



Causal Graph ODE: Continuous Treatment Effect Modeling in Multi-agent Dynamical Systems

Zijie Huang^{†*}

Computer Science Department
University of California, Los Angeles
Los Angeles, California, USA
zjehuang@cs.ucla.edu

Jinwoo Baik

Computer Science Department
University of California, Los Angeles
Los Angeles, California, USA
jbaik23@g.ucla.edu

Yizhou Sun

Computer Science Department
University of California, Los Angeles
Los Angeles, California, USA
yzsun@cs.ucla.edu

Jeehyun Hwang^{*}

Computer Science Department
University of California, Los Angeles
Los Angeles, California, USA
hwanggh96@gmail.com

Weitong Zhang

Computer Science Department
University of California, Los Angeles
Los Angeles, California, USA
weightzero@g.ucla.edu

Quanquan Gu

Computer Science Department
University of California, Los Angeles
Los Angeles, California, USA
qgu@cs.ucla.edu

Junkai Zhang^{*}

Computer Science Department
University of California, Los Angeles
Los Angeles, California, USA
jkzhang@g.ucla.edu

Dominik Wodarz

Department of Ecology, Behavior & Evolution
University of California, San Diego
San Diego, California, USA
dwodarz@ucsd.edu

Wei Wang

Computer Science Department
University of California, Los Angeles
Los Angeles, California, USA
weiwang@cs.ucla.edu

ABSTRACT

Real-world multi-agent systems are often dynamic and continuous, where the agents co-evolve and undergo changes in their trajectories and interactions over time. For example, the COVID-19 transmission in the U.S. can be viewed as a multi-agent system, where states act as agents and daily population movements between them are interactions. Estimating the counterfactual outcomes in such systems enables accurate future predictions and effective decision-making, such as formulating COVID-19 policies. However, existing methods fail to model the continuous dynamic effects of treatments on the outcome, especially when multiple treatments (e.g., "stay-at-home" and "get-vaccine" policies) are applied simultaneously. To tackle this challenge, we propose Causal Graph Ordinary Differential Equations (CAG-ODE), a novel model that captures the continuous interaction among agents using a Graph Neural Network (GNN) as the ODE function. The key innovation of our model is to learn time-dependent representations of treatments and incorporate them into the ODE function, enabling precise predictions of potential outcomes. To mitigate confounding bias, we further propose two domain adversarial learning-based objectives, which enable our model to learn balanced continuous representations that are not affected by treatments or interference. Experiments on

two datasets (i.e., COVID-19 and tumor growth) demonstrate the superior performance of our proposed model.¹

CCS CONCEPTS

- Computing methodologies → *Supervised learning by regression; Causal reasoning and diagnostics; Temporal reasoning; Probabilistic reasoning; Spatial and physical reasoning*.

KEYWORDS

Dynamical System, Neural ODE, Causal Inference, Graph Neural Networks

ACM Reference Format:

Zijie Huang[†], Jeehyun Hwang, Junkai Zhang, Jinwoo Baik, Weitong Zhang, Dominik Wodarz, Yizhou Sun, Quanquan Gu, and Wei Wang. 2024. Causal Graph ODE: Continuous Treatment Effect Modeling in Multi-agent Dynamical Systems. In *Proceedings of the ACM Web Conference 2024 (WWW '24)*, May 13–17, 2024, Singapore, Singapore. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3589334.3648148>

1 INTRODUCTION

Many real-world multi-agent systems are dynamic and continuous, where agents (nodes) interact and exhibit complex behaviors over time. This results in time-evolving node trajectories and dynamic interaction edges. An example is the spread of COVID-19 in the U.S., where states act as agents and daily migration patterns across states form interaction edges [14, 29]. Estimating the counterfactual outcomes over time in such systems are crucial for various applications, such as formulating effective policies and designing medical treatment plans [2, 3, 40]. This can achieve more accurate predictions than non-causal methods by considering the influence of

^{*}Equally contributed. [†] Corresponding author.



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

¹Our code implementation can be found at <https://github.com/Jun-Kai-Zhang/CAG-ODE.git>.

biased confounders. Confounders are variables that have influences on treatments and outcomes. For example, the health status of the residents in each state (confounders) can impact their level of adherence to the state's policies (treatments), which can influence future confirmed cases/deaths (outcomes). Non-causal methods only learn the statistical associations between treatments and outcomes from observational data, which can have non-uniform treatment distributions across confounder values, potentially leading to incorrect predictions such as taking vaccines can increase the number of confirmed cases for each state. Furthermore, causal inference for multi-agent dynamical systems enables effective decision-making by addressing causal questions such as "What if we remove a policy at a specific time" or "What if we change the order of different policies". Therefore, it serves as a promising tool for policymakers.

Traditionally, the standard approach for causal inference over time is randomized controlled trials (RCTs) [5], which can be very costly to obtain and can raise some ethical problems [3, 40]. Thus, researchers have turned to using observational data and employed methods like linear regression [35], recurrent neural networks (RNNs) [4, 25], and Transformers [31] to estimate counterfactual outcomes with time dependencies. However, causal inference for multi-agent dynamical systems presents unique challenges.

One is that most existing methods [3, 40] assume that nodes are independent, meaning their trajectories are determined solely by their own treatments. Some [18] considers the influence of neighboring nodes but only assumes static interactions among them, which fails to capture situations such as daily population travel patterns between states in the context of COVID-19.

In causal terms, influences of neighboring nodes can be categorized into two parts: 1.) time-dependent neighborhood confounding, where a node's treatment and outcome may be confounded by the covariates of its neighbors. For example, if cases in neighboring states rise (covariate), a state may implement a vaccine policy (treatment) that affects future confirmed cases/deaths (outcome). 2.) time-dependent interference, where the outcome of a node can be influenced by the treatments of its neighbors. For example, a state may have reduced future cases/deaths (outcome) if neighboring states have implemented a vaccine policy (covariates), as higher vaccination rates within the population flow network give stronger protection. As the interaction edges evolve along with node trajectories, the challenges lie in predicting the neighbors of each node (edges) and then addressing the time-dependent neighborhood confounding and interference issues.

Another challenge is that current methods lack the ability to capture the continuous and dynamic effects of multiple treatments on such systems. For instance, the impact of a "stay-at-home" policy may be most significant during its initial implementation, and when a "get-vaccine" policy is subsequently introduced, the combined effect of these policies can result in a different outcome. Existing studies often focus on a single treatment [18, 40] or simply append fixed multi-hot treatment representations when a node receives them. These fixed treatment representations fail to differentiate the influences of the same treatment administered at different times.

To tackle these challenges, we propose a novel causal inference framework: the **Causal Graph Ordinary Differential Equations** (CAG-ODE) to estimate the continuous counterfactual outcome of

a multi-agent dynamical system in the presence of multiple treatments and time-varying confounding and interference. Building upon the recent success of graph ordinary differential equations (ODE) in capturing the continuous interaction among agents [13, 14, 17, 27], our key innovation is to learn time-dependent representations of simultaneous treatments and incorporate them into the ODE function to accurately account for their causal effects on the system. As nodes and edges are jointly evolving, we utilize two coupled treatment-induced ODE functions to account for their respective dynamics. To mitigate confounding bias, we further design two adversarial learning losses, which enable our model to learn balanced continuous trajectory representations unaffected by treatments or interference. Experiments on both real and simulated datasets demonstrate the effectiveness of our proposed model. The primary contributions of this paper can be summarized as follows:

- We propose CAG-ODE to estimate continuous counterfactual outcomes in multi-agent systems with evolving interaction edges and multiple treatments.
- CAG-ODE features a novel treatment fusing module that can capture the dynamic effects of treatment over time and the combined effect of multiple treatments.
- Our method achieves the state-of-art results in counterfactual estimation across varying systems, and can serve as a promising tool for policymakers.

2 PRELIMINARIES AND RELATED WORK

2.1 Graph Neural Networks (GNNs)

Graph Neural Networks (GNNs) are a class of neural networks that operate on graph-structured data by passing local messages [24, 42, 43]. They have been extensively employed in various applications such as node classification, link prediction, and recommendation systems [12, 16]. GNNs have shown to be efficient for approximating pair-wise node interactions and achieved accurate predictions for multi-agent dynamical systems [22, 38]. The majority of existing studies propose discrete GNN-based simulators where they take the node features at time t as input to predict the node features at time $t+1$. To further capture the long-term temporal dependency for predicting future trajectories, some work utilizes recurrent neural networks such as RNN, LSTM, or self-attention mechanism to make predictions at time $t+1$ based on the historical trajectory sequence [10, 11, 39]. However, they restrict themselves to learning a one-step state transition function. Therefore, when we successively apply these one-step simulators to previous predictions in order to generate the rollout trajectories, error accumulates and impairs the prediction accuracy, especially for long-range prediction.

2.2 Graph Ordinary Differential Equations for Continuous Multi-agent Dynamical Systems

The dynamics of a multi-agent system can be captured by a series of nonlinear first-order ordinary differential equations (ODEs) [13, 15, 28, 36], which describe how the states of N dependent variables co-evolve over continuous time: $\dot{z}_i^t := \frac{dz_i^t}{dt} = g(z_1^t, z_2^t \dots z_N^t)$. Here $z_i^t \in \mathbb{R}^d$ denotes the state variable for agent i at timestamp t and g denotes the ODE function that drives the system to move forward. Given the initial states $z_1^0, \dots z_N^0$ for all agents and the ODE function

g , a numerical ODE solver such as Runge-Kutta [32] can be used to evaluate \mathbf{z}_i^T at any desired time T using Eqn (1):

$$\mathbf{z}_i^T = \mathbf{z}_i^0 + \int_{t=0}^T g(\mathbf{z}_1^t, \mathbf{z}_2^t, \dots, \mathbf{z}_N^t) dt. \quad (1)$$

To model the interactions among agents, recent studies [13, 14, 34, 44] propose using a GNN as the ODE function g which is learned from observational data. Such GraphODE framework follows an encoder-processor-decoder architecture. The encoder computes latent initial states for all agents based on historical observations. The GNN-based ODE function then predicts the latent trajectories starting from the learned initial states. Finally, a decoder extracts the predicted dynamic features. To regularize the generated trajectories, GraphODE frameworks often adopt a variational autoencoder (VAE) structure [21], where the encoder samples initial states from approximated posterior distributions. GraphODEs are promising in making long-range predictions and can handle irregularly-sampled observations effectively [13, 44].

2.3 Causal Inference Over Time

Time-dependent causal inference methods mainly differ in how they deal with confounding. They differ from traditional statistical time series analysis [1, 23, 45] which we do not discuss in this paper. Traditionally, many statistical tools that are applied, such as marginal structural models (MSMs) [35] utilize the inverse probability of treatment weighting (IPTW). Recently, representation learning-based balancing approaches are proposed, which learn representations that are not predictable of the treatments to ensure unbiased outcome prediction [4, 31]. However, one major limitation is that they are discrete methods, which can offer poor performance on continuous systems such as the spread of COVID-19. There are a series of works [3, 7, 9, 40] that estimate the continuous counterfactual outcomes through neural ODEs or neural controlled differential equations (CDEs). Despite their success, they assume that nodes are independent of each other, regardless of their interactions. One recent work [18] proposed to parameterize the ODE function with a GNN for multi-agent settings. However, this model cannot handle evolving graph structures and the effect of multiple treatments.

3 PROBLEM DEFINITION

We consider a dynamical system of N agents as an evolving interaction graph $\mathcal{G}^t = \{\mathcal{V}, \mathcal{E}^t\}$, where nodes $\mathcal{V} = \{v_1, v_2, \dots, v_N\}$ are agents and \mathcal{E}^t are the weighted edges among them, denoting agents' dynamic interaction that changes over time. Each node is associated with time-varying causal characteristics, which we introduce in the following along with the causal inference framework.

We follow the longitudinal causal inference setting for predicting future potential outcomes as in [37]. We denote the observational data at timestamp t as $(\mathbf{X}^t, \mathbf{W}^t, \mathbf{A}^t, \mathbf{Y}^t)$, where $\mathbf{X}^t \in \mathbb{R}^{N \times d_1}$ represents the time-varying covariates (e.g., the health status of residents) of N agents. $\mathbf{W}^t \in \mathbb{R}^{N \times N}$ represents the weighted adjacency matrix, whose element $w_{i \rightarrow j} \in \mathbb{R}$ is the weight of the directed edge that points from node i to node j and may be asymmetric. $\mathbf{A}^t \in \{0, 1\}^{N \times K}$ are time-dependent treatments, where $\mathbf{A}_{kj}^t = 1$ denotes the k^{th} treatment assigned to node i at timestamp t , and K

is the number of heterogeneous treatments. $\mathbf{Y}^t \in \mathbb{R}^{N \times d_2}$ is the time-dependent outcome, such as the number of confirmed cases in each state, which can be part of \mathbf{X}^t . The historical observations up to time t is represented as $\mathcal{H}^t = \{\bar{\mathbf{X}}^t, \bar{\mathbf{W}}^t, \bar{\mathbf{A}}^t, \bar{\mathbf{Y}}^t\}$, where $\bar{\mathbf{X}}^t, \bar{\mathbf{W}}^t, \bar{\mathbf{A}}^t, \bar{\mathbf{Y}}^t$ contain all $\mathbf{X}^{t^-}, \mathbf{W}^{t^-}, \mathbf{A}^{t^-}, \mathbf{Y}^{t^-}$ ($t^- \leq t$). We aim to predict the unbiased potential outcomes $\mathbb{E}(\mathbf{Y}^{t^+}(\mathbf{A}^{t^+} = a) | \mathcal{H}^t)$ under any treatment assignment a^2 . Here, a is the dynamic treatment trajectory (e.g. sequences of state policies). As only one of the potential outcome trajectories is observed for each treatment assignment, we refer to the unobserved potential outcomes as counterfactuals [4, 40].

To make potential outcomes identifiable from observational data, we follow three standard assumptions [4, 18, 40] below:

Assumption 1: Consistency. The potential outcome is equal to the observed factual outcome if $\mathbf{A}^t = a^t: \mathbf{Y}^{t^+}(\mathbf{A}^t = a^t) = \mathbf{Y}^{t^+}$.

Assumption 2: Overlap. At any time point t^+ , there is some positive probability of treatment assignment regardless of the historical observation: $0 < P(\mathbf{A}^{t^+} = a | \mathcal{H}^t) < 1, \forall \mathcal{H}^t, t < t^+$.

The last assumption defines unconfoundedness (strong ignorability) in dynamical systems. We first define the interference effects caused by neighbors' treatments of node i as $\mathbf{G}_i^t = \sum_{j \in \mathcal{N}_i} \frac{1}{|\mathcal{N}_i|} \mathbf{A}_j^t \in \mathbb{R}^K$, which is the proportion of treated nodes in node i 's neighbors for each treatment type. We refer to \mathbf{G}_i^t as interference summary, which assumes that a node is only influenced by treatments of its immediate neighbors as in previous studies [18, 19, 30].

Assumption 3: Strong Ignorability for Multi-Agent Dynamical Systems. Given the historical observations, the potential outcome trajectory is independent of the treatments and interference summary: $\mathbf{Y}^{t^+}(\mathbf{A}^t = a) \perp \mathbf{A}^{t^+}, \mathbf{G}^{t^+} | \mathcal{H}^t, \forall a, t$.

It ensures that it is sufficient to only condition on the historical observations and graph sequences up to t to block all backdoor paths so as to estimate the potential outcome in the future. With these three assumptions, the potential outcome trajectory can be identified as:

$$\mathbb{E}(\mathbf{Y}^{t^+}(\mathbf{A}^t = a) | \mathcal{H}^t) = \mathbb{E}(\mathbf{Y}^{t^+} | \mathbf{A}^{t^+}, \mathbf{G}^{t^+}, \mathcal{H}^t).$$

This enables us to estimate the potential outcomes by training a machine learning model using observational data, and to use the same model to predict counterfactual outcomes given new treatment trajectories.

4 THE PROPOSED MODEL: CAG-ODE

In this section, we present Causal Graph ODE (CAG-ODE) to predict continuous counterfactual outcomes for multi-agent dynamical systems with evolving interaction edges and dynamic multi-treatment effects. Following the framework of GraphODEs [13, 14, 18, 34, 44], CAG-ODE adopts the encoder-ODE generative model-decoder architecture described in Sec. 2.2 to capture the continuous interaction among agents. As nodes and edges are jointly evolving, we utilize two coupled ODE functions [14] for the evolution of nodes and edges respectively. Contrary to GraphODEs, CAG-ODE can perform causal reasoning by injecting treatment effects into the ODE functions, which we call *treatment-induced coupled graph ODE*. The multi-treatment effects are captured by a novel treatment fusing module that assigns temporal weights to the treatments using an

²The potential outcome can also be formalized using *do* operation [33]

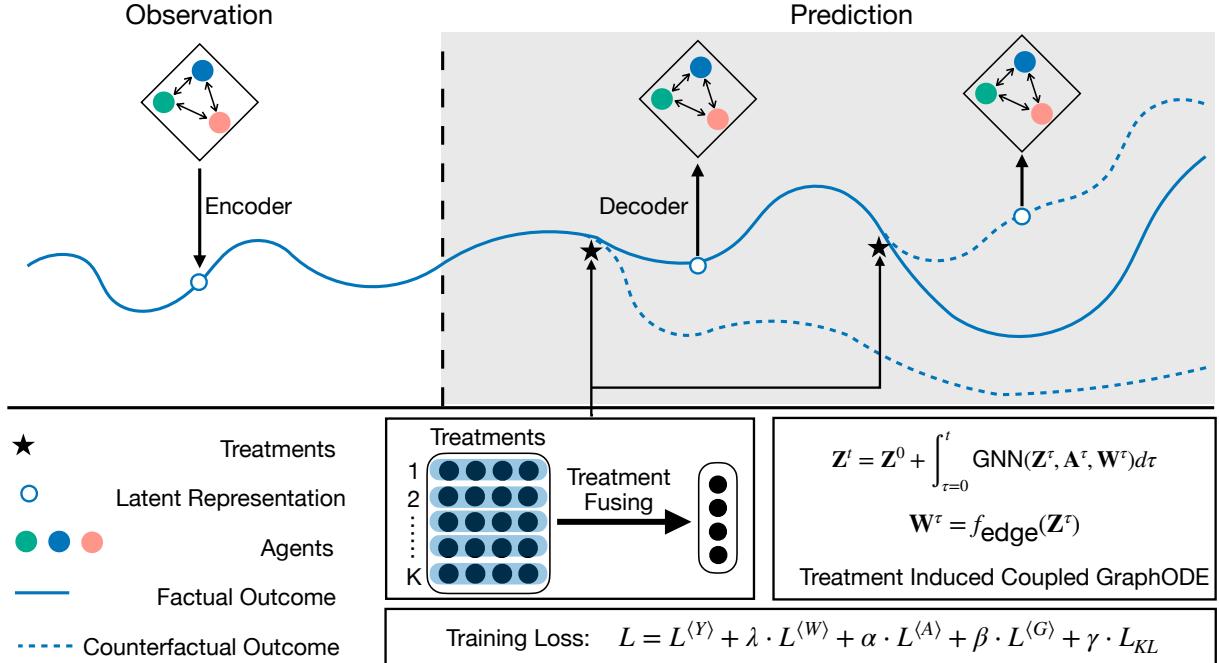


Figure 1: Overall Framework of CAG-ODE. The encoder first computes the latent initial states. Then the treatment-induced coupled ODE functions predict the continuous trajectories over time. Treatment representations learned through the fusing module are incorporated into the ODE functions to enable counterfactual prediction. Finally, the decoder outputs the predicted dynamics. Treatment and interference balancing losses are designed to ensure unbiased counterfactual predictions.

attention mechanism. As time-dependent confounders can result in a biased distribution of treatment assignments and imbalanced interferences due to the evolving graph structure, CAG-ODE utilizes two adversarial learning losses to ensure unbiased estimations of counterfactual outcomes. The overall framework is depicted in Figure 1. We now discuss each module in detail.

4.1 Spatial-Temporal Initial State Encoder

The encoder of CAG-ODE infers the posterior distributions from the historical observations and then samples the latent initial states from them. It follows the architecture described in [14]. As the evolution of different nodes is mutually influenced, we calculate the initial states for all nodes simultaneously considering their interactions over time. The initial states of edges are derived from the initial states of nodes.

Dynamic Node Representation Learning. We construct a graph to represent the spatial-temporal structure of multi-agent dynamical systems, with each node corresponding to an agent's observation at a particular timestamp. There are two types of edges: spatial edges at the same timestamp and temporal edges across different timestamps. The spatial edges are formed according to the adjacency matrices, denoted as $w_{i(t) \rightarrow j(t)}$. For the temporal edges, we only consider edges from an agent's own previous observations to later observations, denoted as $w_{i(t) \rightarrow i(t')}$, where $t' = t + 1$.

The latent representations of observations are learned from this spatial-temporal graph through an attention mechanism approach.

The propagation among L GNN layers is depicted in Equation(2).

$$\begin{aligned}
 \mathbf{h}_{i(t')}^l &= \mathbf{h}_{i(t')}^l + \sigma \left(\sum_{j(t) \in \mathcal{N}_{i(t')}} e_{j(t) \rightarrow i(t')}^l \times \mathbf{W}_o \hat{\mathbf{h}}_{j(t)}^{l-1} \right), \\
 e_{j(t) \rightarrow i(t')}^l &= \mathbf{w}_{j(t) \rightarrow i(t')} \times \alpha_{j(t) \rightarrow i(t')}^l, \\
 \alpha_{j(t) \rightarrow i(t')}^l &= \left(\mathbf{W}_k \hat{\mathbf{h}}_{j(t)}^{l-1} \right)^T \left(\mathbf{W}_q \mathbf{h}_{i(t')}^{l-1} \right) \cdot \frac{1}{\sqrt{d}}, \\
 \hat{\mathbf{h}}_{j(t)}^{l-1} &= \mathbf{h}_{j(t)}^{l-1} + \text{TE}(t - t'), \\
 \text{TE}(\Delta t)_{2i} &= \sin \left(\frac{\Delta t}{10000^{2i/d}} \right), \quad \text{TE}(\Delta t)_{2i+1} = \cos \left(\frac{\Delta t}{10000^{2i/d}} \right).
 \end{aligned} \tag{2}$$

Here, $\mathbf{h}_{i(t)}^l$ represents the agent i at time t from layer l . The attention score $e_{j(t) \rightarrow i(t')}^l$ is defined as the product of edge weights $\mathbf{w}_{j(t) \rightarrow i(t')}$ and affinity score $\alpha_{j(t) \rightarrow i(t')}^l$, which is computed using the representations of the sender and receiver nodes. Additionally, we incorporate temporal embedding, denoted as TE, into the sender node's representation to establish temporal distinction. Then, the final representation is obtained from the L layer as $\mathbf{h}_{i(t)} = \mathbf{h}_{i(t)}^L$.

Sequence Representation Learning. Then, we employ self-attention to compute the sequence representation of observed temporal information for each node, where $\hat{\mathbf{h}}_{i(t)} = \mathbf{h}_{i(t)} + \text{TE}(t)$.

$$\mathbf{u}_i = \frac{1}{N} \sum_{t=1}^T (\mathbf{a}_i^T \hat{\mathbf{h}}_{i(t)} \hat{\mathbf{h}}_{i(t)}), \quad \mathbf{a}_i = \tanh \left(\left(\frac{1}{N} \sum_{t=1}^T \hat{\mathbf{h}}_{i(t)} \right) \mathbf{W}_a \right). \tag{3}$$

Finally, the mean and variance of the posterior distribution is obtained through a neural network f_{dist} from the sequence representation \mathbf{u}_i .

$$\mathbf{z}_i^0 \sim q_\phi(\mathbf{z}_i^0 | \mathcal{H}^0) = \mathcal{N}(\boldsymbol{\mu}_{\mathbf{z}_i^0}, \boldsymbol{\sigma}_{\mathbf{z}_i^0}^2), \boldsymbol{\mu}_{\mathbf{z}_i^0}, \boldsymbol{\sigma}_{\mathbf{z}_i^0} = f_{\text{dist}}(\mathbf{u}_i).$$

Next, the latent initial state for an edge is given by $\mathbf{z}_{i \rightarrow j}^0 = f_{\text{edge}}([\mathbf{z}_i^0, \mathbf{z}_j^0])$, where f_{edge} is parameterized by a neural network and $[\cdot, \cdot]$ is concatenation operation.

4.2 Treatment Fusing

To conduct causal inference with CAG-ODE, we propose to inject the dynamic effects of multiple treatments into the ODE function. Treatments can have time-varying effects in multi-agent dynamical systems and they can occur simultaneously, resulting in a combined effect. To model such complex behaviors, we propose a novel treatment fusing module that assigns temporal weights to multiple treatments through an attention mechanism. The temporal weight of treatment at timestamp t is dependent on both the start time of each treatment and the occurrence of other treatments as shown in Eqn (4). Let $\mathbf{e}_k \in \mathbb{R}^K$ be the one-hot representation of treatment k . We first add it with the temporal encoding TE[14, 41] to account for the time elapsed since the start of the treatment t' . Here $\mathbf{A}_{ik}^t \in \{0, 1\}$ is an indicator showing whether treatment k would be applied to agent i at timestamp t . Therefore the computed treatment representation $\hat{\mathbf{o}}_{ik}^t$ becomes zero when $\mathbf{A}_{ik}^t = 0$, to ensure computational efficiency. A contraction matrix \mathbf{W}_q is then used to transform this sparse representation into a more compact form.

$$\begin{aligned} \hat{\mathbf{o}}_{ik}^t &= \mathbf{A}_{ik}^t \mathbf{e}_k + \text{TE}(t - t') \mathbb{1}[\mathbf{A}_{ik}^t = 1], \quad \mathbf{o}_{ik}^t = \mathbf{W}_q \hat{\mathbf{o}}_{ik}^t, \\ \text{TE}(\Delta t)_{2i} &= \sin\left(\Delta t/M^{2i/d}\right), \\ \text{TE}(\Delta t)_{2i+1} &= \cos\left(\Delta t/M^{2i/d}\right), M = 10000. \end{aligned} \quad (4)$$

To account for the combined effect of simultaneous treatments, we compute the combined treatment representation as a weighted sum of all in-effect treatments at timestamp t (Eqn 5). We first compute an attention vector \mathbf{m}_i^t as the tanh-transformed average of all the treatment representations, $\hat{\mathbf{o}}_{ij}^t$. Each treatment's weight is derived from the dot product of its representation and \mathbf{m}_i^t , thereby integrating each treatment's influence into \mathbf{o}_i^t .

$$\mathbf{o}_i^t = \frac{1}{K} \sum_k \left(\mathbf{m}_i^{t \top} \hat{\mathbf{o}}_{ik}^t \hat{\mathbf{o}}_{ik}^t \right), \mathbf{m}_i^t = \tanh\left(\left(\frac{1}{K} \sum_k \mathbf{o}_{ik}^t\right) \mathbf{W}_m\right). \quad (5)$$

The fusing operation has a time complexity of $O(K)$ if having K treatments and therefore is able to scale up to larger systems.

4.3 Treatment-Induced GraphODE

We use two coupled ODEs to predict the latent trajectories for nodes and edges respectively, accounting for their co-evolution [14]. We incorporate the learned treatment representations into the ODEs to enable counterfactual predictions in the future. Specifically, the co-evolution of nodes and edges is depicted in Eqn 6. The co-evolution depends on all historical information implicitly as \mathbf{Z}_t embeds the trajectories up to time t . $\tilde{\mathbf{W}}_A^t = \mathbf{D}^{-1} \mathbf{W}_A^t$ is the normalized adjacency matrix and \mathbf{D} is the diagonal degree matrix defined as

$\mathbf{D}_{ii} = \sum_j \mathbf{W}_{Aij}^t \cdot f_{\text{e}}, f_{\text{self}}, f_{\text{edge2value}}$ are all implemented as Multi-Layer Perceptrons (MLPs). To incorporate the treatment effect into the function, we use a linear transformation \mathbf{W} to merge the latent states of nodes \mathbf{Z}^t and the treatment representation \mathbf{O}^t . In this way, the latent trajectories of agents are affected not only by their own past trajectories and treatments but also by the trajectories and treatments of their interacting agents.

$$\begin{aligned} \frac{d\mathbf{Z}^t}{dt} &= \sigma\left(\tilde{\mathbf{W}}_A^t \mathbf{W}[\mathbf{Z}^t, \mathbf{O}^t]\right) - \mathbf{Z}^t + \mathbf{Z}^0, \\ \frac{dz_{i \rightarrow j}^t}{dt} &= f_e\left([\mathbf{z}_i^t, \mathbf{z}_j^t]\right) + f_{\text{self}}\left(\mathbf{z}_{i \rightarrow j}^t\right), \\ \mathbf{W}_{Aij}^t &= f_{\text{edge2value}}\left(\mathbf{z}_{i \rightarrow j}^t\right), \quad \tilde{\mathbf{W}}_A^t = \mathbf{D}^{-1} \mathbf{W}_A^t. \end{aligned} \quad (6)$$

4.4 Outcome Prediction

Given the treatment representations, the ODE functions, the latent initial states for nodes and edges, and the latent trajectories for all agents can be determined using any black-box ODE solver. Finally, we compute the predicted trajectories for each agent and their interactions based on the decoding likelihoods in Eqn (7), where f_{decN} and f_{decE} are node and edge decoding functions respectively. They output the means of the normal distributions $p(\mathbf{y}_i^t | \mathbf{z}_i^t)$ and $p(\mathbf{w}_{i \rightarrow j}^t | \mathbf{z}_i^t)$, which we treat as the predicted values from our model.

$$\mathbf{y}_i^t \sim p(\mathbf{y}_i^t | \mathbf{z}_i^t) = f_{\text{decN}}(\mathbf{z}_i^t), \mathbf{w}_{i \rightarrow j}^t \sim p(\mathbf{w}_{i \rightarrow j}^t | \mathbf{z}_i^t) = f_{\text{decE}}(\mathbf{z}_i^t). \quad (7)$$

We implemented all of our decoders using two-layer fully connected neural networks. The node feature decoder's input dimension matches the latent state dimension d , while the output dimension is one, reflecting our outcome of interest. The edge decoder's input dimension is $2d$ and the output dimension is 1. The treatment decoder also has an input dimension equal to the latent state's dimension d . However, its output dimension matches the number of distinct treatments, predicting the probability of each treatment being chosen. Lastly, the interference decoder's input dimension is the sum of the latent state dimension and the treatment embedding dimension, i.e. $2d$. Its output dimension mirrors the number of treatment options. For all decoders, the latent hidden dimension is half of their respective input dimensions.

We calculate the reconstruction loss of model predictions for nodes $\hat{\mathbf{Y}}_i^t$ and edges $\hat{\mathbf{w}}_{i \rightarrow j}^t$ as:

$$L^{(Y)} = \frac{1}{N} \frac{1}{T} \sum_t \|\mathbf{Y}^t - \hat{\mathbf{Y}}^t\|_2^2, \quad L^{(W)} = \frac{1}{N^2} \frac{1}{T} \sum_t \|\mathbf{W}_A^t - \hat{\mathbf{W}}_A^t\|_F^2.$$

4.5 Domain Adversarial Learning

In observational data, treatment assignments are not randomized but are biased based on time-varying confounder values. This can lead to increased variance and bias in counterfactual estimation [40]. In multi-agent dynamical systems, unbalanced interference from neighboring agents further exacerbates this effect and alters the state of each agent. To obtain an unbiased counterfactual prediction, we need to ensure that the distribution of latent representation trajectories is invariant to treatments and interference [18]. This guarantees that the treatments cannot be inferred from the latent trajectory representations and that the interference is not

predictable when the treatment is combined with the latent representation.

To achieve this, we incorporate two adversarial learning losses into the optimization objective function and use gradient reversal layers for the implementation.

Treatment Balancing The treatment combinations \hat{A}^t can be predicted using a decoder from the latent state z_i^t . Formally, $\hat{A}_{ik}^t = \Phi_A(r(z_i^t))$, where Φ_A is a neural network attempting to recover treatments from the latent state z_i^t , and the gradient reversal layer, denoted by r , reverses the sign of gradient during back-propagation. The treatment balancing can be expressed as the maximization of the following loss term through the construction of min-max games:

$$L^{(A)} = -\frac{1}{N} \frac{1}{T} \frac{1}{K} \sum_{i=1}^N \sum_{t=1}^T \sum_{k=1}^K \sum_{j \in \{0,1\}} \mathbb{1}[(A_{ik}^t = j)] \log(\Phi_A^{j,k}(r(z_i^t))),$$

where $\Phi_A^{j,k}$ represents the logits of $d_A(\cdot)$ for predicting j on k -th treatment. Note that we achieve treatment balancing by letting the latent representations z_i^t not be predictable for each individual treatment. This is because the representation of multiple treatments is essentially a linear combination of individual treatments. If each individual treatment is not predictable based on z_i^t , then it is also impossible to use such representation to predict when multiple treatments occur together.

Interference Balancing Similar to treatment balancing, the interference prediction can be represented as $\hat{G}_i^t = \Phi_G(r([Z_i^t, A_i^t]))$, where d_G denotes a neural network designed to estimate interference. As interference is a continuous variable, we employ continuous domain adversarial training to accomplish interference balancing. By incorporating a gradient reversal layer, interference balancing can be achieved by minimizing the following loss term:

$$L^{(G)} = \frac{1}{N} \frac{1}{T} \frac{1}{K} \sum_{i=1}^N \sum_{t=1}^T \|\Phi_G(r([z_i^t, o_i^t])) - G_i^t\|_2^2 d.$$

Overall Loss The overall training objective is defined as the weighted summation of node reconstruction loss, edge reconstruction loss, treatment balancing loss, and interference balancing loss. Since we follow the VAE framework, we also incorporate a KL divergence loss to add regularization towards the sampled initial states, which is defined as: $L_{KL} = KL \left[\prod_{i=1}^N q_\phi(z_i^0 | \mathcal{H}^0) \parallel p(Z^0) \right]$. Therefore, the overall training loss is formalized as:

$$L = L^{(Y)} + \lambda L^{(W)} + \alpha L^{(A)} + \beta L^{(G)} + \gamma L_{KL}.$$

5 EXPERIMENTS

5.1 Experiment Setup

5.1.1 Datasets and Experiment Configuration. We evaluate the performance of our model using two datasets: 1.) The **COVID-19 dataset**, which captures the daily COVID-19 trends of U.S. states from April.12.2020 to Dec.31.2020. The daily population flows among states are represented as dynamic edges. Treatments are state-level COVID-19 policies. We ask the model to predict the daily confirmed cases in each state. 2.) The **Tumor Growth simulation dataset** [8], which describes the tumor growth dynamics in different regions of patients, where they may receive differing

treatments. We aim to predict the tumor volumes in each region. Additional details about the datasets can be found in Appendix A.

We predict trajectory rollouts across varying lengths and use Root Mean Square Error (RMSE) as the evaluation metric. Specifically, we train our model in a sequence-to-sequence setting where we split the trajectory of each training sample into two parts $[t_1, t_K]$ and $[t_{K+1}, t_T]$. We condition the model on the first part of observations and predict the second part. To fully utilize the data points within each trajectory, we generate training and validation samples by splitting each trajectory into several chunks using a sliding window. Details can be found in Appendix B.

5.1.2 Baselines and Model Variants. We conduct a comparative analysis of our model with three baseline models: one non-causal continuous multi-agent baseline CG-ODE [14], and two causal models: TE-CDE [40] and COVID-POLICY [29]. TE-CDE [40] is a causal model that employs continuous-time differential equations to capture temporal event dependencies. COVID-Policy [29] is another causal model designed specifically for assessing the impact of public health policies on COVID-19 outcomes. To further analyze the performance of our model, we also compare variants of our model. Each variant excludes a specific component to assess its individual impact on performance. The variants include models without treatment balancing, interference balancing, both components or the attention module.

5.1.3 Training Details. We employ the AdamW optimizer, as proposed in the study by Loshchilov et al. [26], to train our model. The initial learning rate is set at $\eta = 0.005$, and the batch size is set as 8 to accommodate memory constraints.

The Graph Neural Network (GNN) used for the encoder has a singular layer with a hidden dimension of 64. Similarly, the GNN that parameterizes the ODE function is also comprised of a single layer. The dimension of the latent state is set at 20, and the dimension for the embedded treatments is 5. We assign a weight of 10 for both the treatment balancing term α and the interference balancing term β . Additionally, the weight designated for the edge reconstruction error λ is set at 0.5.

5.2 Performance Evaluation

We evaluate the performance of our model, CAG-ODE, as well as the baselines using Root Mean Square Error (RMSE) across different prediction lengths. The results are shown in Table 1 and Table 2, reporting the factual and counterfactual outcomes respectively. As the COVID-19 is a real-world dataset that does not have counterfactual outcomes, we evaluate only the Tumor Growth dataset in Table 2. To ensure consistent comparison, we align the prediction periods of all models with weekly intervals on the COVID-19 dataset, similar to the statistical baselines derived from their official weekly submissions to the CDC, as done in [14]. To assess the accuracy of short-term and long-term predictions, the prediction lengths for the COVID-19 and Tumor Growth datasets are set to 7, 14, 21 days and 14, 21, and 28 days, respectively. We include longer-range predictions on the Tumor-Growth dataset in Appendix C

Factual Outcome Predictions. Table 1 shows that our model, CAG-ODE, consistently outperforms the baseline models across all prediction lengths for both datasets. This underscores the effectiveness of our model in capturing the dynamic interactions

Table 1: Root Mean Square Error (RMSE) for factual outcome evaluation across prediction lengths (the duration for which predictions are made). For the COVID-19 dataset, we report the mean and standard deviation accuracy with multiple runs.

Dataset Prediction Length	Covid-19			Tumor Growth		
	7-days	14-days	21-days	14-days	21-days	28-days
CG-ODE	4063 \pm 68	4454 \pm 100	4659 \pm 63	18.37	21.00	24.58
TE-CDE	7999 \pm 212	7470 \pm 289	6832 \pm 243	55.45	55.38	71.23
COVID-POLICY	4008 \pm 44	4128 \pm 60	3963 \pm 59	20.07	25.93	29.29
CAG-ODE	3710 \pm 29	3925 \pm 44	3933 \pm 40	10.91	10.82	14.84
w/o $L^{(G)}$	3800 \pm 60	3987 \pm 40	3990 \pm 49	15.57	16.28	16.62
w/o $L^{(A)}$	3840 \pm 35	4100 \pm 53	4069 \pm 49	17.90	14.69	20.19
w/o $L^{(G)}, L^{(A)}$	3793 \pm 23	4089 \pm 79	3953 \pm 38	17.28	16.72	24.36
w/o attention	3867 \pm 61	3958 \pm 31	4256 \pm 55	18.91	17.55	34.45

among objects, especially over longer time periods. Comparing our model with TE-CDE, we observe a performance gap that highlights the benefits of incorporating interference balancing and spatial correlation in the model. Additionally, our model outperforms the COVID-POLICY model, indicating its broader generalizability across different types of data due to modeling dynamic interactions. Furthermore, our model exhibits proficiency in both short-term and long-term predictions. For instance, it achieves promising results for 21-day predictions on the COVID-19 dataset and 28-day predictions on the Tumor Growth simulation dataset. The analysis of our model variants further emphasizes the importance of each component in the model. Particularly, the model variant excluding the attention module has the weakest performance, indicating the significance of our time-embedding attention module in effectively representing the treatment.

Counterfactual Outcome Predictions. In the context of a multi-agent dynamical system, the total number of possible treatments for all nodes is $O(K \times 2N)$, making it infeasible to enumerate all treatment combinations. To assess the robustness of each model to counterfactual treatment scenarios, we perform an experiment where we randomly flip a certain percentage of observed treatments. In Table 2, we evaluate the performance when 25%, 50%, and 75% of all observed treatments in each experiment are randomly flipped. The purpose of this experiment is to examine the robustness of the models to counterfactual treatment scenarios, and since CG-ODE does not incorporate causal modeling, it is excluded from this experiment. CAG-ODE outperforms others by a wide margin across all settings. These findings collectively demonstrate the superiority of our proposed model, CAG-ODE, in capturing the dynamics of multi-agent systems and making accurate predictions across different time horizons. We additionally include the visualization of the learned balanced latent representations in Section 5.4.

5.3 Case Study about COVID-19 Policies

We conduct a case study to show the impact of different treatments, e.g., COVID-19 related policies, on the COVID-19 dataset as shown in Figure 2. Specifically, we consider four different policy intervention methods and report the resulting average changes in the number of daily confirmed cases across all states in the U.S.

First, we focus on the removal of policies in three states that have the highest number of announced policies during the time frame of the COVID-19 dataset. By masking out these policies,

we observe an increase in the average number of confirmed cases across states in the future. This increase is attributed to both in-state disease spread and population flow to other states. The removal of policies exacerbates the spread of COVID-19 over an extended period, as shown in Figure 2(a), indicating that our model captures the dynamic interference resulting from agents' interactions.

We then explore the effect of changing the starting time of a specific policy for all states. We changed the "No Public Gatherings" policy starting time for each state to be 15 days earlier, 15 and 30 days later respectively. As shown in Figure 2(b) when announcing the policy earlier, we observe a decrease in the average number of daily confirmed cases in the future, while announcing the policy later leads to an increase. This intuitive outcome highlights the capability of our model to capture the causal relationships between policy interventions and COVID-19 spread.

Next, we analyze the impact of the top three most frequent policies across all states by removing them separately. As shown in Figure 2(c), the "Public Gatherings" policy has the largest effect in reducing the spread of COVID-19, even though the most frequent policy is "Emergency Funds". This demonstrates the potential of our model in assisting policymakers to identify the relative importance of each policy over time.

Finally, we study the effects of different orders in policy announcements, specifically focusing on the simultaneous or closely timed announcements of "No Public Gatherings" and "No Traveler from Outside States" policies. We change the announcement dates for the two policies in each state to mimic three scenarios shown in Figure 2 (d). We found that initializing the announcement of "No Public Gatherings" early generally contributes to a reduction in the spread of COVID-19 compared with "No Traveler from Outside States". We further analyzed the daily population flow during the given time frame and found that the majority of population flows are within the same states, indicating that residents of each state pose a high risk of virus transmission compared to people from other states. These insights suggest prioritizing the earlier announcement of the "No Public Gatherings" policy over the "No Traveler from Outside States" policy can better mitigate the spread of COVID-19.

These case study results demonstrate the effectiveness of our model CAG-ODE in capturing the complex interactions between treatments, disease spread, and population flow, providing valuable insights for policymakers in making informed decisions.

Table 2: Root Mean Square Error (RMSE) for counterfactual Outcome evaluation on the Tumor Growth dataset with treatment flipping ratio. Treatment F.R. (Treatment Flipping Ratio) represents the ratio of treatments that are flipped.

Prediction Length Treatment F.R.	14-days			21-days			28-days		
	0.25	0.5	0.75	0.25	0.5	0.75	0.25	0.5	0.75
TE-CDE	95.61	103.2	100.8	98.65	103.0	97.93	118.3	124.0	121.4
COVID-POLICY	21.32	22.37	23.31	26.63	26.83	27.00	32.01	32.16	32.21
CAG-ODE	17.23	16.98	16.96	18.64	18.84	18.85	19.91	19.88	19.87
w/o $L^{(G)}$	20.62	20.53	20.51	19.70	19.60	19.55	21.10	21.41	21.38
w/o $L^{(A)}$	22.17	22.35	22.35	20.19	20.10	20.09	20.83	21.14	21.15
w/o $L^{(G)}, L^{(A)}$	19.78	19.75	19.71	19.34	19.29	19.27	21.31	21.40	21.34
w/o attention	19.09	18.37	18.13	22.16	21.78	21.65	27.70	27.44	27.38

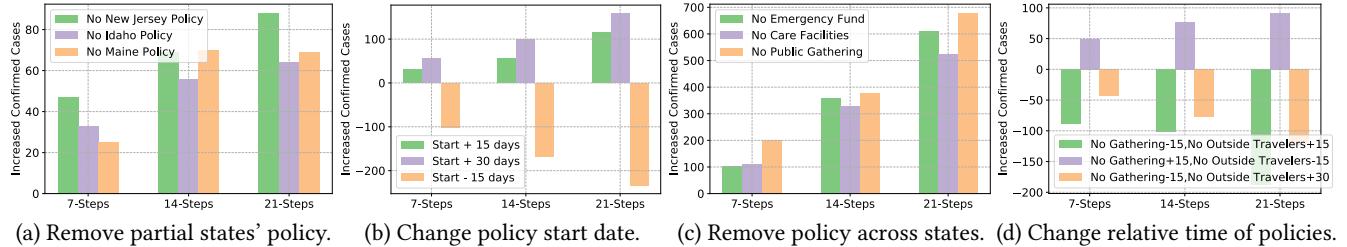


Figure 2: Case Study for changing different policies on the COVID-19 dataset.

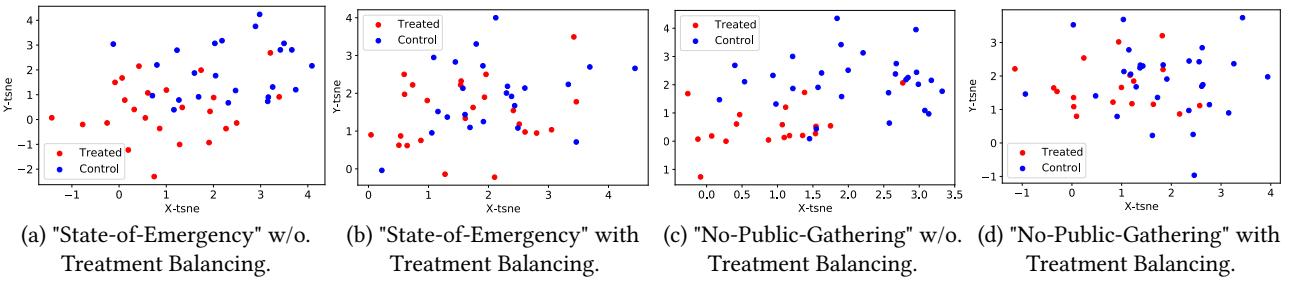


Figure 3: Treatment Balancing Visualization on the COVID-19 Dataset.

5.4 Visualization of Learned Balanced Representations

To further understand the effect of treatment balancing loss in CAG-ODE, we visualize the 2-D T-SNE projections of the latent representations of nodes on the COVID-19 dataset, i.e. z_i^t as shown in Figure 3. Specifically, we visualize the latent node representations under two different treatments: "State-of-Emergency" and "No-Public-Gathering". Under each treatment (policy), we use different colors to denote whether a node receives such treatment (treated) or not (control). As shown in Figure 3(a) and (c), the distributions of the learned representations are more distinguishable between the two groups, compared with Figure 3(b) and (d) which have the treatment balancing loss. This indicates that CAG-ODE indeed learns balanced latent representations by employing the treatment balancing loss.

6 CONCLUSION

In this paper, we introduce the causal graph ODE (CAG-ODE) as a model for estimating continuous counterfactual outcomes in multi-agent-dynamical systems with evolving interaction edges and dynamic multi-treatment effects. Our model builds upon existing

GraphODEs and incorporates causal reasoning for multi-agent dynamical systems. We propose a novel treatment fusing module that captures the dynamic effects of multiple treatments occurring simultaneously. Through extensive experiments on both the real-world and the simulated datasets, we demonstrate the superior performance of our model across various prediction settings, validating its effectiveness. Furthermore, we leverage our model to analyze policy effects analysis on the COVID-19 dataset, providing valuable insights for policymakers.

7 ACKNOWLEDGEMENT

This work was partially supported by NSF 2106859, 2200274, 2312501, NEC and NSF 2211557, NSF 1937599, NSF 2119643, NSF 2303037, NSF 20232551, NASA, SRC JUMP 2.0 Center, Cisco research grant, Picsart Gifts, and Snapchat Gifts. We would like to thank Song Jiang for his valuable discussion throughout this project.

REFERENCES

- [1] Guangji Bai, Chen Ling, and Liang Zhao. 2022. Temporal Domain Generalization with Drift-Aware Dynamic Neural Networks. *arXiv preprint arXiv:2205.10664* (2022).
- [2] Alexis Bellot and Mihaela Van Der Schaar. 2021. Policy analysis using synthetic controls in continuous-time. In *International Conference on Machine Learning*. PMLR, 759–768.
- [3] Ioana Bica, Ahmed M Alaa, James Jordon, and Mihaela van der Schaar. 2020. Estimating counterfactual treatment outcomes over time through adversarially balanced representations. *International Conference on Learning Representations* (2020).
- [4] Ioana Bica, Ahmed M. Alaa, James Jordon, and Mihaela van der Schaar. 2020. Estimating counterfactual treatment outcomes over time through adversarially balanced representations. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.
- [5] Thomas C Chalmers, Harry Smith Jr, Bradley Blackburn, Bernard Silverman, Biruta Schroeder, Dinah Reitman, and Alexander Ambroz. 1981. A method for assessing the quality of a randomized control trial. *Controlled clinical trials* 2, 1 (1981), 31–49.
- [6] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. 2018. Neural Ordinary Differential Equations. In *Advances in Neural Information Processing Systems* 31. 6571–6583.
- [7] Edward De Brouwer, Javier Gonzalez, and Stephanie Hyland. 2022. Predicting the impact of treatments over time with uncertainty aware neural differential equations.. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 4705–4722.
- [8] Changran Geng, Harald Paganetti, and Clemens Grassberger. 2017. Prediction of treatment response for combined chemo-and radiation therapy for non-small cell lung cancer patients using a bio-mathematical model. *Scientific reports* 7, 1 (2017), 13542.
- [9] Daehoon Gwak, Gyuhyeon Sim, Michael Poli, Stefano Massaroli, Jaegul Choo, and Edward Choi. 2020. Neural ordinary differential equations for intervention modeling. *arXiv preprint arXiv:2010.08304* (2020).
- [10] Ehsan Hajiramezanal, Arman Hasanzadeh, Krishna Narayanan, Nick Duffield, Mingyuan Zhou, and Xiaoning Qian. 2019. Variational Graph Recurrent Neural Networks. In *Advances in Neural Information Processing Systems* 32. 10701–10711.
- [11] Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. 2020. Heterogeneous Graph Transformer. In *Proceedings of the 2020 World Wide Web Conference*.
- [12] Zijie Huang, Zheng Li, Haoming Jiang, Tianyu Cao, Hanqing Lu, Bing Yin, Karthik Subbian, Yizhou Sun, and Wei Wang. 2022. Multilingual Knowledge Graph Completion with Self-Supervised Adaptive Graph Alignment. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [13] Zijie Huang, Yizhou Sun, and Wei Wang. 2020. Learning Continuous System Dynamics from Irregularly-Sampled Partial Observations. In *Advances in Neural Information Processing Systems*.
- [14] Zijie Huang, Yizhou Sun, and Wei Wang. 2021. Coupled graph ode for learning interacting system dynamics. In *The 27th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (SIGKDD)*.
- [15] Zijie Huang, Yizhou Sun, and Wei Wang. 2023. Generalizing graph ode for learning complex system dynamics across environments. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 798–809.
- [16] Zijie Huang, Daheng Wang, Binxuan Huang, Chenwei Zhang, Jingbo Shang, Yan Liang, Zhengyang Wang, Xian Li, Christos Faloutsos, Yizhou Sun, and Wei Wang. 2023. Concept2Box: Joint Geometric Embeddings for Learning Two-View Knowledge Graphs. In *Findings of the Association for Computational Linguistics: ACL 2023*. Association for Computational Linguistics, 10105–10118.
- [17] Zijie Huang, Wanjia Zhao, Jingdong Gao, Ziniu Hu, Xiao Luo, Yadi Cao, Yizhou Sun, and Wei Wang. 2024. TANGO: Time-Reversal Latent GraphODE for Multi-Agent Dynamical Systems. <https://openreview.net/forum?id=H7R0z6V9fR>
- [18] Song Jiang, Zijie Huang, Xiao Luo, and Yizhou Sun. 2023. CF-GODE: Continuous-Time Causal Inference for Multi-Agent Dynamical Systems. In *29th SIGKDD Conference on Knowledge Discovery and Data Mining*.
- [19] Song Jiang and Yizhou Sun. 2022. Estimating Causal Effects on Networked Observational Data via Representation Learning. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 852–861.
- [20] Yuhao Kang, Song Gao, Yunlei Liang, Mingxiao Li, and Jake Kruse. 2020. Multi-scale Dynamic Human Mobility Flow Dataset in the U.S. during the COVID-19 Epidemic. *Scientific Data* 1–13 (2020). Issue 390.
- [21] Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes.. In *ICLR*.
- [22] Thomas Kipf, Ethan Fetaya, Kuan-Chieh Wang, Max Welling, and Richard Zemel. 2018. Neural Relational Inference for Interacting Systems. *arXiv preprint arXiv:1802.04687* (2018).
- [23] Thomas Kipf, Ethan Fetaya, Kuan-Chieh Wang, Max Welling, and Richard Zemel. 2018. Neural Relational Inference for Interacting Systems. *arXiv preprint arXiv:1802.04687* (2018).
- [24] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR'17*.
- [25] Bryan Lim. 2018. Forecasting treatment responses over time using recurrent marginal structural networks. *advances in neural information processing systems* 31 (2018).
- [26] Ilya Loshchilov and Frank Hutter. [n. d.]. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.
- [27] Xiao Luo, Haixin Wang, Zijie Huang, Huiyu Jiang, Abhijeet Sadashiv Gangar, Song Jiang, and Yizhou Sun. 2023. CARE: Modeling Interacting Dynamics Under Temporal Environmental Variation. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- [28] Xiao Luo, Jingyang Yuan, Zijie Huang, Huiyu Jiang, Yifang Qin, Wei Ju, Ming Zhang, and Yizhou Sun. 2023. HOPE: high-order graph ODE for modeling interacting dynamics. In *Proceedings of the 40th International Conference on Machine Learning (ICML'23)*. 16 pages.
- [29] Jing Ma, Yushun Dong, Zheng Huang, Daniel Miettchen, and Jundong Li. 2022. Assessing the Causal Impact of COVID-19 Related Policies on Outbreak Dynamics: A Case Study in the US. In *Proceedings of the ACM Web Conference 2022 (WWW '22)*. 2678–2686.
- [30] Jing Ma, Mengting Wan, Longqi Yang, Jundong Li, Brent Hecht, and Jaime Teevan. 2022. Learning causal effects on hypergraphs. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1202–1212.
- [31] Valentyn Melnychuk, Dennis Frauen, and Stefan Feuerriegel. 2022. Causal Transformer for Estimating Counterfactual Outcomes. In *International Conference on Machine Learning, ICML 2022, 17–23 July 2022, Baltimore, Maryland, USA (Proceedings of Machine Learning Research, Vol. 162)*. PMLR, 15293–15329.
- [32] Schöber Michael, Särkkä Simo, and Philipp Hennig. 2019. A probabilistic model for the numerical solution of initial value problems. In *Statistics and Computing*. 99–122.
- [33] Judea Pearl. 2009. *Causality*. Cambridge university press.
- [34] Michael Poli, Stefano Massaroli, Junyoung Park, Atsushi Yamashita, Hajime Asama, and Jinkyoo Park. 2019. Graph neural ordinary differential equations. *arXiv preprint arXiv:1911.07532* (2019).
- [35] James M Robins, Miguel Angel Hernan, and Babette Brumback. 2000. Marginal structural models and causal inference in epidemiology. , 550–560 pages.
- [36] Yulia Rubanova, Ricky T. Q. Chen, and David K Duvenaud. 2019. Latent Ordinary Differential Equations for Irregularly-Sampled Time Series. In *Advances in Neural Information Processing Systems* 32. 5320–5330.
- [37] Donald B Rubin. 1978. Bayesian inference for causal effects: The role of randomization. *The Annals of statistics* (1978), 34–58.
- [38] Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter W. Battaglia. 2020. Learning to Simulate Complex Physics with Graph Networks. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13–18 July 2020, Virtual Event*. 8459–8468.
- [39] Aravind Sankar, Yanhong Wu, Liang Gou, Wei Zhang, and Hao Yang. 2020. DySAT: Deep Neural Representation Learning on Dynamic Graphs via Self-Attention Networks. In *WSDM'20*.
- [40] Nabeel Seedat, Fergus Imrie, Alexis Bellot, Zhaozhi Qian, and Mihaela van der Schaar. 2022. Continuous-Time Modeling of Counterfactual Outcomes Using Neural Controlled Differential Equations. In *International Conference on Machine Learning*. PMLR, 19497–19521.
- [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems* 30. 5998–6008.
- [42] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. *ICLR'18* (2018).
- [43] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How Powerful are Graph Neural Networks?. In *ICLR'19*.
- [44] Chengxi Zang and Fei Wang. 2020. Neural dynamics on complex networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 892–902.
- [45] Yunhao Zhang and Junchi Yan. 2023. Crossformer: Transformer Utilizing Cross-Dimension Dependency for Multivariate Time Series Forecasting. In *International Conference on Learning Representations*.

A APPENDIX

A DATASET DESCRIPTION

A.1 COVID-19 Dataset

Our experiments used the dataset provided by the Johns Hopkins Coronavirus Resource Center (JHU)³ from April 12th to December 31st, 2020. That is, we consider 264 time points, with each point representing one day. The dataset contains a comprehensive range of information, but for our experiments, we focus on up to 7 specific features. These features include the daily counts of confirmed cases, deaths, recovered cases, active cases, incident rate (cases per 100,000 people), mortality rate (calculated as the number of recorded deaths multiplied by 100 divided by the number of cases), and testing rate (total test results per 100,000 people). It's worth noting that while the JHU dataset provides cumulative data for confirmed, deaths, recovered, and active cases, our experiments and models specifically consider the daily increases in these features (e.g., the number of new cases reported each day).

To capture dynamic interaction edges, we use a temporal mobility flow network among a selection of 47 states based on COVID-19 USFlows [20].

The treatments are represented as statewide policies that aim to combat the spread of COVID-19. We identify 58 different statewide policies enacted throughout 2020, from the data given by the Department of Health & Human Services⁴. Each state enacted around 20 of these policies during the time period of April 2020 to December 2020, where the dataset provides the start and end dates of each enacted policy. In our model, treatments are encoded such that for each time point, the value is either 1 or 0 depending on whether the particular policy is enacted (for a given state) at that time or not, respectively.

Overall, the model receives input data for a total of 264 time points, covering each of the 47 states, and includes 7 features. Alongside this, the model is also provided with treatments and a mobility graph. Prior to being used as input for our models, the data is normalized. However, when calculating the test loss for comparison with other baseline models, the output is unnormalized. Our goal is to predict either the number of confirmed cases or the number of deaths for a future period of 7, 14, or 21 days.

A.2 Tumor Growth Dataset

We extend the state-of-the-art pharmacokinetic-pharmacodynamic (PK-PD) model of tumor growth proposed by [8] to simulate a more complex scenario where multiple tumor regions within a single patient interact with each other. The original model characterizes patients suffering from non-small cell lung cancer and models the evolution of their tumor under the combined effects of chemotherapy and radiotherapy. For a detailed description of the original model, we refer the readers to the original paper [8]. In our extended model, we incorporate two new terms: an interference term and a neighborhood covariate term. The volume of the tumor in

³<https://coronavirus.jhu.edu/about/how-to-use-our-data>

⁴<https://catalog.data.gov/dataset/covid-19-state-and-county-policy-orders-9408a>

region i at $t + 1$ days after diagnosis is modeled as follows:

$$\begin{aligned}
 V_i(t+1) = & \underbrace{\left(1 + \rho \log \left(\frac{K}{V_i(t)}\right)\right)}_{\text{Tumor Growth}} - \underbrace{\beta_c C_i(t)}_{\text{Chemotherapy}} - \underbrace{\left(\alpha_r d_i(t) + \beta_r d_i(t)^2\right)}_{\text{Radiotherapy}} + \underbrace{e_{it}}_{\text{Noise}} \Big) V_i(t) \\
 & + \underbrace{\left(\frac{1}{N_i} \sum_{j \in N_i} \iota_c C_j(t)\right)}_{\text{Chemotherapy Interference}} + \underbrace{\left(\frac{1}{N_i} \sum_{j \in N_i} (\iota_r d_j(t) + \iota_r d_j(t)^2)\right)}_{\text{Radiotherapy Interference}} \\
 & + \underbrace{\left(\frac{1}{N_i} \sum_{j \in N_i} \kappa V_j(t)\right)}_{\text{Neighborhood Covariates}}
 \end{aligned}$$

where the parameters $K, \rho, \beta_c, \alpha_r, \beta_r$ are sampled from the prior distributions described in [8], and $e_{it} \sim N(0, 0.012)$ is a noise term that accounts for randomness in the tumor growth. The prior means for β_c and α_r are adjusted to create three patient subgroups $S(i) \in \{1, 2, 3\}$ as described in [3]. The chemotherapy drug concentration follows an exponential decay with a half-life of 1 day. The time-varying confounding is introduced by modeling chemotherapy and radiotherapy assignment as Bernoulli random variables, with probabilities p_c and p_r depending on the tumor diameter. For more details, we refer the reader to the paper [3]. For our newly defined interference and neighborhood covariate terms, we set the hyperparameters ι_c and ι_r to 0.01, and κ to 0.001. These values were carefully chosen to reflect the strength of the interference and neighborhood covariates in the dataset. The number of tumors in each patient N_i is fixed to 15, and for each tumor region, the number of edges connected between the tumor regions is defined randomly from the range of 22 to 45. For additional experiments shown in Appendix C, the number of tumor regions for each patient is fixed to 5, and the number of edges ranges from 6 to 10. The dataset is input into our model similar to the COVID-19 dataset. Both chemotherapy and radiotherapy are encoded into 0 or 1 value depending on whether it was applied at a specific time point. The input data consists of 60-time points with 4 features, which include tumor volume, patient type, and the two treatments. The data is normalized for model input but unnormalized for test loss calculation. The model's objective is to predict tumor volume for future periods of 14, 21, or 28 days. We also create a longer-range dataset with 120-time points, which is described in C.

B DATA SPLITTING

We train our model in a sequence-to-sequence setting, where we split the trajectory of each training sample into two parts $[t_1, t_K]$ and $[t_{K+1}, t_T]$. We condition the model on the first part of observations and predict the second part. To fully utilize the data points within each trajectory, we generate training and validation samples by splitting each trajectory into several chunks using a sliding window with three hyperparameters: the observation length and prediction length for each sample, and the interval between two consecutive chunks (samples). We summarize the procedure in Algorithm 1, where K is the number of trajectories and d is the input feature dimension. For both datasets, we set the observation length to be 7 and the interval to be 3. We ask the model to make predictions at varying lengths for evaluation.

Prediction Length Treatment F.R.	35-days			49-days			63-days		
	0.25	0.5	0.75	0.25	0.5	0.75	0.25	0.5	0.75
CAG-ODE	35.00	34.68	34.50	34.12	37.92	37.88	46.71	47.28	48.12

Table 3: Root Mean Square Error (RMSE) for counterfactual outcome evaluation on the longer-range Tumor Growth dataset with treatment flipping ratio: 25%, 50%, 75%.

Algorithm 1: Data Splitting Procedure.

Input: Original Training trajectories $X_{\text{input}} \in \mathbb{R}^{K \times N \times T \times d}$;
 Observation length O ; Prediction length M ; Interval I ;
 Trajectory length T .
Output: Training samples after splitting X_{train} .

```

1 sample_length = O + M;
2 num_chunk = (T - sample_length) // interval + 1;
3 for i in range (0,K) do
4   for j in range(0,num_chunk,I) do
5     Generate the split training sample as
6     Xinput[i, :, j : j + sample_length, :]
7     Add the training sample to the training set Xtrain.
8 end

```

C LONGER-RANGE PREDICTION FOR THE TUMOR GROWTH DATASET

In Table 3, we evaluate the performance of our model. An extended version of the simulation dataset with a range of 120 days was used in the experiment. The considered prediction lengths are 35, 49, and 63 days.

As anticipated, the prediction errors exhibit a moderate increase with longer prediction lengths, as the added duration poses a greater challenge for accurate predictions. Note that the prediction error remains relatively stable when the treatment flipping ratio is increased from 25% to 75%. This observation suggests that the utilization of treatment balancing and interference balancing techniques

effectively mitigates the risk of overfitting to confounding factors, ensuring CAG-ODE's robustness.

D MODEL IMPLEMENTATION DETAILS

We use the fourth-order Runge-Kutta method from the `torchdiffeq` python package [6] as the ODE solver, for solving the ODE systems on a time grid that is five times denser than the observed time points. We also utilize the Adjoint method described in [6] to reduce memory use.

E LIMITATIONS

One limitation of CAG-ODE is that when inferring the future trajectories of nodes, we simply assume that all nodes are connected and jointly infer such edge evolution. This would bring huge computational costs when generalized to large-scale dynamical systems. In the future, we will consider more efficient sampling methods to accelerate the edge inference procedure to scale up our model. Another line of future work would be how to model more complex multiple treatment effects, including competing, hierarchical relationships.

F BROADER IMPACTS

Our work significantly enhances the performance of causal inference over multi-agent dynamical systems, which can potentially benefit a wide range of fields including public health, biology, physics, and robotics. Our work also advances the recent study of continuous graphODE for modeling multi-agent system dynamics, providing an efficient tool for further research on AI for science.