

Learning Structural Causal Models from Ordering: Identifiable Flow Models

Minh Khoa Le¹, Kien Do¹, Truyen Tran¹

¹ Applied Artificial Intelligence Institute (A2I2), Deakin University, Australia
{minh.le, k.do, truyen.tran}@deakin.edu.au

Abstract

In this study, we address causal inference when only observational data and a valid causal ordering from the causal graph are available. We introduce a set of flow models that can recover component-wise, invertible transformation of exogenous variables. Our flow-based methods offer flexible model design while maintaining causal consistency regardless of the number of discretization steps. We propose design improvements that enable simultaneous learning of all causal mechanisms and reduce abduction and prediction complexity to linear $O(n)$ relative to the number of layers, independent of the number of causal variables. Empirically, we demonstrate that our method outperforms previous state-of-the-art approaches and delivers consistent performance across a wide range of structural causal models in answering observational, interventional, and counterfactual questions. Additionally, our method achieves a significant reduction in computational time compared to existing diffusion-based techniques, making it practical for large structural causal models.

Introduction

Deep neural networks are highly expressive and learnable, but are inherently associative, making it difficult for them to capture causal relationships. This limitation can lead to inaccurate predictions in fields where causality is crucial (LeRoy 2004; Russo and Williamson 2007; Nguyen et al. 2023). Among the efforts to alleviate this problem, a promising direction dubbed causal representation learning (Schölkopf et al. 2021) is to integrate neural networks within the framework of Structural Causal Models (SCMs) (Pearl 2009). SCMs are principled way to answer observational, interventional, and counterfactual questions, but learning them from data remains challenging. This paper focuses on efficient learning of SCMs from only observational data and causal ordering, leveraging deep neural networks to model causal relationships in complex systems.

Some previous methods such as (Sánchez-Martin, Rateike, and Valera 2022) require a fully observed causal graph, which can be infeasible in real-world settings. Others (Javaloy, Martin, and Valera 2023; Khemakhem et al. 2021) use Autoregressive Normalizing Flows (Papamakarios, Pavlakou, and Murray 2017; Durkan et al. 2019), which

restrict model design to be monotonic and require additional regularization to scale to multiple layers. Sanchez and Tsafaris (2022) propose a counterfactual estimation method using diffusion models and classifier guidance, but it only considers bivariate causal graphs and lacks theoretical analysis. (Chao et al. 2024) generalize previous diffusion-based methods but require a fully observed causal graph and a separate deep neural network for each observed variable, resulting in slow sequential inference and a large number of parameters.

We propose a identifiable flow models that requires only observational data and a valid causal ordering. Our approach is designed to represent SCMs and ensure causal consistency through its structure. A parallel sped-up design can answer observational, interventional, and counterfactual questions with computational complexity scaling linearly with the number of model layers, independent of the causal graph’s node count. This scalability enables efficient handling of large causal models, making our approach practical and effective for diverse applications.

We empirically demonstrate that our method outperforms competing approaches across a wide range of synthetic and real datasets, excelling in estimating both the mean and overall shape of interventional and counterfactual distributions. The experiments confirm that our parallel architecture is not only scalable but also maintains high performance as the complexity and size of the datasets increase.

Our key contributions are:

1. We prove the identifiability of flow models for learning SCMs from observational data and causal ordering.
2. We introduce novel model designs enabling parallel abduction and approximated prediction, removing autoregressive constraints and significantly reducing computational and memory requirements.
3. We validate our methods’ effectiveness and performance on diverse synthetic and real-world datasets.

Related Work

Recent advances in deep generative models (DGMs) have found their way into learning Structural Causal Models (SCMs). Karimi et al. (2020) propose a conditional variational autoencoder (VAE) for each Markov factorization implied by the causal graph. VAEs on graphs are also studied assuming certain design constraints but have yet to

achieve empirical success (Zecevic et al. 2021; Sánchez-Martin, Rateike, and Valera 2022). (Kocaoglu et al. 2017) use generative adversarial network for learning a causal implicit generative model for a given causal graph. (Geffner et al. 2022) propose an autoregressive-flow based non-linear additive noise end-to-end framework for causal discovery and inference. (Pawlowski, Coelho de Castro, and Glocker 2020) introduce several generative models to learn SCMs, but these are not guaranteed to learn the true causal mechanism, as multiple models can produce the same observational distribution. Another DGM class, autoregressive normalizing flows, has also been suggested (Khemakhem et al. 2021). (Javaloy, Martin, and Valera 2023) generalize this by considering a class of triangular monotonically increasing maps that are identifiable up to invertible, component-wise transformations. However, this model design is restricted to monotonic functions and requires additional regularization to scale to multiple layers. (Scetbon et al. 2024) view SCMs as a fixed-point problem over causally ordered variables, infer causal ordering from data and use it to develop a fixed-point SCM via an attention-based autoencoder.

Diffusion models represent another competitive class of DGMs. Sanchez and Tsaftaris (2022) use diffusion models for counterfactual estimation in bivariate graphs where an image class is the parent of an image. However, this approach requires training a separate classifier (Dhariwal and Nichol 2021) for *do*-interventions, lacks theoretical guarantees, and shows poor performance for more complex images. (Chao et al. 2024) offer both interventional and counterfactual inferences, but only guarantee identifiability with a fully observed causal graph. Their method requires a separate neural network for each causal node and sequential inference, making it computationally expensive and memory-intensive for large causal graphs. In contrast, our proposed flow models are flexible in design while maintaining causal consistency. Our approach reduces computational and memory complexity, enhances scalability by avoiding separate neural networks for each causal node, and eliminates sequential inference, making it more practical for large causal graphs.

Alternative to DGMs, Ordinary Differential Equations (ODEs) have been used to describe deterministic SCMs, which is often unrealistic (Mooij, Janzing, and Schölkopf 2013). (Peters, Bauer, and Pfister 2022) introduce the Causal Kinetic Model, a collection of ODEs requiring parent values at each time step. (Hansen and Sokol 2014) illustrate a causal interpretation of SDE, show how to apply interventions to a SDE. (Wang, Jennings, and Gong 2023) combine neural SDE with variational inference to model causal structure in continuous-time time-series data..

Preliminaries

Structural Causal Models (SCMs)

Given a directed acyclic graph (DAG) graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ representing the causal relationships between d endogenous variables $x = \{x^1, \dots, x^d\}$, a SCM (Pearl 2009) \mathcal{M} associated with \mathcal{G} is a set of structural equations $x^i = f^i(x^{\text{pa}_i}, u^i)$ for all $i \in \{1, \dots, d\}$ that characterize how each

node x^i in \mathcal{V} is generated from its parent nodes $x^{\text{pa}_i} := \{x^j \mid \text{the directed edge } (j, i) \in \mathcal{E}\}$ and the corresponding exogenous variable u^i via a deterministic function f^i . Usually $u = \{u^1, \dots, u^d\}$ are assumed to be jointly independent, i.e., $p(u^1, \dots, u^d) = \prod_{i=1}^d p(u^i)$. This makes the SCM \mathcal{M} Markovian, leading to the factorization $p(x^1, \dots, x^d) = \prod_{i=1}^d p(x^i \mid x^{\text{pa}_i})$. Since \mathcal{G} is acyclic, we can specify a causal ordering π of all nodes such that if node x^j is a parent of node x^i then $\pi(j) < \pi(i)$.

For deep neural networks to answer causal questions, we can treat u as latent exogenous variables, encode exogenous variables to latent spaces $z = \tilde{f}_{\text{encode}}(x, x^{\text{pa}})$ as the abduction step, and decode back $x = \tilde{f}_{\text{decode}}(z, x^{\text{pa}})$ as the prediction step. This structure enables VAEs, GANs and Normalizing Flows to learn deep SCMs.

Causal consistency A mapping T between variables u and x is deemed causally consistent with structural causal model \mathcal{M} if it shares the same causal dependencies with the \mathcal{M} . It means their Jacobian matrices have zero values in the same positions, i.e., $\nabla_u T(u) \equiv I + \sum_{i=1}^{\text{diam}(A)} A^i$ and $\nabla_x T^{-1}(x) \equiv I - A$, where I is the identity matrix and A is the adjacency matrix of the causal graph.

Triangular Monotonically Increasing (TMI) Maps for Identifiable SCMs

A function $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a monotone increasing triangular map if:

$$f(x) = \begin{bmatrix} f_1(x_1) \\ f_2(x_1, x_2) \\ \vdots \\ f_d(x_1, \dots, x_d) \end{bmatrix}, \quad (1)$$

where each $f_i : \mathbb{R} \rightarrow \mathbb{R}$ is monotone increasing (or decreasing) with respect to x_i for any $x_{1:i-1}$. In case $\mu = \nu \cdot f$ where μ, ν are strictly positive density and f is a TMI map, then f is equivalent to the Knothe–Rosenblatt (KR) transport almost everywhere (Jaini, Selby, and Yu 2019).

Identifiability refers to recovering ground truth latent factors. (Xi and Bloem-Reddy 2023) show that nonlinear independent component analysis (ICA) models with generator functions that are TMI maps, and fully supported latent distributions with independent components, are identifiable up to invertible, component-wise transformations. This means we can recover the model up to an invertible, component-wise transformation of the true latent factors. In causal representation learning, on top of identifying the latent representation, the causal graph encoding their relations must also be identifiable. (Javaloy, Martin, and Valera 2023) further show that not only can the model isolate the exogenous variables, it also shares the functional dependencies with true structural equations. However, due to the TMI assumption, their model design must fulfill the monotonic requirements.

Diffusion and Flow Models

In continuous-time diffusion models, the stochastic process of generating data from a Gaussian prior can be represented

via a backward SDE (Song et al. 2021):

$$dx = (f(x, t) - g^2(t) \nabla_x \log p_t(x)) dt + g(t) d\bar{w} \quad (2)$$

where \bar{w} denotes the reverse-time Wiener process. The probability flow ODE (PF ODE) which shares the same marginal probability densities $p_t(x_t)$ as the SDE above is expressed as follow:

$$dx = \left(f(x, t) - \frac{1}{2} g(t)^2 \nabla_x \log p_t(x) \right) dt, \quad (3)$$

The PF ODE enables efficient mappings in both directions, from the data to the prior and vice versa. The PF ODE can be modeled via the score matching (Hyvärinen and Dayan 2005; Song and Ermon 2019) or flow matching framework (Lipman et al. 2023; Liu, Gong, and Liu 2023). In the latter case, the velocity $v_\theta(x, t)$ of the PF ODE is learned by minimizing the mean square error w.r.t. the target velocity $v(x_0, x_1, t)$ over time t sampled uniformly in $[0, 1]$:

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_t \mathbb{E}_{x_0, x_1} \left[\|v_\theta(x_t, t) - v(x_0, x_1, t)\|^2 \right] \quad (4)$$

To enable fast sampling, the data and prior distributions can be connected through a ‘‘stochastic interpolant’’ $X_t = (1-t)X_0 + tX_1$ (Albergo, Boffi, and Vanden-Eijnden 2023; Liu, Gong, and Liu 2023). In this case, $v(x_0, x_t, t) = x_1 - x_0$ and the PF ODE is called the Rectified Flow.

Method

In this section, we first prove that set of flow models are an identifiable and flexible choice in learning SCMs given only observational data and causal ordering. For ensuring identifiability we assume that SCMs are Markovian, acyclic with diffeomorphic structural equations. Then, we propose a compact model design that can do fast inference (prediction and abduction) in parallel.

Identifiable Causal Flow Models for Learning of SCMs with Ordering

We introduce a flow-based model for learning SCMs with ordering that is identifiable. Our approach involves constructing a set of flows, one for each node, so that i) the distribution of exogenous variables is a factorized distribution with support over the entire space \mathbb{R}^d , and ii) the mapping from endogenous variables to exogenous variables is a TMI map. This allows us to leverage a theoretical result from (Xi and Bloem-Reddy 2023) (Proposition 5.2) to prove for identifiability.

Let $p(u)$ represent a jointly factorized distribution of all d exogenous variables $u = \{u^1, \dots, u^d\}$, i.e., $p(u) = \prod_{i=1}^d p(u^i)$ where $p(u^i)$ is a Gaussian distribution in \mathbb{R} . We assume that while the structural causal model (SCM) \mathcal{M} is unknown, the causal ordering π among the nodes is known. We represent each node $x^i \in \mathbb{R}$ as the value at time $t = 1$ of the initial value problem (IVP) below:

$$dz_t^i = v^i(z_t^i, u^{<\pi_i}, t) dt, \quad z_0^i = u^i \quad (5)$$

where z_t^i denotes the state at time t , $u^{<\pi_i} := \{u^j | \pi_j < \pi_i\}$ is the set of nodes with lower orders than u^i according

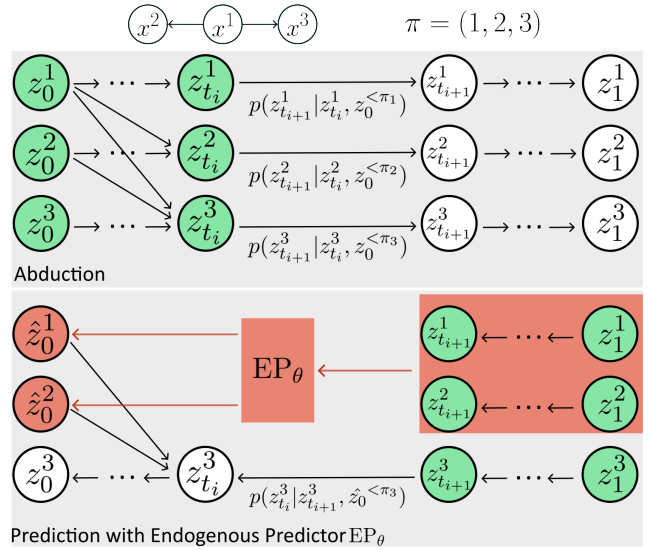


Figure 1: Parallel Causal Flow Model (P-CFM) of a simple SCM of 3 nodes (Top), unrolled into abduction (Middle) and prediction processes (Bottom). Green nodes are known, white nodes are to be calculated, and orange nodes are approximated. In abduction, all $\left\{ p \left(z_{t_{i+1}}^j | z_{t_i}^j, z_0^{<\pi_j} \right) \right\}_{j=1}^d$ can be calculated with one forward pass. In prediction, we show an example of how $p \left(z_{t_i}^3 | z_{t_{i+1}}^3, \hat{z}_0^{<\pi_3} \right)$ can be calculated by approximating z_0^1, z_0^2 using the endogenous predictor EP_θ given $z_{t_{i+1}}^1, z_{t_{i+1}}^2, z_1^1, z_1^2$. All $\left\{ p \left(z_{t_i}^j | z_{t_{i+1}}^j, \hat{z}_0^{<\pi_j} \right) \right\}_{j=1}^d$ can be calculated simultaneously.

to π . The solution at time t of the above IVP can be expressed as $z_t^i = z_0^i + \int_0^t v^i(z_\tau^i, u^{<\pi_i}, \tau) d\tau$, which means z_t^i can be regarded as a function of the initial value z_0^i . Let $f^i(u^i, u^{<\pi_i}, t) = f^i(z_0^i, u^{<\pi_i}, t) := z_t^i$ denote the representation of node i at time t , and let $f(u, t) := (f^1(u^1, u^{<\pi_1}, t), \dots, f^d(u^d, u^{<\pi_d}, t))$ represent the state of all nodes at time t . Below, we show that $f(u, t)$ is a TMI map of u satisfying two key properties: *monotonicity* and *triangularity*, as stated in Theorem 1.

Theorem 1. *Let $f(u, t)$ be the solution of the set of initial value problems (IVPs) for all nodes at time t , with the IVP for node i is described in Eq 5. If the velocity function $v^i(z_t^i, u^{<\pi_i}, t)$ is continuous w.r.t. t and Lipschitz continuous w.r.t. z_t^i for all $t \in (0, 1)$, $u^{<\pi_i} \in \mathbb{R}^{\pi_i-1}$, and $i \in \{1, \dots, d\}$, then f is a triangular monotonically increasing (TMI) map of u .*

Proof. First, we prove that $f^i(u^i, u^{<\pi_i}, t)$ is a monotonically increasing function of u^i . Since $u^{<i}$ is constant w.r.t. u^i and t , we can simplify the notation by denoting $\hat{f}^i(u^i, t) := f^i(u^i, u^{<\pi_i}, t)$ and consider \hat{f}^i as a function of u^i and t only. Suppose there exists two initial values a, b of u_i such that $a > b$ (or equivalently, $\hat{f}^i(a, 0) > \hat{f}^i(b, 0)$) but $\hat{f}^i(a, t) \leq \hat{f}^i(b, t)$. Since $\hat{f}^i(a, t)$ and $\hat{f}^i(b, t)$ are contin-




Figure 2: An illustration of the case where $\hat{f}^i(u^i, t)$ is not a monotonically increasing function of u^i for all $t \in (0, 1)$. Since $\hat{f}^i(u^i, t)$ is continuous w.r.t. t , we can find $\gamma \in (0, t]$ such that $\hat{f}^i(a, \gamma) = \hat{f}^i(b, \gamma) = \zeta$. This leads to two distinct solutions of the IVP starting at $z_\gamma^i = \zeta$, which contradicts the Picard-Lindelof theorem.

uous for all $t \in (0, 1)$ (due to the continuity of v^i w.r.t. t and z_t^i), there must exist $\gamma \in (0, t]$ such that $\hat{f}^i(a, \gamma) = \hat{f}^i(b, \gamma) = \zeta$ (as illustrated in Fig. 2). This implies that $\hat{f}^i(a, t), \hat{f}^i(b, t)$ are two distinct solutions of the IVP:

$$dz_t^i = v^i(z_t^i, u^{<\pi_i}, t) dt, \quad z_\gamma^i = \zeta. \quad (6)$$

with $t \in (\gamma, 1)$. This contradicts the Picard-Lindelof theorem, which states the IVP in Eq. 6 has a unique solution if $v^i(z_t^i, u^{<\pi_i}, t)$ is continuous w.r.t. t and Lipschitz continuous w.r.t. z_t^i (Simmons 2016; Chen et al. 2018). This means $\hat{f}^i(a, t) > \hat{f}^i(b, t)$ if $a > b$, meaning $\hat{f}^i(u^i, t)$ is a monotonically increasing function of u^i . Proving that f is a triangular map is straightforward since, by design, f^i is a function of nodes with orders lower than or equal to π_i . Consequently, f is a TMI map of u . \square

Let v_θ^i be a parameterized model of v^i . To ensure that $v_\theta^i(z_t^i, u^{<\pi_i}, t)$ is continuous w.r.t. t and Lipschitz continuous w.r.t. z_t^i , we model it using a feed-forward neural network with finite weights and 1-Lipschitz continuous activation functions (e.g., sigmoid, tanh, ReLU, ELU). A detailed explanation for this can be found in Appendix. The input to this network is the concatenation of $z_t^i, u^{<\pi_i}$, and t .

Since $u^{<\pi_i}$ is typically unknown and often varies during training $v_\theta^i(z_t^i, u^{<\pi_i}, t)$, we instead model $v_\theta^i(z_t^i, x^{<\pi_i}, t)$ in practice and consider this velocity in subsequent discussions. This velocity can be viewed as a reparameterization of $v_\theta^i(z_t^i, u^{<\pi_i}, t)$ because $x^{<\pi_i} = f(u^{<\pi_i}, 1)$.

We note that our method can be applied when the causal graph \mathcal{G} is available. In this case, we simply replace $x^{<\pi_i}$ in $v_\theta^i(z_t^i, x^{<\pi_i}, t)$ with x^{pa_i} , ensuring that our method remains consistent with the SCM \mathcal{M} .

Since v_θ^i represents the velocity of an ODE that maps $p(u^i) = \mathcal{N}(0, 1)$ to $p(x^i)$, we learn v_θ^i by first constructing a reference diffusion process between $p(x^i)$ and $p(u^i)$, and then aligning v_θ^i with the velocity of the PF ODE corresponding to this diffusion process. In this work, we select $Z_t^i = (1-t)U^i + tX^i$ as the reference process due to the sampling efficiency of its PF ODE. This choice results in the

following training objective for all v_θ^i ($i \in \{1, \dots, d\}$):

$$\mathcal{L}(\theta) = \sum_{i=1}^d \mathbb{E}_t \mathbb{E}_{x^i, u^i} \|v_\theta^i(z_t^i, x^{<\pi_i}, t) - (x^i - u^i)\|_2^2 + \lambda \|\theta\|_2^2 \quad (7)$$

In Eq. 7, we include a weight regularization term to ensure that θ is finite. We refer to our method as a *Causal Flow Model* (CFM). After training $v_\theta^i(z_t^i, x^{<\pi_i}, t)$ for all $i \in \{1, \dots, d\}$, we can compute x^i from u^i and $x^{<\pi_i}$ as $x^i = f^i(u^i, x^{<\pi_i}, 1)$ by solving the ODE in Eq. 5 forward in time. Similarly, we can compute u^i from x^i and $x^{<\pi_i}$ as $u^i = (f^i)^{-1}(x^i, x^{<\pi_i}, 0)$ by solving this backward in time. These two steps correspond to the prediction and abduction steps.

Efficient Causal Flow Models

Efficient Abduction with Masked Autoregressive Velocity nEural Network (MAVEN) To approximate the velocity v for all nodes while reducing time and memory complexity, we propose the MAVEN architecture. It comprises two key components: a Masked Autoregressive Neural Network (MADE $_\theta$) with a fixed ordering (Germain et al. 2015) and a simple feed-forward neural network h_θ . MADE $_\theta$ encodes x into $c \in \mathbb{R}^{d \times k}$, where each vector c^i contains node-specific encoded information dependent only on preceding variables. The feed-forward network then predicts node velocity by taking a concatenated vector of $[z_t^i, c^i, t, i]$.

$$c^i = \text{MADE}_\theta(x^{<\pi_i}) \quad (8)$$

$$v_t^i = h_\theta([z_t^i, c^i, t]) \quad (9)$$

$$v_\theta(z_t, x, t) = [h_\theta([z_t^1, c^1, t]), \dots, h_\theta([z_t^d, c^d, t])]^\top \quad (10)$$

We enforce the flow of each node to follow a linear path from the endogenous variable x^i to exogenous variable u^i . This is achieved through the prior ODE defined as $dx^i = (u^i - x^i) dt$. v_θ can be trained using gradient descent by solving the regression problem:

$$\theta^* = \text{argmin}_\theta \int_0^1 \mathbb{E} [\|v_\theta(z_t, x, t) - (u - x)\|_2^2] dt. \quad (11)$$

We call the proposed method as **Sequential-Causal Flow Model** (S-CFM). It optimizes the abduction process by reducing computational complexity from $O(nd)$ to $O(n)$, enabling more efficient processing of large-scale causal networks. See Abduction in Fig. 1 for an example of 3-node causal graph.

Next, we propose a method that reduce complexity of prediction step from $O(nd)$ to $O(n)$, enabling parallel prediction.

Endogenous predictor During prediction, we encounter challenges due to limited information about endogenous variables, necessitating a sequential, node-by-node prediction approach. At each time step t , the i -th endogenous variable x^i is deterministically derived from current and preceding exogenous variables $u^{<\pi_i} := \{u^j | \pi_j \leq \pi_i\}$. By

Algorithm 1: P-CFM Model Training.

Input: Observational data \mathcal{X} , causal ordering π , exogenous distribution $p(u)$

- 1: **while** not converged **do** ▷ Training S-CFM
- 2: $z_0 = x \sim \mathcal{X}$; $z_1 = u \sim p(u)$; $t \sim \text{Unif}[0, 1]$;
 $z_t = (1 - t) \cdot z_0 + t \cdot z_1$
- 3: Update parameters of velocity neural network v_θ , by minimizing the following loss:

$$\|v_\theta(z_t, x, t) - (z_1 - z_0)\|^2$$

4: **end while**

- 5: **while** not converged **do** ▷ Training Endogenous Predictor

- 6: $t \sim \text{Unif}[0, 1]$; $z_0 = x \sim \mathcal{X}$
- 7: Sample z_t, z_1 given z_0 from S-CFM
- 8: Update parameters of endogenous predictor neural network EP_θ , by minimizing the following loss:

$$\|\text{EP}_\theta(z_t, z_1, t) - z_0\|^2$$

9: **end while**

leveraging a neural network, specifically a MADE, denoted as $\text{EP}_\theta(z_t, u, t)$, we can accurately predict these variables. The predicted endogenous variables then serve as input for MAVEN to estimate velocity at each time step, effectively acting as a distillation mechanism for the flow model.

The endogenous predictor is trained by solving the following optimization problem:

$$\theta^* = \underset{\theta}{\text{argmin}} \int_0^1 \mathbb{E} [\|\text{EP}_\theta(z_t, u, t) - x\|^2] dt. \quad (12)$$

We introduce the Parallel-Causal Flow Model (P-CFM) by integrating S-CFM with an endogenous predictor. The training of the predictor leverages the abduction process, sampling from x with linear complexity. This approach ensures efficient training, enabling scalability and effective handling of larger datasets with minimal computational overhead.

See Fig. 1 for an example of P-CFM on a 3-node SCM. The training of P-CFM is summarized in Alg. 1, the abduction step in Alg. 3 and the prediction step in Alg. 2.

Do operator To enable intervention and counterfactual calculations, we modify Pearl’s *do*-operator $do(x^i = \alpha)$ differently for S-CFM and P-CFM. In S-CFM, we replace the i -th causal mechanism f^i with a constant function, $f_I^i = \alpha$. For P-CFM, we adopt the backtracking counterfactual approach (Von Kügelgen, Mohamed, and Beckers 2023), which modifies the exogenous distribution $p(u)$ while maintaining causal mechanisms f . An intervention $do(x^i = \alpha)$ updates $p(u)$ by restricting plausible u values to those yielding the intervened value α . The intervened SCM is defined as $\mathcal{M}^I = (f, p_I(u))$, with $p_I(u)$ density determined by:

$$p_I(u) = \delta(\{u^i : f_i(x^{\text{pa}_i}, u^i) = \alpha\}) \cdot \prod_{j \neq i} p(u^j). \quad (13)$$

During the prediction step, the endogenous predictor uses intervened exogenous values to predict endogenous val-

ues. Traditional *do*-operators require finding specific set (u^i, x^{pa_i}) that causes the intervened value α , necessitating an additional abduction step that can increase complexity from $O(1)$ to $O(nd)$ during iterative inference. In our approach where the abduction step can be run in parallel, this operator is computationally cheaper with a complexity of $O(n)$. As a result, our overall complexity for observational, interventional and counterfactual queries are $O(n)$. More specifically, these queries are detailed as follows:

- The observation is sampled from $p(x)$ as $x = \text{Predict}(u)$ for $u \sim p(u)$.
- The interventional sampling from $p(x \mid do(x^i = \alpha))$ involves the steps: $x = \text{Predict}(u)$ for $u \sim p(u)$, followed by $x^i = \alpha$, $u^i = \text{Abduct}(x^i)^i$, and $x^{\text{int}} = \text{Predict}(u)$
- The counterfactual sampling from $p(x^{\text{cf}} \mid do(x^i = \alpha), x^f)$ involves the steps: $u = \text{Abduct}(x^f)$, $x^{i,f} = \alpha$, $u^i = \text{Abduct}(x^f)^i$, followed by $x^{\text{cf}} = \text{Predict}(u)$.

Algorithm 2: Predict(u)

Input: u , sequence of times $0 = t_1 < t_2 < \dots < t_{N-1} < t_N = 1$

- 1: $z_1 = u$
 - 2: **for** $n = N$ **to** 2 **do**
 - 3: $\hat{z}_0 = \text{EP}_\theta(z_1, z_{t_n}, t_n)$ ▷ Predict endogenous values
 - 4: $v = v_\theta(z_{t_n}, \hat{z}_0, t_n)$
 - 5: $z_{t_{n-1}} = z_{t_n} - v \cdot (t_n - t_{n-1})$
 - 6: **end for**
 - 7: **Return** z_0
-

Algorithm 3: Abduct(x)

Input: x , sequence of times $0 = t_1 < t_2 < \dots < t_{N-1} < t_N = 1$

- 1: $z_0 = x$
 - 2: **for** $n = 1$ **to** $N - 1$ **do**
 - 3: $v = v_\theta(z_{t_n}, z_0, t_n)$
 - 4: $z_{t_{n+1}} = z_{t_n} + v \cdot (t_{n+1} - t_n)$
 - 5: **end for**
 - 6: **Return** z_1
-

Experiments

We evaluate the performance of the proposed sequential and parallel CFM on both synthetic and real-world dataset, as well as a major speed up by parallelization.

Settings

Implementation For MAVEN and endogenous predictor EP_θ , we use a MADE with three hidden layers [256, 256, 256] and ELU activation, and a fully connected neural network with the same layers and activation. We use the Adam optimizer with a learning rate of 0.001 and apply a decay factor of 0.95 to the learning rate if it remains at

SCM	Metric	S-CFM	P-CFM	CausalNF-NSF	CausalNF-MAF	VACA
		($\times 10^{-2}$)	($\times 10^{-2}$)	($\times 10^{-2}$)	($\times 10^{-2}$)	($\times 10^{-2}$)
Triangle NADD	Obs. MMD	0.67 _{0.07}	<i>0.65</i> _{0.15}	0.43 _{0.08}	2.03 _{0.14}	2.37 _{0.25}
	Int. MMD	3.21 _{0.34}	<i>3.16</i> _{0.24}	1.86 _{0.09}	13.31 _{0.29}	15.53 _{1.24}
	CF. MSE	6.53 _{0.34}	<i>9.59</i> _{3.53}	41.90 _{11.74}	9.91 _{1.02}	146.73 _{16.29}
Simpson NLIN2	Obs. MMD	0.57 _{0.05}	<i>0.61</i> _{0.05}	0.62 _{0.05}	3.54 _{0.21}	2.44 _{0.22}
	Int. MMD	0.58 _{0.01}	<i>0.60</i> _{0.04}	0.61 _{0.01}	4.34 _{0.10}	2.55 _{0.32}
	CF. MSE	0.72 _{0.18}	<i>0.73</i> _{0.18}	1.98 _{0.86}	1.29 _{0.20}	15.28 _{1.25}
Diamond NLIN	Obs. MMD	0.36 _{0.02}	<i>0.38</i> _{0.06}	0.44 _{0.04}	2.26 _{4.15}	5.23 _{0.35}
	Int. MMD	<i>0.41</i> _{0.05}	<i>0.41</i> _{0.06}	0.34 _{0.02}	4.49 _{8.95}	26.64 _{0.63}
	CF. MSE	<i>0.04</i> _{0.04}	0.03 _{0.01}	16.06 _{2.07}	197.27 _{438.22}	60.62 _{3.45}
Y NADD	Obs. MMD	0.46 _{0.05}	0.46 _{0.05}	<i>0.55</i> _{0.07}	1.01 _{0.05}	1.68 _{0.08}
	Int. MMD	0.48 _{0.03}	<i>0.51</i> _{0.04}	1.75 _{0.70}	1.56 _{0.13}	4.44 _{0.20}
	CF. MSE	28.58 _{1.04}	<i>29.67</i> _{0.84}	37.31 _{3.32}	30.27 _{0.78}	32.45 _{0.72}
LargeBD NLIN	Obs. MMD	<i>0.40</i> _{0.03}	0.39 _{0.02}	0.42 _{0.04}	0.97 _{0.06}	1.34 _{1.14}
	Int. MMD	0.41 _{0.02}	0.48 _{0.03}	<i>0.43</i> _{0.03}	0.99 _{0.03}	0.98 _{0.06}
	CF. MSE	0.02 _{0.00}	<i>0.03</i> _{0.00}	0.04 _{0.00}	0.04 _{0.00}	0.71 _{0.01}
LargeLadder NLIN	Obs. MMD	0.63 _{0.01}	<i>0.64</i> _{0.00}	0.66 _{0.04}	0.65 _{0.05}	-
	Int. MMD	0.64 _{0.01}	<i>0.65</i> _{0.00}	0.66 _{0.00}	<i>0.65</i> _{0.00}	-
	CF. MSE	<i>42.71</i> _{3.12}	46.79 _{6.67}	72.63 _{5.4}	38.46 _{1.30}	-
LargeLadder NADD	Obs. MMD	0.58 _{0.06}	<i>0.61</i> _{0.02}	0.62 _{0.03}	0.71 _{0.05}	-
	Int. MMD	0.62 _{0.00}	<i>0.63</i> _{0.00}	0.62 _{0.00}	0.68 _{0.00}	-
	CF. MSE	27.85 _{1.39}	<i>29.09</i> _{0.29}	36.23 _{3.49}	37.42 _{1.94}	-

Table 1: Mean and standard deviation of five SCMs with non-linear (NLIN) and non-additive (NADD) structural equations. Bold indicates the best result; italic indicates the second-best.

a plateau for more than 10 epochs, batch size of 2048, and train for 900 epochs. Inference process is done in 50 steps.

Synthetic datasets We consider synthetic SCMs various in the number of nodes and edges as they allow us to have direct access to the true observational, interventional and counterfactual distributions to evaluate methods. We study two main classes of structural equations:

>*Non-linear additive noise* (NLIN): **Simpson**, a 4-node SCM simulating a Simpson’s paradox that is difficult to approximate. **Diamond**, a 4-node SCM with 4 edges arranged in a diamond shape. **Large Backdoor** (LargeBD), a 9-node non-Gaussian noise with sparse causal graph.

>*Non-additive noise* (NADD): **Triangle**, a 3-node SCM with confounding graph in which x^1 is the parent of x^2 , x^3 and x^2 causes x^3 . **Y**, a 4-node Y-shaped SCM.

We randomly sample 50,000/5,000/5,000 samples for training, evaluation and testing, respectively. See appendix for more details on the structural equations and all results.

Baselines For the synthetic experiments, we compare our methods with two recently state-of-the-art methods Causal Normalizing Flows (CausalNF) (Javaloy, Martin, and Valera 2023) and VACA (Sánchez-Martin, Rateike, and Valera 2022). For CausalNF, we use two types of flow architectures: Masked Autoregressive Flow (MAF) (Papamakarios, Pavlakou, and Murray 2017) and Neural Spline Flow (NSF) (Durkan et al. 2019). As a requirements of VACA, we use

fully observed causal graph for training in testing VACA. Due to the differences in settings, we do not show comparison between our methods and Diffusion-based Causal Models (DCM) (Chao et al. 2024), which requires fully observed causal graph. When only the causal ordering is observed, training the DCM requires either multiple random selections of parents for each node or full causal graph discovery, which are both expensive. We re-implement these methods based on provided code. For a fair comparison, each model uses the same budget for optimization. For the real datasets, we compare with CausalNF-MAF, CausalNF-NSF, CAREFL (Khemakhem et al. 2021), regression function from an additive noise model (ANM) (Hoyer et al. 2008) and standard ridge regression model (Linear SCM).

Metrics To evaluate the methods on the three-layer causal hierarchy, we report the Maximum Mean Discrepancy (MMD) (Gretton et al. 2012) between the true and the estimated observational and interventional distributions. We do not evaluate the Average Treatment Effect (ATE) as in (Javaloy, Martin, and Valera 2023), as the metric does not provide insights into the distribution of treatment effects within the population. For counterfactual estimation, we report the mean squared error (MSE) between the actual and estimated counterfactual values.




Figure 3: Inference time (in seconds) against number of nodes d for 10 samples with 50 discretization steps on **Observational**, **Interventional**, and **Counterfactual** queries, for data generated using linear structural equations, ranging from 5 to 50 nodes.

Results on Synthetic Datasets

Table 1 reports the experimental results on five synthetic SCMs for observational, interventional and counterfactual questions. Each model’s performance is averaged over five random initializations of parameters and training datasets. VACA demonstrates poor performance across all datasets, which can be attributed to the limitations of the encoder-decoder architectures and the constraints of the graph neural network. CausalNF-NSF and CausalNF-MAF show competitive results compared to ours. However, due to the affine constraints of the Masked Autoregressive Flow, CausalNF-MAF only learns an affine transformation between exogenous and endogenous distributions, performing well in cases with additive structural equations but failing in non-additive cases where the endogenous distribution is not Gaussian. CausalNF-NSF closely approximates observational and interventional distributions in both additive and non-additive cases but struggles with answering counterfactual questions.

Our methods exhibit a significantly higher and more consistent performance compared to existing methods across the majority of queries, over both additive and non-additive relationships. Specifically, P-CFM, despite approximating endogenous values x , demonstrates a superior performance relative to previous methods.

Scalability Fig. 3 demonstrates the scalability of P-CFM against previous Diffusion-based methods such as DCM, and our own S-CFM, all using the same number of parameters. DCM’s inference time grows linearly with the number of nodes d . S-CFM significantly reduces counterfactual generation time through parallel abduction. In contrast, P-CFM maintains a nearly constant inference time regardless of d . This result highlights that P-CFM is not only accurate but also efficient, making it suitable for larger problems.

Results on Real Dataset

We evaluate S-CFM on real-world setting on the electrical stimulation interventional fMRI data (Thompson et al. 2020), experimental setup is obtained from (Khemakhem et al. 2021). We do not evaluate P-CFM since it does not show much advantages in case of only two nodes. The fMRI data includes time series from the Cingulate Gyrus (CG) and Heschl’s Gyrus (HG) for 14 patients with medically refractory epilepsy. The underlying simplified causal structure is a

Algorithm	Median Abs. Error
S-CFM	0.5829 _{5.0e-2}
CausalNF-NSF	0.6018 _{2.0e-2}
CausalNF-MAF	0.6025 _{3.0e-2}
CAREFL	0.5983 _{3.0e-2}
ANM	0.6746 _{4.8e-8}
Linear SCM	0.6045 _{0.0}

Table 2: Interventional prediction’s median absolute error on fMRI data. The results for CAREFL, ANM, and Linear SCM are obtained from Chao et al. (2023)

bivariate graph with $CG \rightarrow HG$. Our task is to predict value of HG given intervened CG.

Table 2 reports the interventional results on fMRI datasets. S-CFM demonstrates slightly better performance compared to the alternatives. This is because the evaluation is based on the absolute error of a single intervention, which does not account for distributional information. Additionally, the causal ordering is consistent with the causal structure in the bivariate case. Nonetheless, the experiment provides a benchmark for comparing various causal inference algorithms using real datasets.

Conclusion

We have introduced a new scalable method of causal learning and inference using flow models: We proved that a flow-based models can learn identifiable SCMs from only observational data and causal ordering. To make the models efficient, we designed a masked autoregressive velocity network and an endogenous predictor that enables parallel inference. We empirically demonstrated that our method outperforms SOTA rivals on a wide range of SCMs over non-additive and non-linear relationships. We showed that our parallel causal flow model is scalable, maintaining a nearly constant inference time regardless of the number of variables. We validated our method on real-world datasets, including an fMRI study, showcasing its practical applicability.

References

- Albergo, M. S.; Boffi, N. M.; and Vanden-Eijnden, E. 2023. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*.
- Chao, P.; Blöbaum, P.; Patel, S. K.; and Kasiviswanathan, S. 2024. Modeling Causal Mechanisms with Diffusion Models for Interventional and Counterfactual Queries.
- Chen, R. T.; Rubanova, Y.; Bettencourt, J.; and Duvenaud, D. K. 2018. Neural ordinary differential equations. *Advances in neural information processing systems*, 31.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- Durkan, C.; Bekasov, A.; Murray, I.; and Papamakarios, G. 2019. Neural spline flows. *Advances in neural information processing systems*, 32.
- Geffner, T.; Antoran, J.; Foster, A.; Gong, W.; Ma, C.; Kiciman, E.; Sharma, A.; Lamb, A.; Kukla, M.; Pawlowski, N.; et al. 2022. Deep end-to-end causal inference. *arXiv preprint arXiv:2202.02195*.
- Germain, M.; Gregor, K.; Murray, I.; and Larochelle, H. 2015. Made: Masked autoencoder for distribution estimation. In *International conference on machine learning*, 881–889. PMLR.
- Gretton, A.; Borgwardt, K. M.; Rasch, M. J.; Schölkopf, B.; and Smola, A. 2012. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1): 723–773.
- Hansen, N.; and Sokol, A. 2014. Causal interpretation of stochastic differential equations.
- Hoyer, P.; Janzing, D.; Mooij, J. M.; Peters, J.; and Schölkopf, B. 2008. Nonlinear causal discovery with additive noise models. *Advances in neural information processing systems*, 21.
- Hyvärinen, A.; and Dayan, P. 2005. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4).
- Jaini, P.; Selby, K. A.; and Yu, Y. 2019. Sum-of-squares polynomial flow. In *International Conference on Machine Learning*, 3009–3018. PMLR.
- Javaloy, A.; Martin, P. S.; and Valera, I. 2023. Causal normalizing flows: from theory to practice. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Karimi, A.-H.; Von Kügelgen, J.; Schölkopf, B.; and Valera, I. 2020. Algorithmic recourse under imperfect causal knowledge: a probabilistic approach. *Advances in neural information processing systems*, 33: 265–277.
- Khemakhem, I.; Monti, R.; Leech, R.; and Hyvarinen, A. 2021. Causal Autoregressive Flows. In Banerjee, A.; and Fukumizu, K., eds., *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, 3520–3528. PMLR.
- Kocaoglu, M.; Snyder, C.; Dimakis, A. G.; and Vishwanath, S. 2017. CausalGAN: Learning causal implicit generative models with adversarial training. *arXiv preprint arXiv:1709.02023*.
- LeRoy, S. 2004. *Causality in economics*. London School of Economics, Centre for Philosophy of Natural and Social Sciences.
- Lipman, Y.; Chen, R. T.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2023. Flow Matching for Generative Modeling. In *The Eleventh International Conference on Learning Representations*.
- Liu, X.; Gong, C.; and Liu, Q. 2023. Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow. In *The Eleventh International Conference on Learning Representations*.
- Mooij, J. M.; Janzing, D.; and Schölkopf, B. 2013. From ordinary differential equations to structural causal models: the deterministic case. *arXiv preprint arXiv:1304.7920*.
- Nguyen, T.; Do, K.; Nguyen, D. T.; Duong, B.; and Nguyen, T. 2023. Causal inference via style transfer for out-of-distribution generalisation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1746–1757.
- Papamakarios, G.; Pavlakou, T.; and Murray, I. 2017. Masked autoregressive flow for density estimation. *Advances in neural information processing systems*, 30.
- Pawlowski, N.; Coelho de Castro, D.; and Glocker, B. 2020. Deep structural causal models for tractable counterfactual inference. *Advances in neural information processing systems*, 33: 857–869.
- Pearl, J. 2009. *Causality*. Cambridge University Press, 2 edition.
- Peters, J.; Bauer, S.; and Pfister, N. 2022. Causal models for dynamical systems. In *Probabilistic and Causal Inference: The Works of Judea Pearl*, 671–690.
- Russo, F.; and Williamson, J. 2007. Interpreting causality in the health sciences. *International studies in the philosophy of science*, 21(2): 157–170.
- Sanchez, P.; and Tsaftaris, S. A. 2022. Diffusion Causal Models for Counterfactual Estimation. In *CLEaR*.
- Sánchez-Martin, P.; Rateike, M.; and Valera, I. 2022. VACA: Designing variational graph autoencoders for causal queries. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 8159–8168.
- Scetbon, M.; Jennings, J.; Hilmkil, A.; Zhang, C.; and Ma, C. 2024. FiP: a Fixed-Point Approach for Causal Generative Modeling. *arXiv preprint arXiv:2404.06969*.
- Schölkopf, B.; Locatello, F.; Bauer, S.; Ke, N. R.; Kalchbrenner, N.; Goyal, A.; and Bengio, Y. 2021. Toward causal representation learning. *Proceedings of the IEEE*, 109(5): 612–634.
- Simmons, G. F. 2016. *Differential equations with applications and historical notes*. CRC Press.
- Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*.

Thompson, W. H.; Nair, R.; Oya, H.; Esteban, O.; Shine, J.; Petkov, C.; Poldrack, R.; Howard, M.; and Adolphs, R. 2020. A data resource from concurrent intracranial stimulation and functional MRI of the human brain. *Scientific data*, 7(1): 258.

Von Kügelgen, J.; Mohamed, A.; and Beckers, S. 2023. Backtracking counterfactuals. In *Conference on Causal Learning and Reasoning*, 177–196. PMLR.

Wang, B.; Jennings, J.; and Gong, W. 2023. Neural structure learning with stochastic differential equations. *arXiv preprint arXiv:2311.03309*.

Xi, Q.; and Bloem-Reddy, B. 2023. Indeterminacy in generative models: Characterization and strong identifiability. In *International Conference on Artificial Intelligence and Statistics*, 6912–6939. PMLR.

Zecevic, M.; Dhami, D. S.; Velickovic, P.; and Kersting, K. 2021. Relating Graph Neural Networks to Structural Causal Models. *CoRR*, abs/2109.04173.