the rich functional expressibility of DNNs, which is practically unconstrained^{31,32}, AgentNet poses minimum prior assumptions about the unknown nature of the target agents. Our model jointly learns the interaction strength that affects *each variable's* transition and overall transition function from observed data in an end-to-end manner without any human intervention or manual operation. This is a critical difference from the conventional approach with GATs, which only assigns a single attention value per agent while our model assigns completely independent attention values for every state variable and employing separate decoders for each of them. We found that our variable-wise neural attention achieves better performances over GATs' single multiplicative attention, and enables more extensive physical interpretation for the first time that was impossible for conventional GATs such as identifying directional forces separately. Also, the visualization and inspection of the inner modules as granted by our framework enables a clear interpretation of the trained model, which also provides insights for process-driven modeling. As a prediction model, a trained AgentNet can generate an individual level of state predictions from desired initial conditions, making AgentNet an outstanding simulator of target systems including even those that exhibit collective behavior that was absent in the training data.

First, we show the spontaneous correspondence between the complex system and the structure of AgentNet by providing formulations of both systems. The capability of AgentNet is thor-

oughly demonstrated here via data from simulated complex systems: cellular automata ³³, the Vicsek model ¹, and the active Ornstein–Uhlenbeck particle (AOUP) ³⁴ model, along with application to real-world data comprising trajectories in a flock of birds ³⁵ containing more than 1800 agents in a single instance, greatly exceeding the previous range of neural network approaches ^{18,36,37} which treated at most several dozens of agents. For the simulated systems, we show that each component of AgentNet learns predictable and tractable parts of the expected transition function by comparing extracted features with ground-truth functions. For the bird flock where the exact analytical expression of the system is completely unknown, AgentNet successfully provides the interaction range of a bird, which is physiologically plausible and coincides with previous behavioral studies about the bird ^{38,39}.

2 System formulation

In this paper, we focus on a general agent-based system consisting of n agents for which the state of each agent until time T is (at least partially) identified and observed. The basic premise of the agent-based system is that the agent with the same state variable follows the same decision rule, and the interaction strength between two agents can be fully expressed by their state variable. This implies that any two agents with the same state variables should be interchangeable without altering the outcome.

We denote the set of all n agents as $A = \{a_1, a_2, \dots, a_n\}$ and the corresponding observed state variables of all agents at time t as $\mathbf{S}^t = \{\mathbf{s}_1^t, \mathbf{s}_2^t, \dots, \mathbf{s}_n^t\}$, where each state consists of k state

Figure 1: Overview of system formulation and the neural network architecture of the proposed AgentNet. The correspondence between the decision rule of agents in a complex system and a forward pass of AgentNet is depicted. In both panels, the state variable of each agent s_i^t interacts with the state variables of other agents s_j^t in R_i with interaction strength α_{ij}^t . The graph attention core learns R_i with transformer architecture by encoding s_i^t into key k_i^t , query q_i^t , and value v_i^t , and then calculates the weighted sum of the values of other agents v_j^t according to the variable-wise attention weight $\alpha_{ij}^{q,t}$ as computed by neural attention. Different from GATs, AgentNet assigns attention value for each state variable, and decode it separately to strictly impose the information of variable-wise interaction strength. Other functions, namely h_{self} and f , can be captured by both encoder and decoder modules.

variables $\mathbf{s}_i^t = \{s_{i,1}^t, s_{i,2}^t, \dots, s_{i,k}^t\}$. In addition, the system might have j number of time-dependent global external variables $\mathbf{u}^t = \{u_1^t, u_2^t, \dots, u_j^t\}$ that affect agent interaction, such as temperature in a thermodynamic system. For simplicity, we abbreviate the set of time series vectors from t to $t - m$, namely $[\mathbf{S}^t, \mathbf{S}^{t-1}, \dots, \mathbf{S}^{t-m}]$, as $\mathbf{S}^{t,m}$.

where

$$\Delta \mathbf{S}^{t+1} = [f(\mathbf{s}_1^{t,m}, \mathbf{S}_1^{t,m}, \mathbf{u}^{t,m}), \dots, f(\mathbf{s}_n^{t,m}, \mathbf{S}_n^{t,m}, \mathbf{u}^{t,m})] \quad (1)$$

and $\mathbf{S}_i^{t,m}$ indicates that the i th agent's state vector $\mathbf{s}_i^{t,m}$ is omitted from $\mathbf{S}^{t,m}$.

In this study, we assume that the system is mainly dominated by pairwise interactions and higher-order interactions are negligible. Alleviation of this assumption will be discussed in the Conclusion. This means that Eq. (5) becomes

$$\Delta \mathbf{s}_i^{t+1} = f(h_{\text{self}}(\mathbf{s}_i^{t,m}, \mathbf{u}^{t,m}), \sum_{i \neq j} h_{\text{pair}}(\mathbf{s}_i^{t,m}, \mathbf{s}_j^{t,m}, \mathbf{u}^{t,m})) \quad (2)$$

where $h_{\text{self}}(\mathbf{s}_i^{t,m}, \mathbf{u}^{t,m})$ denotes self-interaction and $h_{\text{pair}}(\mathbf{s}_i^{t,m}, \mathbf{s}_j^{t,m}, \mathbf{u}^{t,m})$ captures the pairwise interaction between the i th and j th agents along with the effect of $\mathbf{u}^{t,m}$. We note that this generalized formulation encompasses the transition functions of various fundamental systems such as the Monte Carlo simulation of the Ising model ⁴⁰, the voter model ⁴¹, systems governed by Newtonian dynamics, and phase space dynamics driven by the Liouville equation.

Although Eq. (2) sums up the interaction with every agent except itself, not every agent is relevant to the transition function of a single agent in a general case. Every agent a_i may have its own interaction range $R_i = \{a_j \in A \mid a_i \text{ interacts with } a_j\}$ that can change depending on the current state of the agent, and only a subset (or possibly the entire set) of agents belonging to R_i should be considered. Furthermore, each state variable might be affected by different interaction strengths, e.g. exerting force F_x and F_y can be generally different. Hence, we define the variable-wise interaction strength function between two agents as $w_{ij}^q(\mathbf{s}_i^{t,m}, \mathbf{s}_j^{t,m}, \mathbf{u}^{t,m}) \geq 0$ that outputs the q -th state variable's interaction magnitude of the i th agent, induced by the j th agent. Now, w_{ij}^q can be separated from the pairwise interaction function $h(\mathbf{s}_i^{t,m}, \mathbf{s}_j^{t,m}, \mathbf{u}^{t,m})$ to explicitly indicate the variable-wise interaction strength between agents, as follows:

$$h_{\text{pair}}(\mathbf{s}_i^{t,m}, \mathbf{s}_j^{t,m}, \mathbf{u}^{t,m}) = \sum_q w_{ij}^q(\mathbf{s}_i^{t,m}, \mathbf{s}_j^{t,m}, \mathbf{u}^{t,m}) v_j^t(\mathbf{s}_j^{t,m}). \quad (3)$$

Note that leftover function v_j conveys information solely from the j th agent without loss of generality. So far, we have decomposed an individual transition function into four parts; variable-wise interaction strength function w^q , leftover function v , self-interaction function h_{self} , and overall function f . We note that our formulation aptly applies to a physical system governed by force dynamics by interpreting w^q as the magnitude of a component of an exerting force vector, while the leftover function vector v contains directional information.

In most cases, the exact analytic forms of all these functions (w^q , v , h_{self} , f) are completely

unknown, and it is infeasible to elicit these functions from observed data alone. Especially, blindness to variable-wise interaction strength function w^q significantly complicates this inverse problem since we have to test every possible combination of neighbor candidates while simultaneously guessing the correct *nonlinear* functional form of v , h_{self} , and f . The problem becomes harder if the system has time-correlation because it expands the range of possibly correlated variable pairs further out in the time dimension. To sum up, many of the current methodologies are not capable of DDM for complex systems without strong prior assumptions regarding the functional form. The proposed framework, AgentNet, successfully tackles this conundrum by employing DNNs to jointly learn all of the aforementioned functions by constructing corresponding neural modules for each of the functions and backpropagating errors from state predictions.

Generally, agent modeling of a complex system aims to identify the transition function of its constituents through time steps, which can be written as

$$\begin{aligned}\mathbf{S}^{t+1} &= \mathbf{S}^t + \Delta \mathbf{S}^{t+1} \\ &= \mathbf{S}^t + F(\mathbf{S}^{t,m}, \mathbf{u}^{t,m})\end{aligned}\tag{4}$$

where m is the maximum lag for the system output and F is an overall function that could be deterministic or stochastic. If we focus on the state difference of an individual agent, we can split the overall function F into individual transition function f and get

$$\Delta \mathbf{s}_i^{t+1} = f(\mathbf{s}_i^{t,m}, \mathbf{S}_i^{t,m}, \mathbf{u}^{t,m})\tag{5}$$

Our formulation of agent-based complex systems is shown in Fig. 1 with corresponding

Table 1: System formulation is applied to simulated model systems. Here, $\mathbf{h} = \sum_j h_j$, $\mathbf{h}_x = \{x \text{ directional component of } \mathbf{h}\}$, $\{\mathbf{h}_y = y \text{ directional component of } \mathbf{h}\}$, and $|R_i|$ denote the number of elements of set R_i . $r(s_i^t, s_j^t)$ represents the distance between two agents' positions, while $\theta(s_i^t, s_j^t)$ represents the respective angle of the j th agent to the i th agent. For AOUP, $t+1$ becomes $t+dt$ since the original model is governed by continuous differential equations. U_0 is a fixed norm of propulsion vector, where $U_0^2 = 2D_a/\tau$.

System	Cellular automata	Vicsek model	Active Ornstein–Uhlenbeck Particle
s_i^t	$\{x_i^t, y_i^t, c_i^t\}$	$\{x_i^t, y_i^t, v_{x,i}^t, v_{y,i}^t\}$	$\{x_i^t, y_i^t\}$
u^t	None	None	R
R_i	$\{a_j \in A r(s_i^t, s_j^t) \leq \sqrt{2}\}$	$\{a_j \in A r(s_i^t, s_j^t) < r_c, \theta(s_i^t, s_j^t) < \theta_c\}$	$A - \{a_i\}$
w^q	$w^c = \begin{cases} 1 & \text{if } a_j \in R_i \\ 0 & \text{if } a_j \notin R_i \end{cases}$	$w^x, w^y = \begin{cases} 1 & \text{if } a_j \in R_i \\ 0 & \text{if } a_j \notin R_i \end{cases}$	$w^x, w^{v_x} \propto -3(x_i^t - x_j^t) \frac{e^{-r(s_i^t, s_j^t)^3/R^3}}{R^3}$ $w^y, w^{v_y} \propto -3(y_i^t - y_j^t) \frac{e^{-r(s_i^t, s_j^t)^3/R^3}}{R^3}$
h_{self}	$\{c_i^t\}$	$\{v_{x,i}^t, v_{y,i}^t\}$	$\{x_i^t, y_i^t\}$
f	$\Delta x_i^{t+1} = 0$ $\Delta y_i^{t+1} = 0$ $\Delta c_i^{t+1} = \delta_{c_i^t=0} \delta_{\mathbf{h}=3} - \delta_{c_i^t=1} (1 - \delta_{\mathbf{h}=2}) (1 - \delta_{\mathbf{h}=3})$	$\Delta x_i^{t+1} = v_{x,i}^{t+1}$ $\Delta y_i^{t+1} = v_{y,i}^{t+1}$ $\Delta v_{x,i}^{t+1} = (\mathbf{h}_x / (R_i + 1)) + \mathcal{N}(0, \sigma)$ $\Delta v_{y,i}^{t+1} = (\mathbf{h}_y / (R_i + 1)) + \mathcal{N}(0, \sigma)$	$\Delta x_i^{t+dt} = v_{x,i}^{t+1}$ $\Delta y_i^{t+dt} = v_{y,i}^{t+1}$ $\Delta v_{x,i}^{t+dt} = \mathbf{h}_x dt + \frac{\sqrt{2\gamma T}}{\gamma} \mathcal{N}(0, \sqrt{dt}) + f_{x,i}$ $\Delta v_{y,i}^{t+dt} = \mathbf{h}_y dt + \frac{\sqrt{2\gamma T}}{\gamma} \mathcal{N}(0, \sqrt{dt}) + f_{y,i}$ $\Delta f_{x,i}^{t+dt} = -dt/\tau f_{x,i}^t + \sqrt{U_0^2 \tau T} \mathcal{N}(0, \sqrt{dt})$ $\Delta f_{y,i}^{t+dt} = -dt/\tau f_{y,i}^t + \sqrt{U_0^2 \tau T} \mathcal{N}(0, \sqrt{dt})$

modules in AgentNet: the value vector of the transformer captures the self-interaction h_{self} and leftover information v ; variable-wise attention weight α^q captures the interaction magnitude w^q ; and the weighted sum along with the decoder corresponds to the overall function f . This formulation can express all of the model systems used in this study, as described in Table 1.

3 Model description

AgentNet is a generalized framework for the data-driven modeling of agent-based complex systems, covering most previous works and reinforced with several modifications. The base module of AgentNet, a graph attention module, is similar to a GAT²⁸, with transformer architecture⁴², where each agent decides its next state by putting information from itself and the attention-weighted sum of other agents together. AgentNet initially operates on a fully-connected graph, implying that it initially assumes every agent as a possible neighbor and gradually learns the true interaction partners and strength through training. Our model first encodes the state variables of agent s^t with an encoder, then passes the information to the transformer which computes the impact from the entire system state s^t , and finally decodes the outcome with a decoder to obtain the state difference.

In most cases, complex systems have diverse characteristics that are difficult to incorporate into a single modeling framework. As a universal framework, AgentNet resolves this diversity by modifying the encoder and decoder and setting a proper optimization function to fit particular system characteristics while maintaining the core module of the network. In this way, AgentNet addresses a variety of system characteristics such as continuity of state variables, stochasticity of

transition function, and memory effects.

First, AgentNet can handle various types of state variables by minimizing cross-entropy for discrete variables and the mean squared error for continuous variables. Second, when the decision rule of a target system is stochastic, there are several ways to construct a neural network with probabilistic output^{43–46}. AgentNet employs a Gaussian neural network⁴⁵ as the decoder of the stochastic AgentNet, which produces means and variances of multiple univariate Gaussian distributions. Lastly, some of the collective phenomena in complex systems appear in non-Markovian settings where past states affect the future state. In this study, we use long short-term memory (LSTM) models for the encoder and decoder of AgentNet to capture (potential) memory effects in the system.

The graph attention module in AgentNet explicitly assigns variable-wise importance α_{ij}^q by first constructing the attention coefficient a_{ij}^q from encoded data $e(\mathbf{s}^t)$ and applying the sigmoid function to normalize the scale. We note that the choice of the sigmoid function is crucial because unlike most previous literature^{24,28,42} where a softmax normalization between agents was used ($\alpha_{ij} = \frac{a_{ij}}{\sum_k \exp(a_{ik})}$), here we aim to infer the absolute variable-wise interaction magnitude without normalization among the agents. Also, further differing from conventional approaches for attention coefficients such as additive⁴⁷ and multiplicative⁴² mechanisms, attention coefficients in AgentNet are calculated by multi-layer perceptrons (MLPs) (Att), which enables much more flexible representations and crucial for capturing complex interactions. We name this attention scheme as *neural* attention (See Supplemenatry Note 2 and Supplementary Figure 2 for the advantages of

neural attention).

By virtue of variable-wise separated decoder, attention weights α_{ij}^q *only affects to* q -th variable, thus one can identify interaction strengths for each variable by visualizing predicted attention weights. Note that this is different from widely-known multi-headed attention since it feeds concatenated output into a single decoder while AgentNet does not concatenate the output and strictly separates each decoder in order to impose a variable-wise transition function for each attention value, not a mixed overall transition function. In short, AgentNet clarifies the variable-wise neural attention scheme from an unknown function of interaction strengths to physically interpretable variable-wise strength. Our study is the first in-depth demonstration of the capability of this form of graph attention scheme, achieved by comparing the attention weights for each variable to the ground-truth interaction strengths in various simulated complex systems.

4 Results

This study utilizes three representative complex systems to demonstrate the capacity of AgentNet, along with one empirical dataset for framework evaluation. Table 2 summarizes the characteristics of the model systems with an escalating level of complexity. All of the code for model training and system simulation has been deposited in ⁴⁸.

Cellular Automata First, we verify AgentNet with an older yet fundamental system with rich phenomena, the cellular automata (CA) model. In the CA model, each cell has its own discrete state, either *alive* or *dead*. Each cell interacts with its eight adjacent neighbors, and the state of

Figure 2: Result of AgentNet for cellular automata. **(A)** Attention weight transition of a single target cell throughout the training. In the initial stage, the model has no information about the interaction range and assigns near-zero values to all of the cells in the system. Attention gets narrowed down to a smaller region as training advances, and finally concentrates on eight surrounding cells, which is the theoretical interaction range. **(B)** Attention weight $\bar{\alpha}^c$ of neighbors and outside cells during 200 epochs of training. The attention weight of neighbor cells increases as training proceeds, while the weight of other cells remains 0. Data is averaged from 100 test samples. **(C)** AgentNet with respect to given alive (left) and dead microstates (right). The total number of alive cells in the neighborhood is denoted by s , which is the sole parameter of the CA decision rule.

each cell evolves according to the following two rules. First, a live cell stays alive if two or three neighbor cells are alive. Secondly, a dead cell becomes alive if exactly three neighbor cells are alive. Thus, the interaction strength of CA can be expressed as an indicator function $\mathbb{1}_{R_i}$ where its value is 1 if $a_j \in R_i$ and 0 otherwise.

We simulate CA data in the form of a 14×14 grid of cells with initially randomized states, and the state of the grid after a single time step becomes the target label for each data. AgentNet for CA receives three state variables from each cell: positions $\mathbf{x}^t$ and $\mathbf{y}^t$, and cell state $\mathbf{c}^t$. The output here is a list of expected probabilities that each cell becomes alive. We use the binary cross-entropy loss function between the AgentNet output and the ground-truth label.

Figure 2 summarizes the results of AgentNet for CA, depicting the cell state attention weight α^c of the target cell (in this case, the 102nd cell) across the entire grid. After 120 epochs, AgentNet quickly *realized* that a vast majority of the cells are irrelevant to the target cell, and thereafter concentrated its attention to a more compact region; Fig. 2B shows that AgentNet gradually learns to focus on neighbor cells only. AgentNet was able to figure out the true interaction range after 200 epochs. The result of the prediction test for unseen cases showed perfect 100% accuracy, as depicted in Fig. 2C.

Vicsek model Next, we validate the capability of AgentNet for a continuous and stochastic system. The Vicsek model ¹ (VM) is one of the earliest and most prominent models to describe an active matter system, where each agent averages the velocity of nearby agents (including itself) to replace its previous velocity. At each time step, every agent updates its position by adding this

Figure 3: Result of AgentNet for the Vicsek model. Initially, agents are randomly distributed in a circular region with radius $R = \sqrt{5}$, without any boundary condition. **(A)** Attention weight visualization of two sample cells, a_1 and a_2 . Both cases show a circular sector of attention distribution with clear boundaries that perfectly matches with ground-truth interaction range. **(B)** Averaged attention weight for variables x and y before and after training. The fully trained AgentNet learned to identify the neighbor agents and ignore the irrelevant others by assigning near-one and near-zero attention weights, respectively. **(C)** Position predictions by AgentNet for the two sample cells a_1 and a_2 . Circles indicate the starting positions of the two particles, with the two heatmaps showing the AgentNet prediction along with the means of predicted distributions (Xs). The model predicts the expected theoretical distribution (crosses) with great precision, even when the given training samples (green and blue stars) are distant from the means of the theoretical distribution.

newly assigned velocity with stochastic noise. In this study, every 300th agent in the simulation interacts with other agents within the range $r_c = 1 \text{ m}$ and viewing angle $\theta_c = 120^\circ$ of its heading direction. This complex interaction range models the limitations of sight range and angle in real organisms such as birds.

The model receives four state variables, positions $\mathbf{x}^t$ and $\mathbf{y}^t$, and velocities $\mathbf{v}_x^t$ and $\mathbf{v}_y^t$, and predicts the positions of the next time step $\mathbf{x}^{t+1}$ and $\mathbf{y}^{t+1}$ in the form of two one-dimensional (1D) Gaussian distributions by optimizing the sum of two negative log-likelihood (NLL) loss functions. Note that each training data provides only a single stochastically sampled value, thus putting AgentNet for VM in the difficult condition of trying to identify the general decision rule with only one sample for each environment.

As a result, AgentNet for VM achieved a NLL loss of -1.365 for the test data, while the theoretically computed NLL loss was -1.524 . We note that other approaches, such as naive MLPs, failed to achieve meaningful prediction and resulted in a NLL loss of around $+1.0$ for the VM. Figure 3A visualizes the x -variable attention weight α^x of two sample agents, a_1 and a_2 . AgentNet for VM accurately learned the interaction boundary of the given VM, which resembles a major sector of the circle. As Fig. 3B shows, the fully trained AgentNet assigns a high value to its x and y -variable attention only for neighbor agents, while the untrained AgentNet has no distinction between neighbor and outside cells. The predicted position distributions for these two sample agents are depicted in Fig. 3C. We observe that AgentNet precisely estimated the ground-truth distribution with true mean, even though the given training data is sampled from a stochastic distribution

and did not match the expected mean value. This shows the capability of AgentNet to learn the general transition rule governing the entire set, rather than merely memorizing every single training datapoint and overfitting them. We also report that AgentNet shows the same outcome with unseen test data.

Active Ornstein–Uhlenbeck Particle Differing from the Vicsek model, some active matter shows a time-correlation of particle positions due to the force inherent in the particles that allows them to move. These systems are generally referred to as self-propelled particles, which can be described by overdamped Langevin equations for the position $\mathbf{x}_i$ of each particle as

$$\gamma \dot{\mathbf{x}}_i = (\mathbf{F}_i^{\text{ext}} + \mathbf{F}_i^{\text{int}}) + \sqrt{2\gamma T} \boldsymbol{\eta}_i + \gamma \mathbf{f}_i, \quad (6)$$

where γ the drag coefficient and T is temperature³⁴. Here, $\mathbf{F}_i^{\text{ext}}$ is the external potential, and $\mathbf{F}_i^{\text{int}} = -\nabla_i V$ is the total force exerted on particle i due to the soft-core potential from other particles, $V = \exp(-|r_{ij}|^3/R^3)$, that depends on relative distance r_{ij} and interaction length R . In this study, we use AOUPs confined in a harmonic potential as an example system, describing the intrinsic propulsion force f_i as an independent Ornstein–Uhlenbeck process as

$$\tau \dot{\mathbf{f}}_i = -\mathbf{f}_i + \sqrt{2D_a} \mathbf{w}_i, \quad (7)$$

where τ is correlation time, D_a is a diffusion constant, and $\mathbf{w}_i$ is a standard Gaussian white noise. As an external potential, we apply a weak harmonic potential $\mathbf{F}_i^{\text{ext}} = -k\mathbf{x}_i$ with spring constant $k = 0.1$ to confine the particles, as broadly assumed and experimentally employed⁴⁹.

Figure 4: Result of AgentNet for AOUP. **(A)** In both panels, eight steps of the test data ($R = 4$) of four particles are drawn with black dots starting from the large black dots. AgentNet predictions of the trajectories in the following 12 steps (right panel) perfectly coincide with the sample trajectories from the true Langevin equation (left panel). 100 samples are drawn in both panels, and the final positions are highlighted with white stars. **(B)** Equilibrium state of a system with $R = 5$, which is unseen at the training stage. A single realization from the true distribution is drawn with final positions marked by red stars (left), while a single sample from the predicted distribution of AgentNet is drawn with final positions marked by blue stars (right). AgentNet for AOUP captures the generalized effect of interaction length R and predicts the collective behavior of the untrained system. **(C)** Exerted x -directional force $F_x^{\text{int},x}$ and x -directional velocity attention $\bar{\alpha}^{v_x}$ shows a strong linear relationship, while single attention value from GAT does not captures any of the force component. Same holds for **(D)**, for the case of y -direction. **(E)** By plotting relative distance r_{ij} versus force and attention, scaled attention shows good coincidence with the force value up to constant factor $c = 0.28$. **(F)** Visualization of $\bar{\alpha}_{19}^{v_x}$ for a single target particle (blue). Attention and force values for (C) to (E) are collected from 100 test samples with $R = 4$.

This model is known to exhibit a collective clustering phenomenon, with the periodicity of the resulting hexagonal pattern known to be approximately $1.4R$ with no $\mathbf{F}^{\text{ext}}$ ⁵⁰.

AgentNet for AOUP adopts an LSTM model as an encoder to enable iterative data generation. The model observes 8 steps of trajectories as input data, and the loss is calculated for the next 12 steps. The model receives four state variables, $\mathbf{x}^t, \mathbf{y}^t, \mathbf{v}_x^t$, and $\mathbf{v}_y^t$, and global variable R ranging from 2.0 to 4.0, and predicts the parameters for four 1D Gaussian distributions, similar to the AgentNet for VM. Note that the internal variable, $\mathbf{f}_i$, which has its own Ornstein–Uhlenbeck dynamics, is not present in the input data and thus the neural network *has to infer* this hidden variable by eight steps of past trajectory.

First, we compare the average displacement error (ADE) and final displacement error (FDE) of our model among 12 predicted steps as in previous works ^{22,23} along with a linear extrapolation and naive LSTM without the graph attention core as baselines. AgentNet for AOUP showed ADE/FDE of 0.041/0.064, while linear extrapolation and LSTM showed much lower performances of 0.210/0.465 and 0.158/0.316, respectively. The performance of our model also exceeds the modern architectures like GAT3+ (GAT with 3-headed attention and transformer architecture), which showed the performance of 0.065/0.087. Figure 4 summarizes the result of AgentNet for AOUP. Our model precisely predicted the future trajectories subject to the past states, as depicted in Fig. 4A where 100 trajectories sampled from the ground-truth Langevin equation and AgentNet for AOUP are drawn. Figure 4B shows that AgentNet is also capable of predicting the untrained region of the global variable R and further exhibits a collective behavior that occurs far beyond the

trained time scale. Since our model can iteratively predict future states indefinitely, we tested our model to predict a total of 42 steps, which is 30 more steps than the model was originally trained for. Surprisingly, our model predicts a precise hexagonal pattern of periodicity 7, which coincides with the theoretical value of periodicity when $R = 5$. This verifies a generalization capability since the model had never been trained in the $R = 5$ condition and yet still properly captured the collective phenomenon, which only occurs at a much longer timescale than its training data had.

Moreover, we demonstrate that the attention α^q corresponds to the internal force $F^{\text{int},q}$, up to a constant factor, as we claimed in system formulation section. Fig. 4C and 4D verifies this by showing the attention for x and y -directional velocity v_x, v_y and the magnitude of corresponding internal force $F^{\text{int},x}, F^{\text{int},y}$, which clearly exhibits a strong linear relationship. This cannot be achieved by a single-valued attention from conventional GAT, which shows a poor agreement with any of the force components. We report that the single attention value from GAT tries to convey the sum or average of each interaction strength. In Fig. 4E, we draw the scaled attention for v_x and the internal force of the x direction $F^{\text{int},x} = \nabla V_{ij} = (-3r_{ij}^2 \exp[-r_{ij}^3/R^3])/R^3$ versus the relative distance to the target particle r_{ij} . Despite a slight disagreement at small r_{ij} , scaled attention with constant factor $c = 0.28$ well matches $F^{\text{int},x}$ and therefore can be considered as a good approximation for interaction magnitude. (See Supplementary Discussion 1 for further investigation on AOUP attention.) AgentNet for AOUP successfully predicted and investigated one of the most complex systems possessing internal potential, external potential, memory effects, and stochastic noises. We note that variables other than v_x also showed similar linear relationships with corresponding forces (results for other variables are reported in ⁴⁸.)

Figure 5: Result of AgentNet for CS. **(A)** Displacement errors of linear extrapolation, naive LSTM, GAT, GAT3+, and AgentNet. AgentNet shows the lowest displacement error compared to the baselines. Here, the final displacement error (FDE) of step n indicates the averaged error of birds for which their trajectories terminated at step n . All of the results are averaged value from three trials. **(B)** Exemplary snapshot of the visualized attention of a single agent (blue circle) from the test dataset. **(C)** Two-dimensional heatmap of averaged attention α^{v_x} and cosine similarity of the velocity with respect to the relative coordinates. We align every bird’s heading direction in the test dataset to the x -axis (blue arrow) and draw cross-sections in the xy -plane (upper panels) and xz -plane (lower panels). Different from the velocity correlation, attention shows more concentrated and strongly directional distributions that coincide with previous literature about the bird’s visual frustum and sight direction. For attention, a contour of the top 0.01% of the attention value is visualized (red, dashed) as well as the direction of the maximum attention value (red arrows).

Chimney swift trajectory Finally, we demonstrate the capability of our framework by predicting the empirical trajectories of a freely behaving flock of chimney swifts (CSs). Bird flocks are renowned for their rich diversity of flocking dynamics, for which models with various mechanisms such as velocity alignment and cohesion have been proposed in the last several decades^{1,16,38}. We employed here a portion of the data from⁵¹, recorded in Raleigh, North Carolina, in 2014. Since half of the trajectories last less than $150f = 5s$ and 80% last less than $300f = 10s$ due to occlusion and the limited sight of the camera, observation data takes the form of a spatiotemporal graph with dynamic nodes where each agent lasts a short period and then disappears. Thus, discarding non-full trajectories as in previous works^{22,23} would significantly reduce the number of birds to consider at a given time step. To handle these disjointed yet entangled pieces of trajectories, we propose a novel inspection method that examines the data at every step of the LSTM to manually connect the hidden states from the past, exclude the nonexistent birds at a certain time, and start a new chain of hidden states from a separate neural network if an agent newly enters the scene. While several previous approaches could handle graphs with dynamic edges^{52–54}, AgentNet is, to the best of our knowledge, the first attempt to deal with dynamic nodes on a spatiotemporal graph (see Supplementary Note 1 for a formal explanation of the inspection scheme).

The number of total birds appearing in each set varied from 300 to 1800, and each trajectory in the set started and ended at different times. The model received state variables that exist at the current time step, produced statistics of three-dimensional position and velocity, and then the sampled states were fed back into the model for the next time prediction. NLL losses were calculated at every LSTM step for existing birds.

Figure 5 summarizes the results of AgentNet for CS. The predictive power of AgentNet is illustrated in Fig. 5A, where linear extrapolation and naive LSTM show mostly similar results while AgentNet shows greatly reduced errors at predicting longer time steps, achieves better performances than GAT and GAT3+. Figure 5B, showing the visualized attention of a typical bird, clearly indicates the near-sighted and forward-oriented nature of the bird’s interaction range. To further verify this interaction range, we averaged the attention values from the first step of predictions according to the relative coordinates of the target bird. The averaged results for α^{v_x} from the entire test set are drawn in Fig. 5C along with the averaged cosine similarity of the velocity, which is a commonly used measure to find the range of interaction. The interaction range projected on the xy -plane coincides with previous literature about biological agents’ visual frustum, which depends on forward-oriented sight and the relative distance from each agent ^{20,35,38,55}. Also, the bird’s z -directional attention is relatively concentrated *downwards*; this predicted attention is physiologically plausible since downward-oriented visual fields are widely reported in various types of birds due to their foraging nature and the blind area from the beak ^{39,56}.

Interestingly, Fig. 5C shows that the velocity correlation on the xy -plane and xz -plane shows no particular directional tendency as attention does. Although many studies employ state correlations between agents to figure out the characteristics of interaction ^{57,58}, correlation might be significantly different from the interaction range itself ⁵⁹. Different from correlations, our model provides a *causal* interaction strength since the attention value is strongly connected to the predictability of future dynamics, which is quite useful for inferring and modeling the microdynamics of individual birds.

Our model with variable-wise attention can further verify important physical insights. For instance, we have found that although the scale is different, the form of attention concentration is surprisingly the same regardless of the directions (Results for other variables are reported in ⁴⁸). This directional homogeneity strongly implies that the bird-bird interaction is more like a near-sighted version of the Vicsek model, differs from the distance-based force models like AOUP which must exhibit directional heterogeneity. In conclusion, AgentNet employed the position and velocity (heading direction) of neighboring birds into its prediction, thereby showing better prediction compared to the non-interactive baseline and qualitatively plausible interaction range.

5 Conclusion

This study proposed AgentNet, a generalized framework for the data-driven modeling of a complex system. We demonstrated the flexibility, capability, and interpretability of our framework with large-scale data from various complex systems. Our framework is *universally* applicable to agent-based systems that are governed by pairwise interactions and for which a sufficient amount of data is available. The proposed framework can infer and visualize variable-wise interaction strength between agents, which could assist researchers in gaining clearer insights into given systems and their dynamics. Furthermore, AgentNet is scalable for an arbitrary number of agents due to the nature of GNNs, thus facilitating free-form simulation of the desired system with any initial condition. Since attention values from our model can be directly interpretable as a variable-wise interaction strength function, we expect that AgentNet will be useful in heterogeneous settings where each state variable interacts with different neighbors.

There are a great number of domains in which AgentNet is anticipated to exhibit its full potential. As we demonstrated via AOUP and CS, the analysis of active matter such as bacterial cells ^{49,60}, animal flocks ^{12–16}, or pedestrian dynamics ^{22,23} may greatly benefit from our approach. Also, since GNNs were originally proposed for data with graph structures, AgentNet may yield data-driven models of both agent and node dynamics of a network by incorporating an adjacency matrix instead of assuming a complete graph. AgentNet can retrieve the underlying graph and interaction strength from data, which encompasses the research fields of epidemic dynamics ⁶¹, network identification, and various inverse Ising problems ⁷. We could further apply different encoders and decoders to improve the performance and include available domain knowledge. For instance, a Gaussian mixture model ⁴⁵ or variational model ²³ that could approximate an arbitrary distribution may be suitable to approximate multimodal or highly irregular distributions.

One limitation of the current work is that AgentNet cannot fully capture three or higher orders of interactions. The pairwise assumption is nearly the only inductive bias we have imposed on our model, which will require modification if the target system is expected to have strong higher-order interaction. First, by increasing the number of message passing layers, GNNs can employ information from further than one-hop neighbors and possibly capture the higher-order interactions among three or more agents. Another way to alleviate the pairwise assumption is to consider higher-order interactions directly in network construction. Applying a GNN with a hypergraph structure ^{62–64}, one of the rapidly growing research areas in machine learning, to AgentNet would be a direct extension of the current study.

We highlight the virtually unbounded scope of the proposed framework in this study, and hope that AgentNet shines a new light on physical modeling and helps researchers in diverse domains delve into their systems in a data-driven manner.

Methods

AgentNet implementations We implemented our AgentNet model with `PyTorch`⁶⁵. The encoder and decoder layers of AgentNet are composed of multi-layer perceptrons (MLPs). The dimension notation such as [32, 16, 1] means that the model consists of three perceptron layers with 32, 16, and 1 neurons in each layer. Also, `dims.` is an abbreviation of dimensions.

All of the encoding layers of AgentNet are composed of [**Input dims, 256, Attention dims**]. Here, input dimensions are chosen as the sum of the number of state variables and additional variables, such as global variables (as in AgentNet for AOUP) or indicator variables (as in AgentNet for CS). The form of the final dimension indicates that each output of the encoder (key, query, and value) will be processed separately. See⁴² for more details about transformer architecture.

With these outputs and (additional) global external variables ($\mathbf{u}$), neural attention is applied to calculate attention value α_{ij} from encoded data $e(\mathbf{s}^t)$. First, the algorithm constructs a_{ij}^q , the attention coefficient for the q -th state variable between agents i and j , and feeds the concatenated vectors into $\text{MLP}(\text{att})$ as

$$a_{ij}^q = \text{Att}(\text{Key}(e(\mathbf{s}^i)), \text{Query}(e(\mathbf{s}^j)), \mathbf{u}) \quad (8)$$

and applies the sigmoid function

$$\alpha_{ij}^q = \frac{1}{1 + \exp(-a_{ij}^q)} \quad (9)$$

where Key, Query, and Att indicate the corresponding MLPs used for transformer architecture and has dimensions. Table ?? shows the implementation details of AgentNet for each target system.

After variable-wise attentions are multiplied to their respective values and averaged, we concatenate the (original target agent's) value and its averaged attention-weighted values (from others) and feed it into the variable-wise separated decoder. Since two tensors are concatenated, the last dimension of this tensor has twice the length of the original dimension of the value tensor. The decoder consists of [**$2 \times \text{value dims.}, 128, \text{output dims.}$**].

In the stochastic setting (VM, AOUP, and CS), the decoded tensor further feeds into other layers to obtain sufficient statistics for the probabilistic distribution. In this paper, those statistics are means and variances of state variables. MLP layers for these values consist of [**output dims., 64, corresponding number of variables**]. For instance, AgentNet for CS has $2 \times 6 = 12$ separate layers to calculate means and variances for 6 state variables.

In the case of a target system with probable time correlations, we adopted long short-term memory (LSTM) as an encoder to capture the correlations⁶⁶. Hidden states and cell states have 128 dims. each and are initialized by additional MLPs that are jointly trained with the main module.

As explained in the main manuscript, AgentNet checks at each time step whether an agent is new and present. When an agent is newly entered, new LSTM hidden states are initialized. Otherwise, hidden states succeed from the previous result.

Baseline implementations For the baseline, we employed a MLP, LSTM, and GAT model where the variable-wise graph attention module is missing. For MLP and LSTM, We doubled the number of layers and neurons of the decoder to compensate for the missing attention module, which its decoder consists of $[2 \times \text{value dims.}, 256, 256, \text{output dims.}]$. For standard GAT²⁸, we left everything the same as AgentNet and replaced variable-wise attention core to original graph-attention core with linear projection matrices of $[\text{Attention dims.}, 128]$ and inner-product attention was used with those 128-dimensional vectors. For GAT3+, we implemented multi-headed attention (with 3 heads) with dimensions of 12 (for AOUP) and 32 (for CS), and every other module is the same as AgentNet. This choice is to (Note that even for GAT and GAT3+, we used $[\text{Input dims.}, 256, \text{Attention dims}]$ dimensions of encoding layers for the key, query and value, instead of linear projection matrices as the original architecture.)

Training scheme All training used 2 to 10 NVIDIA TITAN V GPUs, with which the longest training for a single model took less than two days. Mish activation function⁶⁷ with a form of $f(x) = x \tanh(\text{softplus}(x))$ and the Adam Optimizer⁶⁸ were used for the construction of models and training. The learning rate was set to 0.0005 and decreased to 70% of the previous value when the test loss remained steady for 30 epochs. In the case of AgentNet for CS, we employed weighted NLL loss for different time steps, in which weights are inversely proportional to the frequency of the sample with a given trajectory length, to resolve the imbalance of available trajectory length.

Table 3 shows further details of the model for each system, including the number of attention heads.

6 References

1. Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I. & Shochet, O. Novel type of phase transition in a system of self-driven particles. *Physical Review Letters* **75**, 1226 (1995).
2. Bongard, J. & Lipson, H. Automated reverse engineering of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences* **104**, 9943–9948 (2007).
3. Schmidt, M. & Lipson, H. Distilling free-form natural laws from experimental data. *science* **324**, 81–85 (2009).
4. Runge, J., Nowack, P., Kretschmer, M., Flaxman, S. & Sejdinovic, D. Detecting and quantifying causal associations in large nonlinear time series datasets. *Science Advances* **5**, eaau4996 (2019).
5. Lu, P. Y., Kim, S. & Soljagic, M. Extracting interpretable physical parameters from spatiotemporal systems using unsupervised learning. *Phys. Rev. X* **10**, 031056 (2020). URL <https://link.aps.org/doi/10.1103/PhysRevX.10.031056>.
6. Champion, K., Lusch, B., Kutz, J. N. & Brunton, S. L. Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences* **116**, 22445–22451 (2019).

Table 2: Models to test the performance of AgentNet and their respective characteristics. The X s indicate the opposite characteristics: discrete, deterministic, Markovian, and simulated data, respectively.

System	Continuity	Stochasticity	Memory effect	Empirical data	Interaction	Remarks
Cellular automata	X	X	X	X	Discrete	-
Vicsek model	✓	✓	X	X	Discrete	Generalization
Active OU particle	✓	✓	✓	X	Continuous	Collective phenomena
Chimney swift flock	✓	✓	✓	✓	Unknown	Missing data

Table 3: Implementation details of the models for sample systems. * : global variable (R), ** : indicator variable

System	Input data dims.	Output data dims.	Attention head dims.	Layer composition
1. Cellular Automata	3	1	16	[32, 32, 16, 1]
2. Vicsek Model	4	2	16	[32, 64, 32, 1]
3. Active Ornstein–Uhlenbeck Particle	4	4	16	[32+1*, 16, 8, 1]
4. Chimney Swift	6 + 1**	6	96	[192, 16, 8, 1]

7. Nguyen, H. C., Zecchina, R. & Berg, J. Inverse statistical problems: from the inverse ising problem to data science. *Advances in Physics* **66**, 197–261 (2017).
8. Gorbachenko, V. I., Lazovskaya, T. V., Tarkhov, D. A., Vasilyev, A. N. & Zhukov, M. V. Neural network technique in some inverse problems of mathematical physics. In *International Symposium on Neural Networks*, 310–316 (Springer, 2016).
9. Wu, T. & Tegmark, M. Toward an artificial intelligence physicist for unsupervised learning. *Physical Review E* **100**, 033311 (2019).
10. Li, H. *et al.* Data-driven quantitative modeling of bacterial active nematics. *Proceedings of the National Academy of Sciences* **116**, 777–785 (2019).
11. Solomatine, D. P. & Ostfeld, A. Data-driven modelling: some past experiences and new approaches. *Journal of hydroinformatics* **10**, 3–22 (2008).
12. Buhl, J. *et al.* From disorder to order in marching locusts. *Science* **312**, 1402–1406 (2006).
13. Puckett, J. G., Kelley, D. H. & Ouellette, N. T. Searching for effective forces in laboratory insect swarms. *Scientific reports* **4**, 4766 (2014).
14. Katz, Y., Tunstrøm, K., Ioannou, C. C., Huepe, C. & Couzin, I. D. Inferring the structure and dynamics of interactions in schooling fish. *Proceedings of the National Academy of Sciences* **108**, 18720–18725 (2011).

15. Ballerini, M. *et al.* Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the national academy of sciences* **105**, 1232–1237 (2008).
16. Bialek, W. *et al.* Statistical mechanics for natural flocks of birds. *Proceedings of the National Academy of Sciences* **109**, 4786–4791 (2012).
17. Battaglia, P. W. *et al.* Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261* (2018).
18. Battaglia, P., Pascanu, R., Lai, M., Rezende, D. J. *et al.* Interaction networks for learning about objects, relations and physics. In *Advances in neural information processing systems*, 4502–4510 (2016).
19. Gilpin, W. Cellular automata as convolutional neural networks. *Physical Review E* **100**, 032402 (2019).
20. Heras, F. J., Romero-Ferrero, F., Hinz, R. C. & de Polavieja, G. G. Deep attention networks reveal the rules of collective motion in zebrafish. *PLoS computational biology* **15**, e1007354 (2019).
21. Schütt, K. T., Tkatchenko, A. & Müller, K.-R. Learning representations of molecules and materials with atomistic neural networks. In *Machine Learning Meets Quantum Physics*, 215–230 (Springer, 2020).
22. Alahi, A. *et al.* Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 961–971 (2016).

23. Gupta, A., Johnson, J., Fei-Fei, L., Savarese, S. & Alahi, A. Social gan: Socially acceptable trajectories with generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2255–2264 (2018).
24. Vemula, A., Muelling, K. & Oh, J. Social attention: Modeling attention in human crowds. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 1–7 (IEEE, 2018).
25. Kipf, T., Fetaya, E., Wang, K.-C., Welling, M. & Zemel, R. Neural relational inference for interacting systems. *arXiv preprint arXiv:1802.04687* (2018).
26. Mrowca, D. *et al.* Flexible neural representation for physics prediction. In *Advances in Neural Information Processing Systems*, 8799–8810 (2018).
27. Seo, S. & Liu, Y. Differentiable physics-informed graph networks. *arXiv preprint arXiv:1902.02950* (2019).
28. Veličković, P. *et al.* Graph attention networks. *arXiv preprint arXiv:1710.10903* (2017).
29. Bapst, V. *et al.* Unveiling the predictive power of static structure in glassy systems. *Nature Physics* **16**, 448–454 (2020).
30. Sanchez-Gonzalez, A. *et al.* Learning to simulate complex physics with graph networks. *arXiv preprint arXiv:2002.09405* (2020).
31. Cybenko, G. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems* **2**, 303–314 (1989).

32. Raghu, M., Poole, B., Kleinberg, J., Ganguli, S. & Dickstein, J. S. On the expressive power of deep neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 2847–2854 (JMLR. org, 2017).
33. Gardner, M. Mathematical games. *Scientific American* **222**, 132–140 (1970).
34. Caprini, L., Hernández-García, E., López, C. & Marconi, U. M. B. A comparative study between two models of active cluster crystals. *Scientific reports* **9**, 1–13 (2019).
35. Evangelista, D. J., Ray, D. D., Raja, S. K. & Hedrick, T. L. Three-dimensional trajectories and network analyses of group behaviour within chimney swift flocks during approaches to the roost. *Proceedings of the Royal Society B: Biological Sciences* **284**, 20162602 (2017).
36. Hoshen, Y. Vain: Attentional multi-agent predictive modeling. In *Advances in Neural Information Processing Systems*, 2701–2711 (2017).
37. Sukhbaatar, S., Fergus, R. *et al.* Learning multiagent communication with backpropagation. In *Advances in Neural Information Processing Systems*, 2244–2252 (2016).
38. Hemelrijk, C. K. & Hildenbrandt, H. Some causes of the variable shape of flocks of birds. *PloS one* **6**, e22479 (2011).
39. Martin, G. & Shaw, J. Bird collisions with power lines: failing to see the way ahead? *Biological Conservation* **143**, 2695–2702 (2010).

40. Landau, D. P. & Binder, K. Phase diagrams and critical behavior of ising square lattices with nearest-, next-nearest-, and third-nearest-neighbor couplings. *Physical Review B* **31**, 5946 (1985).
41. Holley, R. A. & Liggett, T. M. Ergodic theorems for weakly interacting infinite systems and the voter model. *The annals of probability* 643–663 (1975).
42. Vaswani, A. *et al.* Attention is all you need. In *Advances in neural information processing systems*, 5998–6008 (2017).
43. Goodfellow, I. *et al.* Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680 (2014).
44. Kingma, D. P. & Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
45. Chung, J. *et al.* A recurrent latent variable model for sequential data. In *Advances in neural information processing systems*, 2980–2988 (2015).
46. Graves, A. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850* (2013).
47. Chorowski, J. K., Bahdanau, D., Serdyuk, D., Cho, K. & Bengio, Y. Attention-based models for speech recognition. In *Advances in neural information processing systems*, 577–585 (2015).
48. <https://github.com/nokpil/agentnet> (2020).

49. Maggi, C. *et al.* Generalized energy equipartition in harmonic oscillators driven by active baths. *Physical Review Letters* **113**, 238303 (2014).
50. Delfau, J.-B., Ollivier, H., López, C., Blasius, B. & Hernández-García, E. Pattern formation with repulsive soft-core interactions: Discrete particle dynamics and dean-kawasaki equation. *Physical Review E* **94**, 042120 (2016).
51. Evangelista, D. J., Ray, D. D., Raja, S. K. & Hedrick, T. L. Data from: Three-dimensional trajectories and network analyses of group behaviour within chimney swift flocks during approaches to the roost (2018). URL <https://doi.org/10.5061/dryad.p68f8>.
52. Ma, Y., Guo, Z., Ren, Z., Tang, J. & Yin, D. Streaming graph neural networks. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 719–728 (2020).
53. Manessi, F., Rozza, A. & Manzo, M. Dynamic graph convolutional networks. *Pattern Recognition* **97**, 107000 (2020).
54. Pareja, A. *et al.* Evolvegcnn: Evolving graph convolutional networks for dynamic graphs. In *AAAI*, 5363–5370 (2020).
55. Smith, K., Ba, S. O., Odobez, J.-M. & Gatica-Perez, D. Tracking the visual focus of attention for a varying number of wandering people. *IEEE transactions on pattern analysis and machine intelligence* **30**, 1212–1229 (2008).
56. Martin, G. Bird collisions: a visual or a perceptual problem. *BOU Proceedings–Climate Change and Birds*. <http://www.bou.org.uk/bouproc-net/ccb/martin.pdf> (2010).

57. Ni, R. & Ouellette, N. Velocity correlations in laboratory insect swarms. *The European Physical Journal Special Topics* **224**, 3271–3277 (2015).
58. Nagy, M., Ákos, Z., Biro, D. & Vicsek, T. Hierarchical group dynamics in pigeon flocks. *Nature* **464**, 890–893 (2010).
59. Cavagna, A. *et al.* Scale-free correlations in starling flocks. *Proceedings of the National Academy of Sciences* **107**, 11865–11870 (2010).
60. Sokolov, A. & Aranson, I. S. Physical properties of collective motion in suspensions of bacteria. *Physical Review Letters* **109**, 248109 (2012).
61. Pastor-Satorras, R., Castellano, C., Van Mieghem, P. & Vespignani, A. Epidemic processes in complex networks. *Reviews of Modern Physics* **87**, 925 (2015).
62. Feng, Y., You, H., Zhang, Z., Ji, R. & Gao, Y. Hypergraph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 3558–3565 (2019).
63. Bai, S., Zhang, F. & Torr, P. H. Hypergraph convolution and hypergraph attention. *arXiv preprint arXiv:1901.08150* (2019).
64. Morris, C. *et al.* Weisfeiler and leman go neural: Higher-order graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 4602–4609 (2019).
65. Paszke, A. *et al.* Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop* (2017).
66. Hochreiter, S. & Schmidhuber, J. Long short-term memory. *Neural computation* **9**, 1735–1780 (1997).

67. Misra, D. Mish: A self regularized non-monotonic neural activation function. *arXiv preprint arXiv:1908.08681* (2019).
68. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

Acknowledgements This research was supported by the Basic Science Research Program through the National Research Foundation of Korea NRF-2017R1A2B3006930. We appreciate Y.J. Baek for providing insights for model demonstration and fruitful discussion.

Competing Interests The authors declare that they have no competing financial interests.

Correspondence Correspondence and requests for materials should be addressed to Hawoong Jeong (hjeong@kaist.edu).

Figures

Figure 1

Overview of system formulation and the neural network architecture of the proposed AgentNet. The correspondence between the decision rule of agents in a complex system and a forward pass of AgentNet is depicted. In both panels, the state variable of each agent s_i^t interacts with the state variables of other agents s_j^t in R_i with interaction strength a_{ij}^t . The graph attention core learns R_i with transformer architecture by encoding s_i^t into key k_i^t , query q_i^t , and value v_i^t , and then calculates the weighted sum of the values of other agents v_j^t according to the variable-wise attention weight a_{ij}^t as computed by neural attention. Different from GATs, AgentNet assigns attention value for each state variable, and decode it separately to strictly impose the information of variable-wise interaction strength. Other functions, namely h_{self} and f , can be captured by both encoder and decoder modules.

Figure 2

Result of AgentNet for cellular automata. (A) Attention weight transition of a single target cell throughout the training. In the initial stage, the model has no information about the interaction range and assigns near-zero values to all of the cells in the system. Attention gets narrowed down to a smaller region as training advances, and finally concentrates on eight surrounding cells, which is the theoretical interaction range. (B) Attention weight $\bar{\alpha}_c$ of neighbors and outside cells during 200 epochs of training. The attention weight of neighbor cells increases as training proceeds, while the weight of other cells remains 0. Data is averaged from 100 test samples. (C) AgentNet with respect to given alive (left) and dead microstates (right). The total number of alive cells in the neighborhood is denoted by s , which is the sole parameter of the CA decision rule.

Figure 3

Result of AgentNet for the Vicsek model. Initially, agents are randomly distributed in a circular region with radius $R = \sqrt{5}$, without any boundary condition. (A) Attention weight visu alization of two sample cells, a1 and a2. Both cases show a circular sector of attention distribution with clear boundaries that perfectly matches with ground-truth interaction range. (B) Averaged attentionweightforvariables x and y beforeandaftertraining. ThefullytrainedAgentNetlearned toidentifytheneighboragentsandignoretheirrelevantothersbyassigningnear-oneandnear-zero attention weights, respectively. (C) Position predictions by AgentNet for the two sample cells a1 and a2. Circles indicate the starting positions of the two particles, with the two heatmaps showing the AgentNet prediction along with the means of predicted distributions (Xs). The model predicts the expected theoretical distribution (crosses) with great precision, even when the given training samples (green and blue stars) are distant from the means of the theoretical distribution.

Figure 4

Result of AgentNet for AOUP.(A)In both panels, eight steps of the test data ($R = 4$) of fourparticlesaredrawnwithblackdotsstartingfromthelargeblackdots. AgentNetpredictionsof thetrajectoriesinthefollowing12steps(rightpanel)perfectlycoincidewiththesampletrajectories from the true Langevin equation (left panel). 100 samples are drawn in both panels, and the final positions are highlighted with white stars. (B) Equilibrium state of a system with $R = 5$, which is unseen at the training stage. A single realization from the true distribution is drawn with final positions marked by red stars (left), while a single sample from the predicted distribution of AgentNet is drawn with final positions marked by blue stars (right). AgentNet for AOUP captures the generalized effect of interaction length R and predicts the collective behavior of the untrained system. (C) Exerted x-directional force $F_{int,x}$ and x-directional velocity attention α^{v_x} shows a strong linear relationship, while single attention value from

GAT does not capture any of the force component. Same holds for (D), for the case of y-direction. (E) By plotting relative distance r_{ij} versus force and attention, scaled attention shows good coincidence with the force value up to constant factor $c = 0.28$. (F) Visualization of $\bar{\alpha} v_x$ for a single target particle (blue). Attention and force values for (C) to (E) are collected from 100 test samples with $R = 4$.

Figure 5

Result of AgentNet for CS. (A) Displacement error of linear extrapolation, naive LSTM, GAT, GAT3+, and AgentNet. AgentNet shows the lowest displacement error compared to the baselines. Here, the final displacement error (FDE) of step n indicates the averaged error of birds for which their trajectories terminated at step n . All of the results are averaged value from three trials. (B) Exemplary snapshot of the visualized attention of a single agent (blue circle) from the test dataset. (C) Two-dimensional heatmap of averaged attention αv_x and cosine similarity of the velocity with respect to the relative coordinates. We align every bird's heading direction in the test dataset to the x-axis (blue arrow) and draw cross-sections in the xy-plane (upper panels) and xz plane (lower panels). Different from the velocity correlation, attention shows more concentrated and strongly directional distributions that coincide with previous literature about the bird's visual frustum and sight direction. For attention, a contour of the top 0.01% of the attention value is visualized (red, dashed) as well as the direction of the maximum attention value (red arrows).

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [NatComPhysSI.pdf](#)