Learning interaction rules from multi-animal trajectories via augmented behavioral models

Anonymous Author(s)
Affiliation
Address
email

Abstract

Extracting the interaction rules of biological agents from moving sequences pose challenges in various domains. Granger causality is a practical framework for analyzing the interactions from observed time-series data; however, this framework ignores the structures of the generative process in animal behaviors, which may lead to interpretational problems and sometimes erroneous assessments of causality. In this paper, we propose a new framework for learning Granger causality from multi-animal trajectories via augmented theory-based behavioral models with interpretable data-driven models. We adopt an approach for augmenting incomplete multi-agent behavioral models described by time-varying dynamical systems with neural networks. For efficient and interpretable learning, our model leverages theory-based architectures separating navigation and motion processes, and the theory-guided regularization for reliable behavioral modeling. This can provide interpretable signs of Granger-causal effects over time, i.e., when specific others cause the approach or separation. In experiments using synthetic datasets, our method achieved better performance than various baselines. We then analyzed multi-animal datasets of mice, flies, birds, and bats, which verified our method and obtained novel biological insights.

1 Introduction

Extracting the interaction rules of real-world agents from data is a fundamental problem in a variety of scientific and engineering fields. For example, animals, vehicles, and pedestrians observe other's states and execute their actions in complex situations. Discovering the directed interaction rules of such agents from observed data will contribute to the understanding of the principles of biological agents' behaviors. Among methods analyzing directed interactions within multivariate time series, Granger causality (GC) [22] is a practical framework for exploratory analysis [48] in various fields, such as neuroscience [63] and economics [3] (see Section 2). Recent methodological developments have focused on inferring GC under nonlinear dynamics (e.g., [71, 30, 79, 54, 42]).

However, the structure of the generative process in biological multi-agent trajectories, which include navigational and motion processes [53] regarded as time-varying dynamical systems (see Section 3.1), is not fully utilized in existing base models of GC including vector autoregressive [26] and recent neural models [71, 30, 79]. Ignoring the structures of such processes in animal behaviors will lead to interpretational problems and sometimes erroneous assessments of causality. That is, incorporating the structures into the base model for inferring GC, e.g., augmenting (inherently) incomplete behavioral models with interpretable data-driven models (see Section 3.2), can solve these problems. Furthermore, since data-driven models sometimes detect false causality that is counterintuitive to the user of the analysis, e.g., introducing architectures and regularization to utilize scientific knowledge (see Sections 3.2 and 4.2) will be effective for a reliable base model of a GC method.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

In this paper, we propose a framework for learning GC from biological multi-agent trajectories via augmented behavioral models (ABM) using interpretable data-driven neural models. We adopt an approach for augmenting incomplete multi-agent behavioral models described by time-varying dynamical systems with neural networks (see Section 3.2). The ABM leverages theory-based architectures separating navigation and motion processes based on a well-known conceptual behavioral model [53], and the theory-guided regularization (see Section 4.2) for interpretable and reliable behavioral modeling. This framework can provide interpretable signs of Granger-causal effects over time, e.g., when specific others cause the approach or separation.

The main contributions of this paper are as follows: (1) we propose a framework for learning Granger causality via ABM, which can extract interaction rules from real-world multi-agent trajectory data; (2) methodologically, for efficient and interpretable learning, the behavioral model leverages theory-based architectures separating navigation and motion processes based on a conceptual behavioral model [53], and the theory-guided regularization for reliable behavioral modeling; and (3) in the experiments, our method achieved better performance than various baselines using synthetic datasets, and obtained new biological insights and verified our method using multiple datasets of mice, birds, bats, and flies. In the remainder of this paper, we describe the background of GC in Section 2. Next, we formulate our ABM in Section 3, and the learning and inference methods in Section 4. We overview related works in Section 5, present experimental results in Section 6, and conclude this paper in Section 7.

2 Granger Causality

GC [22] is one of the most popular and practical approaches to infer directed causal relations from observational multivariate time series data. Although the classical GC is defined by linear models, here we introduce a more recent definition of [71] for non-linear GC. Consider p stationary time-series $\boldsymbol{x} = \{\boldsymbol{x}^1, ...\boldsymbol{x}^p\}$ across timesteps $t = \{1, ..., T\}$ and a non-linear autoregressive function g_j, such that

$$\boldsymbol{x}^j_{t+1} = g_j(\boldsymbol{x}^1_{\leq t}, ..., \boldsymbol{x}^p_{\leq t}) + \boldsymbol{\varepsilon}^j_t, \tag{1}$$

where $\boldsymbol{x}^j_{\leq t} = (..., \boldsymbol{x}^j_{t-1}, \boldsymbol{x}^j_t)$ denotes the present and past of series j and $\boldsymbol{\varepsilon}^j_t$ represents independent noise. We then consider that variable $\boldsymbol{x}^i$ does not Granger-cause variable $\boldsymbol{x}^j$, denoted as $\boldsymbol{x}^i \nrightarrow \boldsymbol{x}^j$, if and only if $g_j(\cdot)$ is constant in $\boldsymbol{x}^i_{\leq t}$. Granger causal relations are equivalent to causal relations in the underlying directed acyclic graph if all relevant variables are observed and no instantaneous (i.e., connections between two variables at the same timestep) connections exist [58]. Many methods for Granger causal discovery, including vector autoregressive [26] and recent deep learning-based approaches [71, 30, 79], can be encapsulated by the following framework. First, we define a function f_θ (e.g., an multilayer perceptrons (MLP) in [71], a linear model in [26]), which learns to predict the next time-step of the test sequence $\boldsymbol{x}$. Then, we fit f_θ to $\boldsymbol{x}$ by minimizing some loss $\mathcal{L}$: $\theta_\star = \text{argmin}_\theta\, \mathcal{L}(\boldsymbol{x}, f_\theta)$. Finally, we apply some fixed function h (e.g., thresholding) to the learned parameters to produce a Granger causal graph estimate for $\boldsymbol{x}$: $\hat{\mathcal{G}}_{\boldsymbol{x}} = h(\theta_\star)$.

Furthermore, we need to differentiate between positive and negative Granger-causal effects (e.g, approaching and separating). Based on the definition of [44], we define the effect sign as follows: if $g_j(\cdot)$ is increasing in all $\boldsymbol{x}^i_{\leq t}$, then we say that variable $\boldsymbol{x}^i$ has a positive effect on $\boldsymbol{x}^j$, if $g_j(\cdot)$ is decreasing in $\boldsymbol{x}^i_{\leq t}$, then $\boldsymbol{x}^i$ has a negative effect on $\boldsymbol{x}^j$. Note that $\boldsymbol{x}^i$ can contribute both positively and negatively to the future of $\boldsymbol{x}^j$ at different delays.

Overall, the causality measures, however elaborate in construction, are simply statistics estimated from a model [69]. If the model inadequately represents the system properties of interest, subsequent analyses based on the model will fail to address the question of interest. The inability of the model to represent key features of interest can cause interpretational problems and sometimes erroneous assessments of causality. Therefore, in our case, incorporating the structures of the generative process in animal behaviors will be required. We thus propose the ABM based on a well-known conceptual model [53] in biological sciences in the next section.

3 Augmented behavioral model

Since ignoring the structures of the generative process causes interpretational problems and sometimes erroneous assessments of causality, we incorporate the structures into the base model of a GC

framework. In this section, we firstly formulate a well-known conceptual behavioral model [53] so that it can be computable. Second, we propose (multi-animal) ABMs with theory-based architectures based on scientific knowledge. Further, we discuss the relation to the existing explainable neural models [2].

3.1 Formulation of a conceptual behavioral model

In movement ecology, which is a branch of biology concerning the spatial and temporal patterns of behaviors of organisms, a coherent framework [53] has been conceptualized to explore the causes, mechanisms, and patterns of movement. For example, two alternative structural representations [53] were proposed to model a new position $\boldsymbol{p}_{t+1}$ from its current location $\boldsymbol{p}_t$ (for details, see Appendix A): the motion-driven case $\boldsymbol{p}_{t+1} = f_U(f_M(\boldsymbol{\Omega}, f_N(\boldsymbol{\Phi}, \boldsymbol{r}_t, \boldsymbol{w}_t, \boldsymbol{p}_t), \boldsymbol{r}_t, \boldsymbol{w}_t, \boldsymbol{p}_t)) + \boldsymbol{\varepsilon}_t$, and the navigation-driven case $\boldsymbol{p}_{t+1} = f_U(f_N(\boldsymbol{\Phi}, f_M(\boldsymbol{\Omega}, \boldsymbol{r}_t, \boldsymbol{w}_t, \boldsymbol{p}_t), \boldsymbol{r}_t, \boldsymbol{w}_t, \boldsymbol{p}_t)) + \boldsymbol{\varepsilon}_t$, where $\boldsymbol{w}_t$ is the internal state, $\boldsymbol{\Omega}$ is the motion capacity, $\boldsymbol{\Phi}$ is the navigation capacity, and $\boldsymbol{r}_t$ is the environmental factors (these are conceptual parameters). f_M, f_N, and f_U are conceptual functions to represent actions of the motion (or planning), navigation, and movement progression processes, respectively.

For efficient learning to use only one time series in most (neural) GC frameworks [71, 30, 44], in this paper, we consider a simple case with homogeneous navigation and motion capacities, and internal states. Moreover, to make the contribution of f_M, f_N, and f_U interpretable after training from the data without assuming either motion-driven or navigation-driven case, one of the simplified processes for agent i is represented by

$$\boldsymbol{x}^i_{t+1} = f^i_U(f^i_N(\boldsymbol{r}^i_t, \boldsymbol{x}^i_t), f^i_M(\boldsymbol{r}^i_t, \boldsymbol{x}^i_t), \boldsymbol{r}^i_t, \boldsymbol{x}^i_t) + \boldsymbol{\varepsilon}^i_t, \tag{2}$$

where $\boldsymbol{x}^i \in \mathbb{R}^d$ includes location $\boldsymbol{p}^i$ and velocity for the agent i. We here consider $\boldsymbol{r}^i \in \mathbb{R}^{(p-1)d_r}$ including $p-1$ other agents' d_r-dimensional information. Such behaviors have been conventionally modeled by mathematical equations such as force- and rule-based models (e.g., reviewed by [75, 46]). Recently, these models have become more sophisticated by incorporating the models into hand-crafted functions representing anticipation (e.g., [28, 50]) and navigation (e.g., [8, 74]).

However, these conventional and recent models are sometimes too simplistic and customized for the specific animals, respectively; thus it is sometimes difficult to define the dynamics of general biological multi-agent systems (i.e., multiple species of animals). Therefore, methods for learning parameters and interaction rules of behavioral models are needed. There have been some researches to estimate specific parameters (and their distributions) of the interpretable behavior models (e.g., [81, 82, 15]), and others to model the parameters and rules in purely data-driven manners (i.e., not interpretable) only for accurate prediction (e.g., [16, 27]). In the proposed framework, we consider flexible data-driven interpretable models to focus on inferring GC for exploratory analysis from the observed data without specific knowledge of the species and obtaining additional data.

Recently, some attempts have been made to explore flexible and interpretable models bridging theory-based and data-driven approaches. For example, a paradigm called theory-guided data science has been proposed [29], which leverages the wealth of scientific knowledge for improving the effectiveness of data-driven models in enabling scientific discovery. For example, scientific knowledge can be used as architectures or regularization terms in learning algorithms in physical and biological sciences (e.g., [59, 20]). In biological multi-agent systems, an approach extract interpretable dynamical information based on physics-based knowledge [18] from multi-agent interaction data, and another approach made a particular module such as observation (e.g., [19]) interpretable in mostly black-box neural models. However, these data-driven models did not sufficiently utilize the above scientific knowledge of multi-animal interactions. In the next subsection, to make the model (e.g., of GC) flexible and interpretable, we propose an ABM with theory-based architectures.

3.2 Augmented behavioral model with theory-based architectures

In this subsection, we propose a ABM using interpretable neural models with theory-based architectures for learning GC from multi-animal trajectories. In general, it is scientifically beneficial if a model mimics the data-generating process well, e.g., because existing scientific insights can be leveraged or revalidated. In our case of GC, additionally, it is expected to eliminate obvious erroneous causality by utilizing existing knowledge, we thus propose a theory-based ABM for learning GC.

Generally, scientific knowledge can be used to influence the architecture of data-driven scientific models. Most design considerations are mainly motivated to simplify the learning procedure, minimize the training loss, and ensure robust generalization performance [29]. In some cases, domain

knowledge can be used designing neural models by decomposing the overall problem into modular sub-problems. For example, in our problem, to describe the overall process of multi-animal behaviors, modular neural models can be learned for different sub-processes, such as the navigation, planning, and movement processes (f_N^i, f_M^i, and f_U^i, respectively) described in Section 3.1. This will help in using the power of learning frameworks while following a high-level organization in the architecture that is motivated by domain knowledge [29]. Specifically, to accurately model the relationships between agents (finally interpreted as causal relationships) with limited information in usual GC settings, we explicitly formulate the f_N^i, f_M^i, and f_U^i and estimate f_N^i and f_M^i from data.

In summary, our base ABM can be expressed as

$$\boldsymbol{x}_t^i = \sum_{k=1}^{K} \left(F_N^{i,t,k}(\boldsymbol{h}_{t-k}^i) \odot F_M^{i,t,k}(\boldsymbol{h}_{t-k}^i) \right) \boldsymbol{h}_{t-k}^i + \boldsymbol{\varepsilon}_t^i, \tag{3}$$

where $\boldsymbol{h}_{t-k}^i \in \mathbb{R}^{d_h}$ is a vector concatenating the self state $\boldsymbol{x}_{t-k}^i \in \mathbb{R}^d$ and all others' state $\boldsymbol{r}_{t-k}^i \in \mathbb{R}^{(p-1)d_r}$, and $\odot$ denotes a element-wise multiplication. K is the order of the autoregressive model. $F_N^{i,t,k}, F_M^{i,t,k} : \mathbb{R}^{d_h} \to \mathbb{R}^{d \times d_h}$ are matrix-valued functions that represent navigation and motion functions, which are implemented by MLPs. For brevity, we omit the intercept term here and in the following equations. The value of the element of $F_N^{i,k}$ is $[-1, 1]$ is like a switching function value, i.e., a positive or negative sign to represent the approach and separation from others. The value of the element of $F_M^{i,k}$ is a positive value or zero, which changes continuously and represents coefficients of time-varying dynamics. Relationships between agents $\boldsymbol{x}^1, ..., \boldsymbol{x}^p$ and their variability throughout time can be examined by inspecting coefficient matrices $\boldsymbol{\Psi}_{\boldsymbol{\theta}_{t,k}}^i = \left(F_N^{i,t,k}(\boldsymbol{h}_{t-k}^i) \odot F_M^{i,t,k}(\boldsymbol{h}_{t-k}^i) \right)$. We separate $\boldsymbol{\Psi}_{\boldsymbol{\theta}_{t,k}}^i$ into $F_N^{i,t,k}(\boldsymbol{h}_{t-k}^i)$ and $F_M^{i,t,k}(\boldsymbol{h}_{t-k}^i)$ for two reasons: interpretability and efficient use of scientific knowledge. The interpretability of two coefficients $F_N^{i,k}$ and $F_M^{i,k}$ contributes to the understanding of navigation and motion planning processes of animals (i.e., signs and amplitudes in the GC effects), respectively. The efficient use of scientific knowledge in the learning of a model enables us to incorporate the knowledge into the model. The effectiveness was shown in the ablation studies in Appendix G. Specific forms of Eq. (3) are described in Appendices D.2 and F.2. The formulation of the model via linear combinations of the interpretable feature $\boldsymbol{h}_{t-k}^i$ for an explainable neural model is related to the self-explanatory neural network (SENN) [2].

3.3 Relation to self-explanatory neural network

SENN [2] was introduced as a class of intrinsically interpretable models motivated by explicitness, faithfulness, and stability properties. A SENN with a link function $g(\cdot)$ and interpretable basis concepts $h(\boldsymbol{x}) : \mathbb{R}^p \to \mathbb{R}^u$ follows the form

$$f(\boldsymbol{x}) = g(\theta(\boldsymbol{x})_1 h(\boldsymbol{x})_1, ..., \theta(\boldsymbol{x})_u h(\boldsymbol{x})_u), \tag{4}$$

where $\boldsymbol{x} \in \mathbb{R}^p$ are predictors; and $\theta(\cdot)$ is a neural network with u outputs (here, we consider the simple case of $d = 1$ and $d_r = 1$). We refer to $\theta(\boldsymbol{x})$ as generalized coefficients for data point $\boldsymbol{x}$ and use them to *explain* contributions of individual basis concepts to predictions. In the case of $g(\cdot)$ being sum and concepts being raw inputs, Eq. (4) simplifies to $f(\boldsymbol{x}) = \sum_{i=1}^{p} \theta(\boldsymbol{x})_i \boldsymbol{x}_i$. In this paper, we regard the movement function f_U^i as $g(\cdot)$ and the function of f_N^i and f_M^i as θ for the following interpretable modeling of f_U^i, f_N^i, and f_M^i. Appendix B presents additional properties SENNs need to satisfy and the learning algorithm, as defined by [2]. Note that our model does not always satisfy the requirements of SENN [2, 44] due to the modeling of time-varying dynamics (see Appendix B). SENN was first applied to GC [44] via generalized vector autoregression model (GVAR): $\boldsymbol{x}_t = \sum_{k=1}^{K} \boldsymbol{\Psi}_{\boldsymbol{\theta}_k}(\boldsymbol{x}_{t-k}) \boldsymbol{x}_{t-k} + \boldsymbol{\varepsilon}_t$, where $\boldsymbol{\Psi}_{\boldsymbol{\theta}_k} : \mathbb{R}^p \to \mathbb{R}^{p \times p}$ is a neural network parameterized by $\boldsymbol{\theta}_k$. $\boldsymbol{\Psi}_{\boldsymbol{\theta}_k}(\boldsymbol{x}_{t-k})$ is a matrix whose components correspond to the generalized coefficients for lag k at timestep t. The component (i, j) of $\boldsymbol{\Psi}_{\boldsymbol{\theta}_k}(\boldsymbol{x}_{t-k})$ corresponds to the influence of $\boldsymbol{x}_{t-k}^j$ on $\boldsymbol{x}_t^i$. However, the SENN model did not use scientific knowledge of multi-element interactions and may cause interpretational problems and sometimes erroneous assessments of causality.

4 Learning with theory-guided regularization and inference

Here, we describe the learning method of the ABM including theory-guided regularization. We first overview the learning method and define the objective function. We then explain the theory-guided

regularization for incorporating scientific knowledge into the learning of the model. Finally, we describe the inference of GC by our method.

4.1 Overview

To mitigate the inference in multivariate time series, Eq. (3) for each agent is summarized as the following expression:

$$\boldsymbol{x}_t = \sum_{k=1}^{K} \left[\left(F_N^{1,t,k}(\boldsymbol{h}_{t-k}^1) \odot F_M^{1,t,k}(\boldsymbol{h}_{t-k}^1) \right) \boldsymbol{h}_{t-k}^1, \ldots, \left(F_N^{p,t,k}(\boldsymbol{h}_{t-k}^p) \odot F_M^{p,t,k}(\boldsymbol{h}_{t-k}^p) \right) \boldsymbol{h}_{t-k}^p \right] + \boldsymbol{\varepsilon}_t, \quad (5)$$

where $\boldsymbol{x}_t$ and $\boldsymbol{\varepsilon}_t$ concatenate the original variables for all p agents. We train our model by minimizing the following penalized loss function with the mini-batch gradient descent

$$\sum_{t=K+1}^{T} \left(\mathcal{L}_{pred}(\hat{\boldsymbol{x}}_t, \boldsymbol{x}_t) + \lambda \mathcal{L}_{sparsity}(\boldsymbol{\Psi}_t) + \gamma \mathcal{L}_{TG}(\boldsymbol{\Psi}_t, \boldsymbol{\Psi}_t^{TG}) \right) + \sum_{t=K+1}^{T-1} \beta \mathcal{L}_{smooth}(\boldsymbol{\Psi}_{t+1}, \boldsymbol{\Psi}_t), \quad (6)$$

where $\{\boldsymbol{x}_t\}_{t=1}^T$ is a single observed time series of length T with d-dimensions and p-agents; $\hat{\boldsymbol{x}}_t$ is the one-step forecast for the t-th time point based on Eq. (5); $\boldsymbol{\Psi}_t \in \mathbb{R}^{pd \times K d_h}$ is defined as a concatenated matrix of $[\boldsymbol{\Psi}^1_{\boldsymbol{\theta}_{t,K}}, \ldots, \boldsymbol{\Psi}^1_{\boldsymbol{\theta}_{t,1}}], \ldots, [\boldsymbol{\Psi}^p_{\boldsymbol{\theta}_{t,K}}, \ldots, \boldsymbol{\Psi}^p_{\boldsymbol{\theta}_{t,1}}]$ in a row; $\boldsymbol{\Psi}_t^{TG}$ is a coefficient determined by the following theory-guided regularization; and $\lambda, \beta, \gamma \geq 0$ are regularization parameters. The loss function in Eq. (6) consists of four terms: (i) the mean squared error (MSE) prediction loss, (ii) a sparsity-inducing penalty term, (iii) theory-guided regularization, and (iv) the smoothing penalty term. The sparsity-inducing term $\mathcal{L}_{sparsity}$ is an appropriate penalty on the norm of $\boldsymbol{\Psi}_t$. Among possible various regularization terms, in our implementation, we employ the elastic-net-style penalty term [83, 55] $\mathcal{L}_{sparsity}(\boldsymbol{\Psi}_t) = \frac{1}{T-K} \left(\alpha \left\| \boldsymbol{\Psi}_t \right\|_1 + (1-\alpha) \left\| \boldsymbol{\Psi}_t \right\|_F^2 \right)$, with $\alpha = 0.5$, based on [44]. Note that other penalties can be also easily adapted to our model. The smoothing penalty term, given by $\mathcal{L}_{smooth}(\boldsymbol{\Psi}_{t+1}, \boldsymbol{\Psi}_t) = \frac{1}{T-K-1} \left\| \boldsymbol{\Psi}_{t+1} - \boldsymbol{\Psi}_t \right\|_F^2$, is the average norm of the difference between generalized coefficient matrices for two consecutive time points. This penalty term encourages smoothness in the evolution of coefficients with respect to time [44]. To avoid overfitting and model selection problems, we eliminate unused factors based on prior knowledge (for details, see Appendices D.2 and F.2).

4.2 Theory-guided regularization

The third term in Eq. (6) is the theory-guided regularization for reliable Granger causal discovery by leveraging regularization with scientific knowledge. Here we utilize theory-based and data-driven prediction results and impose penalties in the appropriate situations as described below. Again, let $\hat{\boldsymbol{x}}_t$ be the prediction from the data. In addition to the data, we prepare some input-output pairs $(\tilde{\boldsymbol{x}}_{t-k \leq t}, \tilde{\boldsymbol{x}}_t)$ based on scientific knowledge. We call them pairs of theory-guided feature and prediction, respectively. In this case, we assume that the theory-guided cause or weight of the ABM $\boldsymbol{\Psi}_t^{TG}$ is uniquely determined. When the difference between $\hat{\boldsymbol{x}}_t$ and $\tilde{\boldsymbol{x}}_t$ is below a certain threshold, we assume that the cause (weight) of $\hat{\boldsymbol{x}}_t$ is equivalent to the cause of $\tilde{\boldsymbol{x}}_t$.

In animal behaviors, the theory-guided prediction utilizes the intuitive prior knowledge such that the agents go straight from the current state if there is no interaction. In this case, $\tilde{\boldsymbol{x}}_t$ includes the same velocity as the previous step and the corresponding positions after going straight. The penalty is expressed as $\mathcal{L}_{TG}(\boldsymbol{\Psi}_t, \boldsymbol{\Psi}_t^{TG}) = \frac{1}{T-K} \exp(\|\boldsymbol{x}_t - \tilde{\boldsymbol{x}}_t\|_2^2/\sigma) \|\boldsymbol{\Psi}'_t\|_F^2$, where $\boldsymbol{\Psi}'_t \in \mathbb{R}^{pd \times K(p-1)d_r}$ is the weight matrix regarding others' information (i.e., eliminating the information of the agents themselves from $\boldsymbol{\Psi}_t$) and σ is a parameter regarding the threshold. Note that here the matrix $\boldsymbol{\Psi}_t^{'TG}$ corresponding to $\boldsymbol{\Psi}_t^{TG}$ is a zero matrix representing no interaction with others (i.e., $\|\boldsymbol{\Psi}'_t - \boldsymbol{\Psi}_t^{'TG}\|_F^2 = \|\boldsymbol{\Psi}'_t\|_F^2$).

Next, we can consider the general cases. All possible combinations of the pairs are denoted as the direct product $\mathcal{H}_0 := L \times M \times \cdots \times M = \{(l, m_1, \ldots, m_p) \mid l \in L \wedge m_1 \in M \wedge \cdots \wedge m_p \in M\}$, where $L = \{1, \ldots, p\}$ and $M = \{-1, 0, 1\}$ if we consider the sign of Granger causal effects (otherwise, $M = \{0, 1\}$). However, if we consider the pairs $(\tilde{\boldsymbol{x}}_{t-k \leq t}, \tilde{\boldsymbol{x}}_t)$ uniquely determined, it will be a considerably fewer number of combinations by avoiding underdetermined problems. We denote the set of the uniquely-determined combinations as $\mathcal{H}_1 \subset \mathcal{H}_0$. We can then impose penalties on the weights: $\mathcal{L}_{TG}(\boldsymbol{\Psi}_t, \boldsymbol{\Psi}_t^{TG}) = \frac{1}{|\mathcal{H}_1|(T-K)} \sum_{l, m_1, \ldots, m_p \in \mathcal{H}_1} \left(\exp(\|\boldsymbol{x}_t - \tilde{\boldsymbol{x}}_t\|_F^2/\sigma) \|\boldsymbol{\Psi}'_t - \boldsymbol{\Psi}_{t,l,m_1,\ldots,m_p}^{'TG}\|_F^2 \right)$, where $\boldsymbol{\Psi}_{t,l,m_1,\ldots,m_p}^{'TG} \in$

$\mathbb{R}^{pd \times K(p-1)d_r}$ is the weight matrix regarding others' information in $\mathbf{\Psi}_t$. In animal behaviors, due to unknown terms, such as inertia and other biological factors, the theory-guided prediction utilizes the only intuitive prior knowledge such that the agents go straight from the current state if there are no interactions (i.e., $|\mathcal{H}_1| = 1$).

4.3 Inference of Granger causality

Once $\mathbf{\Psi}_t$ is trained, we quantify strengths of Granger-causal relationships between variables by aggregating matrices $\mathbf{\Psi}_t$ across all K, d, d_r, t into summary statistics. Although most neural GC methods [71, 54, 30, 79, 44] did not provide an obvious way for handling multi-dimensional time series (i.e., $d > 1$), our main problems include two- or three-dimensional positional and velocity data for each animal. Therefore, we compute the norm with respect to spatial dimensions d, d_r, and the sign of the GC separately. That is, we aggregate the obtained generalized coefficients into matrix $\boldsymbol{S} \in \mathbb{R}^{p \times p}$ as follows:

$$S_{i,j} = \text{sign} \left\{ \max_{K+1 \le t \le T} \left(\max_{1 \le k \le K} \left(\underset{\substack{q=1,\dots,d_r \\ u=1,\dots,d}}{\text{median}} \left(\mathbf{\Psi}_{i,j} \right) \right) \right) \right\} \max_{K+1 \le t \le T} \left(\max_{1 \le k \le K} \left(\| \left(\mathbf{\Psi}_{i,j} \right)_{t,k} \|_F \right) \right), \quad (7)$$

where $\mathbf{\Psi}_{i,j} \in \mathbb{R}^{(T-K) \times K \times d \times d_r}$ is computed by reshaping and concatenating $\mathbf{\Psi}_t$ over $K+1 \le t \le T$. $\| \left(\mathbf{\Psi}_{i,j} \right)_{t,k} \|_F$ is the Frobenius norm of the matrix $\left(\mathbf{\Psi}_{i,j} \right)_{t,k} \in \mathbb{R}^{d \times d_r}$ for each t, k. If we do not consider the sign of Granger causal effects, we ignore the coefficient of the signed function. If we investigate the GC effects over time, we eliminate max function among t. Note that we only consider off-diagonal elements of adjacency matrices and ignore self-causal relationships. Intuitively, $S_{i,j}$ are statistics that quantify the strength of the Granger-causal effect of $\boldsymbol{x}^i$ on $\boldsymbol{x}^j$ using magnitudes of generalized coefficients. We expect $S_{i,j}$ to be close to 0 for non-causal relationships and $S_{i,j} \gg 0$ if $\boldsymbol{x}^i \to \boldsymbol{x}^j$. Note that in practice $S_{i,j}$ is not binary-valued, as opposed to the ground truth, which we want to infer, because the outputs of $\mathbf{\Psi}_{i,j}$ are not shrunk to exact zeros. Therefore, we need a procedure deciding for which variable pairs $S_{i,j}$ are significantly different from 0. To infer a binary matrix of GC relationships, we use a heuristic threshold. For the detail, see Appendix C.1.

5 Related work

Methods for nonlinear GC. Initial work for nonlinear GC methods focused on time-varying dynamic Bayesian networks [68], regularized logistic regression with time-varying coefficients [33], and kernel-based regression models [45, 67, 38]. Recent approaches to inferring Granger-causal relationships leverage the expressive power of neural networks [49, 77, 71, 54, 30, 42, 79] and are often based on regularized autoregressive models. Methods using sparse-input MLPs and long short-term memory to model nonlinear autoregressive relationships have been proposed [71], followed by a more sample efficient economy statistical recurrent unit (eSRU) architecture [30]. Other researchers proposed a temporal causal discovery framework that leverages attention-based convolutional neural networks [54] and a framework to interpret signs of GC effects and their variability through time building on SENN [2]. However, the structure of time-varying dynamical systems in multi-animal trajectories was not fully utilized in the above models.

Information-theoretic analysis for multi-animal motions. In this topic, most researchers have adopted transfer entropy (TE) and its variants and have analyzed them in terms of e.g., information cascades rather than causal discovery among animals. In the pioneering work, [76] analyzed information cascades among artificial collective motions using (conditional) TE [39, 40]. [62] applied variants of conditional mutual information to identify dynamical coupling between the trajectories of foraging meerkats. TE has been used to study the response of schools of zebrafish to a robotic replica of the animal [9, 35], to infer leadership in pairs of bats [57] and simulated zebrafish [10], and to identify interactions in a swarm of insects (Chironomus riparius) [41]. Local TE (or pointwise TE) [65, 39] has been used to detect local dependencies at specific time points in a swarm of soldier crabs [73], teams of simulated RoboCup agents [11], and a school of fish [13]. Since biological collective motions are intrinsically time-varying dynamical systems, we compared our methods and local TE in our experiments.

Other Biological multi-agent motion analysis. Previous studies have investigated leader-follower relationships. For example, the existences of the leadership have been investigated via the correlation in movement with time delay (e.g., [52, 64]) and via global physical (e.g., [4]) and statistical

properties [56]. Meanwhile, methods for data-driven biological multi-agent motion modeling have been intensively investigated for pedestrian (e.g., [1, 24]), vehicles (e.g., [5, 61, 70]), animals [16, 27], and athletes (e.g., [80, 36]). In most of these methods, the agents are assumed to have the full observation of other agents to achieve accurate prediction. In contrast, some researches have modeled partial observation in real-world multi-agent systems [25, 37, 23, 17, 18, 19, 21]. However, the above approaches required a large amount of training data and would not be suitable for application to multi-animal trajectory datasets that can only be measured in small quantities.

6 Experiments

The purpose of our experiments is to validate the proposed methods for application to real-world multi-animal moving trajectories, which have usually a smaller amount of sequences and no ground truth of the interaction rules. Thus, for verification of our methods, we first compared their performances to infer the Granger causality to those in various baselines using two synthetic datasets with ground truth: nonlinear oscillator (Kuramoto model) and boid model simulation datasets. We used the same ABM as applied to real-world multi-animal trajectory datasets: mice, birds, bats, and flies. To demonstrate the applicability to the multi-element dynamics other than multi-animal trajectories, we validated our method using the Kuramoto dataset (the results are shown in Appendix E). Each method was trained only on one sequence according to most neural GC frameworks [71, 30, 44]. The hyperparameters of the models were determined by validation datasets in the synthetic data experiments (for the details, see Appendices D and F). The common training details, (binary) inference methods, computational resources, and the amount of computation are described in Appendix C.

6.1 Synthetic datasets

For verification of our method, we compared the performances to infer the GC to those in various baselines using two synthetic datasets with ground truth. To compare with various baselines of GC methods, we tackled problems where the true causality is not changed over time. Here, we compared our methods (ABM) to 5 baseline methods: economic statistical recurrent unit (eSRU) [30]; amortized causal discovery (ACD) [42]; and GVAR [44], and simple baselines such as linear GC and local TE modified from [79, 42]. Except for ACD [42], most baselines did not provide an obvious way for handling multi-dimensional time series, whereas our main problems include two- or three-dimensional trajectories for each animal. Therefore, we modified the baselines except for ACD to compute norms with respect to spatial dimensions (2 or 3) for comparability with the proposed method. Note that the interpretations of the relationships estimated by ACD and Local TE are different from other methods, thus the sign of the relationship could not be investigated (we denote such by N/A in Table 1).

We investigated the continuously-valued inference results: the norms of relevant weights, scores, and strengths of GC relationships. We compared these scores against the true structures using areas under receiver operating characteristic (AUROC) and precision-recall (AUPRC) curves. We also evaluated thresholded inference results: accuracy (Acc) and balanced accuracy (BA) scores. For the inference methods of a binary matrix of GC relationships, see Appendix C.1. For all evaluation metrics, we only considered off-diagonal elements of adjacency matrices, ignoring self-causal relationships.

Boid model. Here, we validated our method on the simulation data using the boid model, which contains five agents moving trajectories (for details, see Appendix F). In this experiment, we set the boids (agents) directed preferences: we ran-

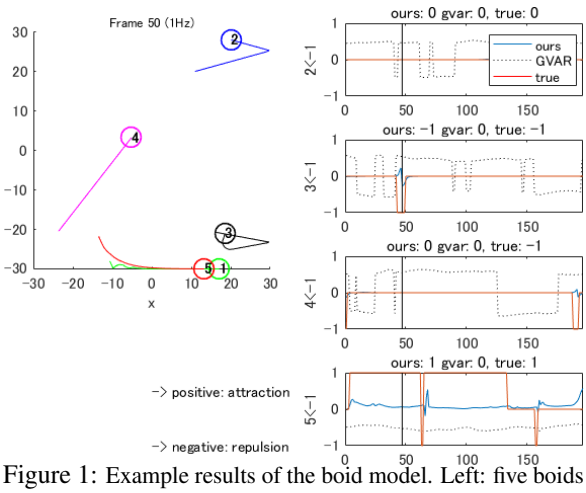

Figure 1: Example results of the boid model. Left: five boids (agents) movements. Trajectories are histories of the movement. Right: the results of our method (blue) and GVAR [44] (black dash) for the relationship between the cause (boid #1) and effects (other boids; i.e., $S_{i,j}$ for $i = 1$ and $j = 2, \ldots, 5$). The binary relationships are described in the upper of the plot. $1, 0$, and -1 indicate attraction, no interaction, and repulsion, respectively. Note that the magnitudes of our method and GVAR [44] were normalized with their maximal values (thus, the values were not be comparable among methods and red ground truth). For the detail, see the main text.

domly set the ground truth relationships $1, 0$, and -1 as the rules of attraction, no interaction, and repulsion, respectively. Figure 1 illustrates that e.g., boid #5 was attracted to boid #1 (i.e., true relationship: 1) and boid #3 avoided boid #1 (true relationship: -1). In this figure, our method detected the changes in the signed relationships whereas the GVAR [44] did not.

The performances were evaluated using $S_{i,j}$ in Eq. (7) throughout time because our method and ground truth were sensitive the sign as shown in Figure 1. Table 1 shows that our method achieved better performance than various baselines. The ablation studies shown in Appendix G reveal that the main two contributions of this work, the theory-guided regularization $\mathcal{L}_{TG}$ and learning navigation function $\boldsymbol{F}_N^k$ and motion function $\boldsymbol{F}_M^k$ separately, improved the performance greatly. These suggest that the utilization of scientific knowledge via the regularization and architectures efficiently worked in the limited data situations. Similarly, the results of the Kuramoto dataset are shown in Appendix E, indicating that our method achieved much better performance than these baselines. Therefore, our method can effectively infer the GC in multi-agent (or multi-element) systems with partially known structures.

	Boid model			
	Bal. Acc.	AUPRC	BA$_{pos}$	BA$_{neg}$
Linear GC	0.487 ± 0.028	0.591 ± 0.169	0.55 ± 0.150	0.530 ± 0.165
Local TE	0.634 ± 0.130	0.580 ± 0.141	N/A	N/A
eSRU [30]	0.500 ± 0.000	0.452 ± 0.166	0.495 ± 0.102	0.508 ± 0.153
ACD [42]	0.411 ± 0.099	0.497 ± 0.199	N/A	N/A
GVAR [44]	0.441 ± 0.090	0.327 ± 0.119	0.524 ± 0.199	0.579 ± 0.126
ABM (ours)	$\mathbf{0.767 \pm 0.146}$	$\mathbf{0.819 \pm 0.126}$	$\mathbf{0.724 \pm 0.189}$	$\mathbf{0.760 \pm 0.160}$

Table 1: Performance comparison on the boid model.

6.2 Multi-animal trajectory datasets

We here analyzed biological multi-agent trajectory datasets of bats, birds, mice, and flies and obtained new biological insights using our framework (for the results of flies, see Appendix H). We used the same ABM as used in the boid dataset. In real-world data, since there was no ground truth, we used the hyperparameters of the boid simulation dataset. As a possible verification method, our method can be verified by investigating whether the GC result follows the hypothesis using mice and flies (Appendix H) datasets, which were controlled based on scientific knowledge. Next, we show that our methods as analytical tools can obtain new insights from birds and bats datasets based on the quantitative results. Compared with the methodologies mentioned in Section 5, our method has advantages for providing local interactions: interpretable signs of Granger-causal effects over time.

Mice for verification. As an application to a hypothesis-driven study, for example, we show the effectiveness of our method using three mice raised in different environments. The hypothesis is that, as is well known (e.g., [51]), when grown in different cages, they are more socially novel to others, thus more frequently attractive and repulsive movements will be observed. In this experiment, we regarded the same/different cage as a group pseudo-label, and confirmed that our method could extract its features in an unsupervised manner. We analyzed the trajectories of three mice in the same and different cages measured at 30 Hz for 5 min each (see also Appendix H). As shown in Figure 2A, our method extracted a significantly larger duration in the same cage than that in the different cages for both movements ($p < 0.05$; p is a statistical p-value), whereas GVAR [44] did not in repulsive movements ($p > 0.05$) but did in attractive movements ($p < 0.05$) and extracted too much interaction. The statistical analysis and videos are presented in Appendix H and the supplementary materials. Our methods characterized the moving behaviors before the contacts with others, which have been previously evaluated (e.g., [72]).

Growing birds. Animals grow while interacting with other individuals, but the directed interaction between young individuals has not been fully investigated as longitudinal (i.e., long period) studies. Here, as an example, we analyzed the flight GPS (two-dimensional) trajectories of three juvenile brown boobies *Sula leucogaster* over six times for 34 days (11.91 ± 0.09 [h] for each day), which were recorded at 1 Hz. We segmented two or three bird trajectories within 1 km and during moving (over 1 km/s) each other, in which interactions were considered to exist, and obtained 25 sequences of length 367 ± 278 [s] (for details, see Appendix H). Results of inferring GC in Figure 2B show that on the first measurement day, the most frequent directed interactions were observed between ID 1 and 2 (particularly two individuals had more repulsions and ID 1 had fewer attractions). On the other hand, in the second and subsequent measurements, it was observed that the most interacting

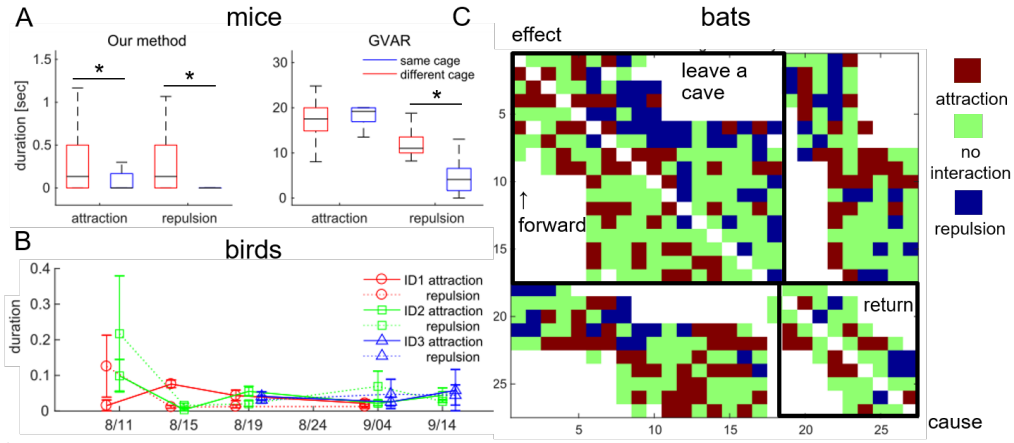

Figure 2: Analyzed results for multiple species of multi-animal trajectories. The results and details are given in the main text and Appendix H. (A) Results of three mice data grown in the different (red) and same (blue) cages. The vertical axis indicates the duration [sec] of their attraction and repulsion during 10-second bins of three interactions. Our method significantly extracted distinctive differences between cages in both movements. (B) Results of the longitudinal two or three birds data. The horizontal axis indicates the measurement date. The GPS trajectories of identified three young brown boobies (red, green, and blue) were analyzed (missing values indicates no measurement). The vertical axis indicates the normalized duration of positive (attraction: solid line) and negative (repulsion: negative) GCs for each bird (i.e., worked as the cause of another one or two birds). Error bar is the standard error among the segment during the movement. (C) Results of the observational 27 bats. The horizontal and vertical axes are the agents of the cause and effect in GC inferred by our method, respectively. The agents were sorted in the order they framed out by leaving and returning to the cave (the groups of leaving and returning were separated). The colors are the signed maximal values of the absolute GC coefficients inferred by our method, i..e, red, green, and blue indicate attraction (1), no interaction (0), and repulsion (-1), respectively.

individuals changed every measurement day. One possible factor of the decrease in the duration of interactions (especially repulsion) may be the habituation with the same individuals. Measurements and analyses over longer periods will reveal the acquisition of social behavior in young individuals.

Wild bats. As an example of an exploratory analysis, we applied our method to three-dimensional trajectories of eastern bent-wing bats *Miniopterus fuliginosus* that left a cave (some bats returned to the cave). Although some multi-animal studies have investigated leader-follower relationships (see also Section 5), those in wild bats are unknown. We used two sequences with 7 and 27 bats of length 237 and 296 frames, respectively, which were obtained via digitizing the videos at 30 Hz. Details of the dataset are given in Appendix H. As a result, among 138 interactions of all 34 individuals within the leaving and returning groups, there were 46 interactions where the locationally-leading (i.e., flying forward) bats repelled the following bats in the same direction, 27 interactions where the leading bats were attracted from the following bats, 65 ones with no interactions (the results of the following bats were not discussed because it was obvious; see also example results of 27 bats in Figure 2C). Since bats can echo-locate other bats in all directions up to a range of approximately 20 m [43, 6], the locationally-leading bats can be influenced by the locationally-following bats in the same direction (if no perception, they cannot be influenced). The results suggest that the groups of flying bats would not show simple leader-follower relationships.

7 Conclusions

We proposed a framework for learning GC from multi-animal trajectories via a theory-based ABM with interpretable neural models. Our method can analyze the biological moving sequence of mice, birds, and bats, and obtained novel biological insights. One possible future research direction is to incorporate other scientific knowledge into the models such as body inertia, which could contribute to further understanding of biological multi-agent behaviors.

For societal impact, our method can be utilized as real-world multi-agent analyses to estimate interaction rules such as in animals, pedestrians, vehicles, and athletes in sports. On the other hand, there are some concerns in our method from the perspectives of negative impact. One is ignoring a fairness problem, i.e., it may include bias e.g., when applied to human data. Solutions for such a problem will improve the applicability of the proposed method in our society.

References

[1] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 961–971, 2016.

[2] D. Alvarez-Melis and T. Jaakkola. Towards robust interpretability with self-explaining neural networks. In *Advances in Neural Information Processing Systems*, pages 7775–7784, 2018.

[3] M. O. Appiah. Investigating the multivariate Granger causality between energy consumption, economic growth and CO_2 emissions in Ghana. *Energy Policy*, 112:198–208, 2018.

[4] A. Attanasi, A. Cavagna, L. Del Castello, I. Giardina, T. S. Grigera, A. Jelić, S. Melillo, L. Parisi, O. Pohl, E. Shen, et al. Information transfer and behavioural inertia in starling flocks. *Nature Physics*, 10(9):691–696, 2014.

[5] M. Bansal, A. Krizhevsky, and A. Ogale. Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. *arXiv preprint arXiv:1812.03079*, 2018.

[6] T. Beleyur and H. R. Goerlitz. Modeling active sensing reveals echo detection even in large groups of bats. *Proceedings of the National Academy of Sciences*, 116(52):26662–26668, 2019.

[7] K. Branson, A. A. Robie, J. Bender, P. Perona, and M. H. Dickinson. High-throughput ethomics in large groups of drosophila. *Nature Methods*, 6(6):451–457, 2009.

[8] C. H. Brighton and G. K. Taylor. Hawks steer attacks using a guidance system tuned for close pursuit of erratically manoeuvring targets. *Nature communications*, 10(1):1–10, 2019.

[9] S. Butail, F. Ladu, D. Spinello, and M. Porfiri. Information flow in animal-robot interactions. *Entropy*, 16(3):1315–1330, 2014.

[10] S. Butail, V. Mwaffo, and M. Porfiri. Model-free information-theoretic approach to infer leadership in pairs of zebrafish. *Physical Review E*, 93(4):042411, 2016.

[11] O. M. Cliff, J. T. Lizier, X. R. Wang, P. Wang, O. Obst, and M. Prokopenko. Quantifying long-range interactions and coherent structure in multi-agent dynamics. *Artificial Life*, 23(1):34–57, 2017.

[12] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology*, 218(1):1–11, 2002.

[13] E. Crosato, L. Jiang, V. Lecheval, J. T. Lizier, X. R. Wang, P. Tichit, G. Theraulaz, and M. Prokopenko. Informative and misinformative interactions in a school of fish. *Swarm Intelligence*, 12(4):283–305, 2018.

[14] E. Demir and B. J. Dickson. fruitless splicing specifies male courtship behavior in drosophila. *Cell*, 121(5):785–794, 2005.

[15] R. Escobedo, V. Lecheval, V. Papaspyros, F. Bonnet, F. Mondada, C. Sire, and G. Theraulaz. A data-driven method for reconstructing and modelling social interactions in moving animal groups. *Philosophical Transactions of the Royal Society B*, 375(1807):20190380, 2020.

[16] E. Eyjolfsdottir, K. Branson, Y. Yue, and P. Perona. Learning recurrent representations for hierarchical behavior modeling. In *International Conference on Learning Representations*, 2017.

[17] K. Fujii and Y. Kawahara. Dynamic mode decomposition in vector-valued reproducing kernel hilbert spaces for extracting dynamical structure among observables. *Neural Networks*, 117:94–103, 2019.

[18] K. Fujii, N. Takeishi, M. Hojo, Y. Inaba, and Y. Kawahara. Physically-interpretable classification of network dynamics for complex collective motions. *Scientific Reports*, 10(3005), 2020.

[19] K. Fujii, N. Takeishi, Y. Kawahara, and K. Takeda. Policy learning with partial observation and mechanical constraints for multi-person modeling. *arXiv preprint arXiv:2007.03155*, 2020.

[20] T. Golany, K. Radinsky, and D. Freedman. Simgans: Simulator-based generative adversarial networks for ecg synthesis to improve deep ecg classification. In *International Conference on Machine Learning*, pages 3597–3606. PMLR, 2020.

[21] C. Graber and A. G. Schwing. Dynamic neural relational inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

[22] C. W. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, pages 424–438, 1969.

[23] L. Guangyu, J. Bo, Z. Hao, C. Zhengping, and L. Yan. Generative attention networks for multi-agent behavioral modeling. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2020.

[24] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi. Social gan: Socially acceptable trajectories with generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2255–2264, 2018.

[25] Y. Hoshen. Vain: Attentional multi-agent predictive modeling. In *Advances in Neural Information Processing Systems 30*, pages 2701–2711, 2017.

[26] A. Hyvärinen, K. Zhang, S. Shimizu, and P. O. Hoyer. Estimation of a structural vector autoregression model using non-gaussianity. *Journal of Machine Learning Research*, 11(5), 2010.

[27] M. J. Johnson, D. K. Duvenaud, A. Wiltschko, R. P. Adams, and S. R. Datta. Composing graphical models with neural networks for structured representations and fast inference. In *Advances in Neural Information Processing Systems 29*, pages 2946–2954, 2016.

[28] I. Karamouzas, B. Skinner, and S. J. Guy. Universal power law governing pedestrian interactions. *Physical review letters*, 113(23):238701, 2014.

[29] A. Karpatne, G. Atluri, J. H. Faghmous, M. Steinbach, A. Banerjee, A. Ganguly, S. Shekhar, N. Samatova, and V. Kumar. Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Transactions on Knowledge and Data Engineering*, 29(10):2318–2331, 2017.

[30] S. Khanna and V. Y. Tan. Economy statistical recurrent units for inferring nonlinear granger causality. In *International Conference on Learning Representations*, 2019.

[31] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

[32] T. Kipf, E. Fetaya, K.-C. Wang, M. Welling, and R. Zemel. Neural relational inference for interacting systems. In *International Conference on Machine Learning*, pages 2688–2697, 2018.

[33] M. Kolar, L. Song, A. Ahmed, E. P. Xing, et al. Estimating time-varying networks. *Annals of Applied Statistics*, 4(1):94–123, 2010.

[34] Y. Kuramoto. Self-entrainment of a population of coupled non-linear oscillators. In *International Symposium on Mathematical Problems in Theoretical Physics*, pages 420–422. Springer, 1975.

[35] F. Ladu, V. Mwaffo, J. Li, S. Macrì, and M. Porfiri. Acute caffeine administration affects zebrafish response to a robotic stimulus. *Behavioural Brain Research*, 289:48–54, 2015.

[36] H. M. Le, Y. Yue, P. Carr, and P. Lucey. Coordinated multi-agent imitation learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1995–2003, 2017.

[37] E. Leurent and J. Mercat. Social attention for autonomous decision-making in dense traffic. *arXiv preprint arXiv:1911.12250*, 2019.

[38] N. Lim, F. d'Alché Buc, C. Auliac, and G. Michailidis. Operator-valued kernel-based vector autoregressive models for network inference. *Machine Learning*, 99(3):489–513, 2015.

[39] J. T. Lizier, M. Prokopenko, and A. Y. Zomaya. Local information transfer as a spatiotemporal filter for complex systems. *Physical Review E*, 77(2):026110, 2008.

[40] J. T. Lizier, M. Prokopenko, and A. Y. Zomaya. Information modification and particle collisions in distributed computation. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 20(3):037109, 2010.

[41] W. M. Lord, J. Sun, N. T. Ouellette, and E. M. Bollt. Inference of causal information flow in collective animal behavior. *IEEE Transactions on Molecular, Biological and Multi-Scale Communications*, 2(1):107–116, 2016.

[42] S. Löwe, D. Madras, R. Zemel, and M. Welling. Amortized causal discovery: Learning to infer causal graphs from time-series data. *arXiv preprint arXiv:2006.10833*, 2020.

[43] Y. Maitani, K. Hase, K. I. Kobayasi, and S. Hiryu. Adaptive frequency shifts of echolocation sounds in miniopterus fuliginosus according to the frequency-modulated pattern of jamming sounds. *Journal of Experimental Biology*, 221(23), 2018.

[44] R. Marcinkevics and J. E. Vogt. Interpretable models for granger causality using self-explaining neural networks. In *1st NeurIPS workshop on Interpretable Inductive Biases and Physically Structured Learning (2020)*, 2020.

[45] D. Marinazzo, M. Pellicoro, and S. Stramaglia. Kernel-granger causality and the analysis of dynamical networks. *Physical review E*, 77(5):056215, 2008.

[46] F. Martinez-Gil, M. Lozano, I. García-Fernández, and F. Fernández. Modeling, evaluation, and scale on artificial pedestrians: a literature review. *ACM Computing Surveys (CSUR)*, 50(5):1–35, 2017.

[47] A. Mathis, P. Mamidanna, K. M. Cury, T. Abe, V. N. Murthy, M. W. Mathis, and M. Bethge. Deeplabcut: markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience*, 2018.

[48] J. M. McCracken. Exploratory causal analysis with time series data. *Synthesis Lectures on Data Mining and Knowledge Discovery*, 8(1):1–147, 2016.

[49] A. Montalto, S. Stramaglia, L. Faes, G. Tessitore, R. Prevete, and D. Marinazzo. Neural networks with non-uniform embedding and explicit validation phase to assess granger causality. *Neural Networks*, 71:159–171, 2015.

[50] H. Murakami, T. Niizato, and Y.-P. Gunji. Emergence of a coherent and cohesive swarm based on mutual anticipation. *Scientific reports*, 7:46447, 2017.

[51] J. Nadler, S. Moy, G. Dold, D. Trang, N. Simmons, A. Perez, N. Young, R. Barbaro, J. Piven, T. Magnuson, et al. Automated apparatus for quantitation of social approach behaviors in mice. *Genes, Brain and Behavior*, 3(5):303–314, 2004.

[52] M. Nagy, Z. Ákos, D. Biro, and T. Vicsek. Hierarchical group dynamics in pigeon flocks. *Nature*, 464(7290):890–893, 2010.

[53] R. Nathan, W. M. Getz, E. Revilla, M. Holyoak, R. Kadmon, D. Saltz, and P. E. Smouse. A movement ecology paradigm for unifying organismal movement research. *Proceedings of the National Academy of Sciences*, 105(49):19052–19059, 2008.

[54] M. Nauta, D. Bucur, and C. Seifert. Causal discovery with attention-based convolutional neural networks. *Machine Learning and Knowledge Extraction*, 1(1):312–340, 2019.

[55] W. B. Nicholson, D. S. Matteson, and J. Bien. Varx-l: Structured regularization for large vector autoregressions with exogenous variables. *International Journal of Forecasting*, 33(3):627–651, 2017.

[56] T. Niizato, K. Sakamoto, Y.-i. Mototake, H. Murakami, T. Tomaru, T. Hoshika, and T. Fukushima. Finding continuity and discontinuity in fish schools via integrated information theory. *PloS One*, 15(2):e0229573, 2020.

[57] N. Orange and N. Abaid. A transfer entropy analysis of leader-follower interactions in flying bats. *The European Physical Journal Special Topics*, 224(17):3279–3293, 2015.

[58] J. Peters, D. Janzing, and B. Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.

[59] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.

[60] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th annual Conference on Computer Graphics and Interactive Techniques*, pages 25–34, 1987.

[61] N. Rhinehart, R. McAllister, K. Kitani, and S. Levine. Precog: Prediction conditioned on goals in visual multi-agent settings. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2821–2830, 2019.

[62] T. O. Richardson, N. Perony, C. J. Tessone, C. A. Bousquet, M. B. Manser, and F. Schweitzer. Dynamical coupling during collective animal motion. *arXiv preprint arXiv:1311.1417*, 2013.

[63] A. Roebroeck, E. Formisano, and R. Goebel. Mapping directed influence over the brain using Granger causality and fMRI. *NeuroImage*, 25(1):230–242, 2005.

[64] D. Sankey, L. O'Bryan, S. Garnier, G. Cowlishaw, P. Hopkins, M. Holton, I. Fürtbauer, and A. King. Consensus of travel direction is achieved by simple copying, not voting, in free-ranging goats. *Royal Society Open Science*, 8(2):201128, 2021.

[65] T. Schreiber. Measuring information transfer. *Physical Review Letters*, 85(2):461, 2000.

[66] J. C. Simon and M. H. Dickinson. A new chamber for studying the behavior of drosophila. *Plos One*, 5(1):e8793, 2010.

[67] V. Sindhwani, H. Q. Minh, and A. C. Lozano. Scalable matrix-valued kernel learning for high-dimensional nonlinear multivariate regression and granger causality. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*, pages 586–595, 2013.

[68] L. Song, M. Kolar, and E. Xing. Time-varying dynamic bayesian networks. *Advances in Neural Information Processing Systems*, 22:1732–1740, 2009.

[69] P. A. Stokes and P. L. Purdon. A study of problems encountered in granger causality analysis from a neuroscience perspective. *Proceedings of the National Academy of Sciences*, 114(34):E7063–E7072, 2017.

[70] C. Tang and R. R. Salakhutdinov. Multiple futures prediction. In *Advances in Neural Information Processing Systems 32*, pages 15398–15408, 2019.

[71] A. Tank, I. Covert, N. Foti, A. Shojaie, and E. Fox. Neural granger causality for nonlinear time series. *arXiv preprint arXiv:1802.05842*, 2018.

[72] P. K. Thanos, C. Restif, J. R. O'Rourke, C. Y. Lam, and D. Metaxas. Mouse social interaction test (most): a quantitative computer automated analysis of behavior. *Journal of Neural Transmission*, 124(1):3–11, 2017.

[73] T. Tomaru, H. Murakami, T. Niizato, Y. Nishiyama, K. Sonoda, T. Moriyama, and Y.-P. Gunji. Information transfer in a swarm of soldier crabs. *Artificial Life and Robotics*, 21(2):177–180, 2016.

[74] K. Tsutsui, M. Shinya, and K. Kudo. Human navigational strategy for intercepting an erratically moving target in chase and escape interactions. *Journal of motor behavior*, 52(6):750–760, 2020.

[75] T. Vicsek and A. Zafeiris. Collective motion. *Physics Reports*, 517(3-4):71–140, 2012.

[76] X. R. Wang, J. M. Miller, J. T. Lizier, M. Prokopenko, and L. F. Rossi. Quantifying and tracing information cascades in swarms. *PloS One*, 7(7):e40084, 2012.

[77] Y. Wang, K. Lin, Y. Qi, Q. Lian, S. Feng, Z. Wu, and G. Pan. Estimating brain connectivity with varying-length time lags using a recurrent neural network. *IEEE Transactions on Biomedical Engineering*, 65(9):1953–1963, 2018.

[78] I. Winkler, D. Panknin, D. Bartz, K.-R. Müller, and S. Haufe. Validity of time reversal for testing granger causality. *IEEE Transactions on Signal Processing*, 64(11):2746–2760, 2016.

[79] T. Wu, T. Breuel, M. Skuhersky, and J. Kautz. Discovering nonlinear relations with minimum predictive information regularization. *arXiv preprint arXiv:2001.01885*, 2020.

[80] S. Zheng, Y. Yue, and J. Hobbs. Generating long-term trajectories using deep hierarchical networks. In *Advances in Neural Information Processing Systems 29*, pages 1543–1551, 2016.

[81] A. Zienkiewicz, D. Barton, M. Porfiri, and M. Di Bernardo. Leadership emergence in a data-driven model of zebrafish shoals with speed modulation. *The European Physical Journal Special Topics*, 224(17):3343–3360, 2015.

[82] A. K. Zienkiewicz, F. Ladu, D. A. Barton, M. Porfiri, and M. Di Bernardo. Data-driven modelling of social forces and collective behaviour in zebrafish. *Journal of Theoretical Biology*, 443:39–51, 2018.

[83] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [Yes]

(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]

2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]