

Deep Structural Causal Models for Tractable Counterfactual Inference

Nick Pawlowski*
Imperial College London
np716@imperial.ac.uk

Daniel C. Castro*
Imperial College London
dc315@imperial.ac.uk

Ben Glocker
Imperial College London
b.glocker@imperial.ac.uk

Abstract

We formulate a general framework for building structural causal models (SCMs) with deep learning components. The proposed approach employs normalising flows and variational inference to enable tractable inference of exogenous noise variables—a crucial step for counterfactual inference that is missing from existing deep causal learning methods. Our framework is validated on a synthetic dataset built on MNIST as well as on a real-world medical dataset of brain MRI scans. Our experimental results indicate that we can successfully train deep SCMs that are capable of all three levels of Pearl’s ladder of causation: association, intervention, and counterfactuals, giving rise to a powerful new approach for answering causal questions in imaging applications and beyond. The code for all our experiments is available at <https://github.com/biomed-mira/deepscm>.

1 Introduction

Many questions in everyday life as well as in scientific inquiry are causal in nature: “How would the climate have changed if we’d had less emissions in the ’80s?”, “How fast could I run if I hadn’t been smoking?”, or “Will my headache be gone if I take that pill?”. None of those questions can be answered with statistical tools alone, but require methods from causality to analyse interactions with our environment (interventions) and hypothetical alternate worlds (counterfactuals), going beyond joint, marginal, and conditional probabilities [1]. Even though these are natural lines of reasoning, their mathematical formalisation under a unified theory is relatively recent [2].

In some statistics-based research fields, such as econometrics or epidemiology, the use of causal inference methods has been established for some time [3, 4]. However, causal approaches have been introduced into deep learning (DL) only very recently [5]. For example, research has studied the use of causality for disentanglement [6, 7], causal discovery [8, 9], and for deriving causality-inspired explanations [10, 11] or data augmentations [12]. Causal DL models could be capable of learning relationships from complex high-dimensional data and of providing answers to interventional and counterfactual questions, although current work on deep counterfactuals is limited by modelling only direct cause-effect relationships [11] or instrumental-variable scenarios [13], or by not providing a full recipe for tractable counterfactual inference [14].

The integration of causality into DL research promises to enable novel scientific advances as well as to tackle known shortcomings of DL methods: DL is known to be susceptible to learning spurious correlations and amplifying biases [e.g. 15], and to be exceptionally vulnerable to changes in the input distribution [16]. By explicitly modelling causal relationships and acknowledging the difference between causation and correlation, causality becomes a natural field of study for improving the transparency, fairness, and robustness of DL-based systems [17, 18]. Further, the tractable inference of deep counterfactuals enables novel research avenues that aim to study causal reasoning on a

*Joint first authors.

per-instance rather than population level, which could lead to advances in personalised medicine as well as in decision-support systems, more generally.

In this context, our work studies the use of DL-based causal mechanisms and establishes effective ways of performing counterfactual inference. Our main contributions are: 1) a unified framework for structural causal models using modular deep mechanisms; 2) an efficient approach to estimating counterfactuals by inferring exogenous noise via variational inference or normalising flows; 3) case studies exemplifying how to apply deep structural causal models and perform counterfactual inference. The paper is organised as follows: we first review structural causal models and discuss how to leverage deep mechanisms and enable tractable counterfactual inference. Second, we compare our work to recent progress in deep causal learning in light of Pearl’s ladder of causation [19]. Finally, we apply deep structural causal models to a synthetic experiment as well as to modelling brain MRI scans, demonstrating the practical utility of our framework in answering counterfactual questions.

2 Deep Structural Causal Models

We consider the problem of modelling a collection of K random variables $\mathbf{x} = (x_1, \dots, x_K)$. By considering causal relationships between them, we aim to build a model that not only is capable of generating convincing novel samples, but also satisfies all three rungs of the causation ladder [19]. The first level, **association**, describes reasoning about passively observed data. This level deals with correlations in the data and questions of the type “*What are the odds that I observe...?*”, which relates purely to marginal, joint, and conditional probabilities. **Intervention** concerns interactions with the environment. It requires knowledge beyond just observations, as it relies on structural assumptions about the underlying data-generating process. Characteristic questions ask about the effects of certain actions: “*What happens if I do...?*”. Lastly, **counterfactuals** deal with retrospective hypothetical scenarios. Counterfactual queries leverage functional models of the generative processes to imagine alternative outcomes for individual data points, answering “*What if I had done A instead of B?*”. Arguably, such questions are at the heart of scientific reasoning (and beyond), yet are less well-studied in the field of machine learning. The three levels of causation can be operationalised by employing structural causal models (SCMs)², recapitulated in the next section.

2.1 Background on structural causal models

A structural causal model $\mathcal{G} := (\mathbf{S}, P(\epsilon))$ consists of a collection $\mathbf{S} = (f_1, \dots, f_K)$ of structural assignments $x_k := f_k(\epsilon_k; \mathbf{pa}_k)$ (called *mechanisms*), where $\mathbf{pa}_k$ is the set of parents of x_k (its *direct causes*), and a joint distribution $P(\epsilon) = \prod_{k=1}^K P(\epsilon_k)$ over mutually independent exogenous noise variables (i.e. unaccounted sources of variation). As assignments are assumed acyclic, relationships can be represented by a directed acyclic graph (DAG) with edges pointing from causes to effects, called the *causal graph* induced by $\mathcal{G}$. Every SCM $\mathcal{G}$ entails a unique joint observational distribution $P_{\mathcal{G}}(\mathbf{x})$, satisfying the causal Markov assumption: each variable is independent of its non-effects given its direct causes. It therefore factorises as $P_{\mathcal{G}}(\mathbf{x}) = \prod_{k=1}^K P_{\mathcal{G}}(x_k | \mathbf{pa}_k)$, where each conditional distribution $P_{\mathcal{G}}(x_k | \mathbf{pa}_k)$ is determined by the corresponding mechanism and noise distribution [1].

Crucially, unlike conventional Bayesian networks, the conditional factors above are imbued with a causal interpretation. This enables $\mathcal{G}$ to be used to predict the effects of *interventions*, defined as substituting one or multiple of its structural assignments, written as ‘ $\text{do}(\dots)$ ’. In particular, a constant reassignment of the form $\text{do}(x_k := a)$ is called an atomic intervention, which disconnects x_k from all its parents and represents a direct manipulation disregarding its natural causes.

While the observational distribution relates to statistical associations and interventions can predict causal effects, SCMs further enable reasoning about *counterfactuals*. These are hypothetical retrospective interventions, given an observed outcome: ‘What would x_i have been if x_j were different, given that we observed $\mathbf{x}$?’ This type of question effectively offers explanations of the data, since we can analyse the changes resulting from manipulating each variable. Counterfactual queries can be mathematically formulated as a three-step procedure [2, Ch. 7]:

1. **Abduction:** Predict the ‘state of the world’ (the exogenous noise, ϵ) that is compatible with the observations, $\mathbf{x}$, i.e. infer $P_{\mathcal{G}}(\epsilon | \mathbf{x})$.

²SCMs are also known as (nonlinear) structural equation models or functional causal models.




Figure 1: Classes of deep causal mechanisms considered in this work. Bi-directional arrows indicate invertible transformations, optionally conditioned on other inputs (edges ending in black circles). Black and white arrowheads refer resp. to the generative and abductive directions, while dotted arrows depict an amortised variational approximation. Here, f_k is the forward model, e_k is an encoder that amortises abduction in non-invertible mechanisms, g_k is a ‘high-level’ non-invertible branch (e.g. a probabilistic decoder), and h_k is a ‘low-level’ invertible mapping (e.g. reparametrisation).

2. **Action:** Perform an intervention (e.g. $\text{do}(x_k := \tilde{x}_k)$) corresponding to the desired manipulation, resulting in a modified SCM $\tilde{\mathfrak{G}} = \mathfrak{G}_{\mathbf{x}; \text{do}(\tilde{x}_k)} = (\tilde{\mathbf{S}}, P_{\tilde{\mathfrak{G}}}(\epsilon|\mathbf{x}))$ [1, Sec. 6.4].
3. **Prediction:** Compute the quantity of interest based on the distribution entailed by the counterfactual SCM, $P_{\tilde{\mathfrak{G}}}(\mathbf{x})$.

With these operations in mind, the next section explores a few options for building flexible, expressive, and counterfactual-capable functional mechanisms for highly structured data.

2.2 Deep mechanisms

In statistical literature (e.g. epidemiology, econometrics, sociology), SCMs are typically employed with simple linear mechanisms (or generalised linear models, involving an output non-linearity). Analysts attach great importance to the regression weights, as under certain conditions these may be readily interpreted as estimates of the causal effects between variables. While this approach generally works well for scalar variables and can be useful for decision-making, it is not flexible enough to model higher-dimensional data such as images. Solutions to this limitation have been proposed by introducing deep-learning techniques into causal inference [8, 14].

We call an SCM that uses deep-learning components to model the structural assignments a *deep structural causal model* (DSCM). In DSCMs, the inference of counterfactual queries becomes more complex due to the potentially intractable abduction step (inferring the posterior noise distribution, as defined above). To overcome this, we propose to use recent advances in normalising flows and variational inference to model mechanisms for composable DSCMs that enable tractable counterfactual inference. While here we focus on continuous data, DSCMs also fully support discrete variables without the need for relaxations (see Appendix C). We consider three types of mechanisms that differ mainly in their invertibility, illustrated in Fig. 1.

Invertible, explicit: Normalising flows model complex probability distributions using transformations from simpler base distributions with same dimensionality [20]. For an observed variable x , diffeomorphic transformation f , and base variable $\epsilon \sim P(\epsilon)$ such that $x = f(\epsilon)$, the output density $p(x)$ can be computed as $p(x) = p(\epsilon)|\det \nabla f(\epsilon)|^{-1}$, evaluated at $\epsilon = f^{-1}(x)$ [21, 22]. For judicious choices of f , the Jacobian ∇f may take special forms with efficiently computable determinant, providing a flexible and tractable probabilistic model whose parameters can be trained via exact maximum likelihood. Furthermore, flows can be made as expressive as needed by composing sequences of simple transformations. For more information on flow-based models, refer to the comprehensive survey by Papamakarios et al. [22]. Note that this class of models also subsumes the typical location-scale and inverse cumulative distribution function transformations used in the reparametrisation trick [23, 24], as well as the Gumbel trick for discrete variable relaxations [25, 26].

Although normalising flows were originally proposed for unconditional distributions, they have been extended to conditional densities [27], including in high dimensions [28, 29], by parametrising the transformation as $x = f(\epsilon; \mathbf{pa}_X)$, assumed invertible in the first argument. In particular, conditional

flows can be adopted in DSCMs to represent invertible, explicit-likelihood mechanisms (Fig. 1a):

$$x_i := f_i(\epsilon_i; \mathbf{pa}_i), \quad p(x_i | \mathbf{pa}_i) = p(\epsilon_i) \cdot |\det \nabla_{\epsilon_i} f_i(\epsilon_i; \mathbf{pa}_i)|^{-1} \Big|_{\epsilon_i = f_i^{-1}(x_i; \mathbf{pa}_i)}. \quad (1)$$

Amortised, explicit: Such invertible architectures typically come with heavy computational requirements when modelling high-dimensional observations, because all intermediate operations act in the space of the data. Instead, it is possible to use arbitrary functional forms for the structural assignments, at the cost of losing invertibility and tractable likelihoods $p(x_k | \mathbf{pa}_k)$. Here, we propose to separate the assignment f_k into a ‘low-level’, invertible component h_k and a ‘high-level’, non-invertible part g_k —with a corresponding noise decomposition $\epsilon_k = (u_k, z_k)$ —such that

$$x_k := f_k(\epsilon_k; \mathbf{pa}_k) = h_k(u_k; g_k(z_k; \mathbf{pa}_k), \mathbf{pa}_k), \quad P(\epsilon_k) = P(u_k)P(z_k). \quad (2)$$

In such a decomposition, the invertible transformation h_k can be made shallower, while the upstream non-invertible g_k maps from a lower-dimensional space and is expected to capture more of the high-level structure of the data. Indeed, a common implementation of this type of model for images would involve a probabilistic decoder, where g_k may be a convolutional neural network, predicting the parameters of a simple location-scale transformation performed by h_k [24].

As the conditional likelihood $p(x_k | \mathbf{pa}_k)$ in this class of models is no longer tractable because z_k cannot be marginalised out, it may alternatively be trained with amortised variational inference. Specifically, we can introduce a variational distribution $Q(z_k | x_k, \mathbf{pa}_k)$ to formulate a lower bound on the true marginal conditional log-likelihood, which will be maximised instead:

$$\log p(x_k | \mathbf{pa}_k) \geq \mathbb{E}_{Q(z_k | x_k, \mathbf{pa}_k)}[\log p(x_k | z_k, \mathbf{pa}_k)] - D_{\text{KL}}[Q(z_k | x_k, \mathbf{pa}_k) \| P(z_k)]. \quad (3)$$

The argument of the expectation in this lower bound can be calculated similarly to Eq. (1):

$$p(x_k | z_k, \mathbf{pa}_k) = p(u_k) \cdot |\det \nabla_{u_k} h_k(u_k; g_k(z_k, \mathbf{pa}_k), \mathbf{pa}_k)|^{-1} \Big|_{u_k = h_k^{-1}(x_k; g_k(z_k, \mathbf{pa}_k), \mathbf{pa}_k)}. \quad (4)$$

The approximate posterior distribution $Q(z_k | x_k, \mathbf{pa}_k)$ can for example be realised by an encoder function, $e_k(x_k; \mathbf{pa}_k)$, that outputs the parameters of a simple distribution over z_k (Fig. 1b), as in the auto-encoding variational Bayes (AEVB) framework [24].

Amortised, implicit: While the models above rely on (approximate) maximum-likelihood as training objective, it is admissible to train a non-invertible mechanism as a conditional implicit-likelihood model (Fig. 1c), optimising an adversarial objective [30–32]. Specifically, a deterministic encoder e_j would strive to fool a discriminator function attempting to tell apart tuples of encoded real data $(x_j, e_j(x_j; \mathbf{pa}_j), \mathbf{pa}_j)$ and generated samples $(f_j(\epsilon_j; \mathbf{pa}_j), \epsilon_j, \mathbf{pa}_j)$.

2.3 Deep counterfactual inference

Now equipped with effective deep models for representing mechanisms in DSCMs, we discuss the inference procedure allowing us to compute answers to counterfactual questions.

Abduction: As presented in Section 2.1, the first step in computing counterfactuals is abduction, i.e. to predict the exogenous noise, ϵ , based on the available evidence, $\mathbf{x}$. Because each noise variable is assumed to affect only the respective observed variable, $(\epsilon_k)_{k=1}^K$ are conditionally independent given $\mathbf{x}$, therefore this posterior distribution factorises as $P_{\Theta}(\epsilon | \mathbf{x}) = \prod_{k=1}^K P_{\Theta}(\epsilon_k | x_k, \mathbf{pa}_k)$. In other words, it suffices to infer the noise independently for each mechanism, given the observed values of the variable and of its parents³.

For invertible mechanisms, the noise variable can be obtained deterministically and exactly by just inverting the mechanism: $\epsilon_i = f_i^{-1}(x_i; \mathbf{pa}_i)$. Similarly, implicit-likelihood mechanisms can be approximately inverted by using the trained encoder function: $\epsilon_j \approx e_j(x_j; \mathbf{pa}_j)$.

Some care must be taken in the case of amortised, explicit-likelihood mechanisms, as the ‘high-level’ noise z_k and ‘low-level’ noise u_k are not independent given x_k . Recalling that this mechanism is

³Note that here we assume full observability, i.e. no variables are missing when predicting counterfactuals. We discuss challenges of handling partial evidence in Section 6.

trained along with a conditional probabilistic encoder, $Q(z_k | e_k(x_k; \mathbf{pa}_k))$, the noise posterior can be approximated as follows, where $\delta_w(\cdot)$ denotes the Dirac delta distribution centred at w :

$$\begin{aligned} P_{\mathfrak{G}}(\epsilon_k | x_k, \mathbf{pa}_k) &= P_{\mathfrak{G}}(z_k | x_k, \mathbf{pa}_k) P_{\mathfrak{G}}(u_k | z_k, x_k, \mathbf{pa}_k) \\ &\approx Q(z_k | e_k(x_k; \mathbf{pa}_k)) \delta_{h_k^{-1}(x_k; g_k(z_k; \mathbf{pa}_k), \mathbf{pa}_k)}(u_k). \end{aligned} \quad (5)$$

Action: The causal graph is then modified according to the desired hypothetical intervention(s), as in the general case (Section 2.1). For each intervened variable x_k , its structural assignment is replaced either by a constant, $x_k := \tilde{x}_k$ —making it independent of its former parents (direct causes, $\mathbf{pa}_k$) and of its exogenous noise (ϵ_k)—or by a surrogate mechanism $x_k := \tilde{f}_k(\epsilon_k; \tilde{\mathbf{pa}}_k)$, forming a set of counterfactual assignments, $\tilde{\mathbf{S}}$. This then defines a counterfactual SCM $\tilde{\mathfrak{G}} = (\tilde{\mathbf{S}}, P_{\mathfrak{G}}(\epsilon | \mathbf{x}))$.

Prediction: Finally, we can sample from $\tilde{\mathfrak{G}}$. Noise variables that were deterministically inverted (either exactly or approximately) can simply be plugged back into the respective forward mechanism to determine the new output value. Notice that this step is redundant for observed variables that are not descendants of the ones being intervened upon, as they will be unaffected by the changes.

As mentioned above, the posterior distribution over (z_k, u_k) for an amortised, explicit-likelihood mechanism does not factorise (Eq. (5)), and the resulting distribution over the counterfactual x_k cannot be characterised explicitly. However, sampling from it is straightforward, such that we can approximate the counterfactual distribution via Monte Carlo as follows, for each sample s :

$$\begin{aligned} z_k^{(s)} &\sim Q(z_k | e_k(x_k; \mathbf{pa}_k)) \\ u_k^{(s)} &= h_k^{-1}(x_k; g_k(z_k^{(s)}; \mathbf{pa}_k), \mathbf{pa}_k) \\ \tilde{x}_k^{(s)} &= \tilde{h}_k(u_k^{(s)}; \tilde{g}_k(z_k^{(s)}; \tilde{\mathbf{pa}}_k), \tilde{\mathbf{pa}}_k). \end{aligned} \quad (6)$$

Consider an uncorrelated Gaussian decoder for images as a concrete example, predicting vectors of means and variances for each pixel of x_k : $g_k(z_k; \mathbf{pa}_k) = (\mu(z_k; \mathbf{pa}_k), \sigma^2(z_k; \mathbf{pa}_k))$. Exploiting the reparametrisation trick, counterfactuals that preserve x_k 's mechanism can be computed simply as

$$u_k^{(s)} = (x_k - \mu(z_k^{(s)}; \mathbf{pa}_k)) \oslash \sigma(z_k^{(s)}; \mathbf{pa}_k), \quad \tilde{x}_k^{(s)} = \mu(z_k^{(s)}; \tilde{\mathbf{pa}}_k) + \sigma(z_k^{(s)}; \tilde{\mathbf{pa}}_k) \odot u_k^{(s)},$$

where $\oslash$ and $\odot$ denote element-wise division and multiplication, respectively. In particular, in the constant-variance setting adopted for our experiments, counterfactuals further simplify to

$$\tilde{x}_k^{(s)} = x_k + [\mu(z_k^{(s)}; \tilde{\mathbf{pa}}_k) - \mu(z_k^{(s)}; \mathbf{pa}_k)].$$

This showcases how true image counterfactuals are able to retain pixel-level details. Typical conditional generative models would output only $\mu(z_k; \tilde{\mathbf{pa}}_k)$ (which is often blurry in vanilla variational auto-encoders [33]), or would in addition have to sample $P(u_k)$ (resulting in noisy images).

3 Related Work

Deep generative modelling has seen a wide range of contributions since the popularisation of variational auto-encoders (VAEs) [24], generative adversarial networks (GANs) [34], and normalising flows [21]. These models have since been employed to capture conditional distributions [27, 29, 32, 35], and VAEs and GANs were also extended to model structured data by incorporating probabilistic graphical models [36–38]. In addition, deep generative models have been heavily used for (unsupervised) representation learning with an emphasis on disentanglement [39–42]. However, even when these methods faithfully capture the distribution of observed data, they are capable of fulfilling only the association rung of the ladder of causation.

Interventions build on the associative capabilities of probabilistic models to enable queries related to changes in causal mechanisms. By integrating a causal graph into the connectivity of a deep model, it is possible to perform interventions with GANs [14] and causal generative NNs [8]. VAEs can also express causal links using specific covariance matrices between latent variables, which however restrict the dependences to be linear [6]. Despite reaching the second rung of the causal ladder, these methods lack tractable abduction capabilities and therefore cannot generate counterfactuals.




Figure 2: Computational graphs of the structural causal models for the Morpho-MNIST example. The image is denoted by x , stroke thickness by t , and image intensity by i . The corresponding causal diagrams are displayed in the top-right corners.

Some machine-learning tasks such as explainability, image-to-image translation, or style transfer are closely related to counterfactual queries of the sort ‘How would x (have to) change if we (wished to) modify y ?’. Here, y could be the style of a picture for style transfer [43], the image domain (e.g. drawing to photo) for image-to-image translation [44], the age of a person in natural images [45] or medical scans [46], or a predicted output for explainability [11]. However, these approaches do not explicitly model associations, interventions, nor causal structure. Potentially closest to our work is a method for counterfactual explainability of visual models, which extends CausalGANs [14] to predict reparametrised distributions over image attributes following an assumed causal graph [10]. However, this approach performs no abduction step, instead resampling the noise of attributes downstream from the intervention(s), and does not include a generative model of imaging data. To the best of our knowledge, the proposed DSCM framework is the first flexible approach enabling end-to-end training and tractable inference on all three levels of the ladder of causation for high-dimensional data.

4 Case Study 1: Morpho-MNIST

We consider the problem of modelling the causal model of a synthetic dataset based on MNIST digits [47], where stroke thickness causes the brightness of the digit: thicker digits are brighter whereas thinner digits are dimmer. This simple dataset allows for examining the three levels of causation in a controlled and measurable environment. We use morphological transformations on MNIST [48] to generate a dataset with known causal structure and access to the ‘true’ process of generating counterfactuals. The SCM for this synthetic dataset is as follows:

$$\begin{aligned} t &:= f_T^*(\epsilon_T^*) = 0.5 + \epsilon_T^*, & \epsilon_T^* &\sim \Gamma(10, 5). \\ i &:= f_I^*(\epsilon_I^*, t) = 191 \cdot \sigma(0.5 \cdot \epsilon_I^* + 2 \cdot t - 5) + 64, & \epsilon_I^* &\sim \mathcal{N}(0, 1). \\ x &:= f_X^*(\epsilon_X^*; i, t) = \text{SetIntensity}(\text{SetThickness}(\epsilon_X^*; t); i), & \epsilon_X^* &\sim \text{MNIST}, \end{aligned} \quad (7)$$

where $\text{SetIntensity}(\cdot; i)$ and $\text{SetThickness}(\cdot; t)$ refer to the operations that act on an image of a digit and set its intensity to i and thickness to t (see Appendix A.1 for details), x is the resulting image, ϵ^* is the exogenous noise for each variable and $\sigma(\cdot)$ is the logistic sigmoid.

We use this setup to study the capabilities of our framework in comparison to models with less causal structure. We adapt the true causal graph from Eq. (7) and model thickness and intensity using (conditional) normalising flows and employ a conditional VAE for modelling the image. In particular, we adopt the causal graphs shown in Fig. 2 and test a fully independent model (Fig. 2a), a conditional decoder model (Fig. 2b), as well as our full causal model (Fig. 2c). All our experiments were implemented within PyTorch [49] using the Pyro probabilistic programming framework [50], and implementation details can be found in Appendices A.2 and B.2.

We quantitatively compare the associative capabilities of all models by evaluating their evidence lower bound (Eq. (3)), log-likelihoods and reconstruction errors as shown in Table 1. We find that performance improves consistently with the model’s capabilities: enabling conditional image generation improves $p(x|t, i)$, and adding a causal dependency between t and i improves $p(i|t)$. Further, we examine samples of the conditional and unconditional distributions in Appendix A.3.1.

The interventional distributions can be directly compared to the true generative process. Figure 3 shows that the densities predicted by our full model after intervening on t closely resemble the

Table 1: Comparison of the associative abilities of the models shown in Fig. 2. The image is denoted by x , thickness by t , and intensity by i . Quantities with $\geq$ are lower bounds. MAE refers to the mean absolute error between pixels of the original image and of its reconstruction.

Model	$\log p(x, t, i) \geq$	$\log p(x t, i) \geq$	$\log p(t)$	$\log p(i t)$	MAE(x, x')
Independent	-5925.26	-5919.14	-0.93	-5.19	4.50
Conditional	-5526.50	-5520.37	-0.93	-5.19	4.26
Full	-5692.94	-5687.71	-0.93	-4.30	4.43




Figure 3: Distributions of thickness and intensity in the true data (left), and learned by the full (centre) and conditional (right) models. Contours depict the observational (red, shaded) and interventional joint densities for $\text{do}(t := f_T(\epsilon_T) + 1)$ (blue, solid) and $\text{do}(t := f_T(\epsilon_T) - 0.5)$ (green, dashed).




Figure 4: Counterfactuals generated by the full model. (left) Counterfactual ‘trajectories’ of two original samples, A and B, as their thickness and intensity are modified, overlaid on the learned joint density $p(t, i)$. (right) Original and counterfactual images corresponding to samples A and B.

true behaviour. The conditional and independent models operate equivalently and are incapable of modelling the relationship between t and i , capturing only their marginal distributions.

Lastly, we examine the full model’s ability to generate counterfactuals. The other two models were omitted as they are incapable of accomplishing interventions, a prerequisite for counterfactual inference. Examples of previously unseen images and generated counterfactuals are shown in Fig. 4. We see that our model is capable of generating convincing counterfactuals that preserve the digit identity while changing thickness and intensity consistently with the underlying causal model.

5 Case Study 2: Brain Imaging

Our real-world application touches upon fundamental scientific questions in the context of medical imaging: how would a person’s anatomy change if particular traits were different? We illustrate with a (simplified) example that our DSCM framework may provide the means to answer such counterfactual




Figure 5: Brain imaging example. Variables are image (x), age (a), sex (s), and brain (b) and ventricle (v) volumes. The counterfactuals show different interventions on the same original brain.

queries, which may enable entirely new research into better understanding the physical manifestation of lifestyle, demographics, and disease. Here, we model the appearance of brain MRI scans given the person’s age and biological sex, as well as brain and ventricle volumes⁴, using population data from the UK Biobank [51]. Ventricle and total brain volumes are two quantities that are closely related to brain age [52] and can be observed relatively easily. We adopt the causal graph shown in Fig. 5a and otherwise follow the same training procedure as for the MNIST experiments.

The learned DSCM is capable of all three levels of the causal hierarchy. We present the analysis of lower levels in Appendix B.3.1 and focus here on counterfactuals, shown in Fig. 5b (more examples in Appendix B.3.2). The difference maps show plausible counterfactual changes: increasing age causes slightly larger ventricles while decreasing the overall brain volume (first column). In contrast, directly changing brain volume has an opposite effect on the ventricles compared to changing age (second column). Intervening on ventricle volume has a much more localised effect (third column), while intervening on the categorical variable of biological sex has smaller yet more diffuse effects. Note how the anatomical ‘identity’ (such as the cortical folding) is well preserved after each intervention.

6 Conclusion

We introduce a novel general framework for fitting SCMs with deep mechanisms. Our deep SCM (DSCM) framework fulfils all three rungs of Pearl’s ladder of causation—in particular, it is the first to enable efficient abduction of exogenous noise, permitting principled counterfactual inference. We demonstrate the potential of DSCMs with two case studies: a synthetic task of modelling Morpho-MNIST digits with a known causal structure and a real-world example with brain MRI.

The ability to correctly generate plausible counterfactuals could greatly benefit a wide variety of possible applications, e.g.: *explainability*, where differences between observed and counterfactual data can suggest causal explanations of outcomes; *data augmentation*, as counterfactuals can extrapolate beyond the range of observed data (e.g. novel combinations of attributes); and *domain adaptation*, since including the source of the data as an indicator variable in the causal model could enable generating counterfactual examples in a relevant target domain.

The proposed method does not come without limitations to be investigated in future work. Like the related approaches, the current setup requires all variables to be observed when computing a counterfactual, which may limit its applicability in certain scenarios. This could be alleviated by imputing the missing data via MCMC or learning auxiliary distributions. Further work should study more closely the dynamic behaviour of deep mechanisms in SCMs. While not observed in our experiments, neural networks may not learn to cleanly separate the roles of its inputs on the output

⁴Ventricles are fluid-filled cavities identified as the dark areas in the centre of the brain.

as expected—which could require custom counterfactual regularisation similar to losses used in image-to-image translation [46] and explainability work [11]. The use of such flexible models also raises questions about the identifiability of the ‘true’ mechanism, as counterfactuals may not be uniquely defined. Lastly, it would be interesting to examine whether this framework can be applied to causal discovery, attempting to uncover plausible causal structures from data.

Broader Impact

Causal inference can be applied to a wide range of applications, promising to provide a deeper understanding of the observed data and prevent the fitting of spurious correlations. Our research presents a methodological contribution to the causal literature proposing a framework that combines causal models and deep learning to facilitate modelling high-dimensional data.

Because of the general applicability of deep learning and causal inference, our framework could have a broad impact of enabling fairer machine learning models explicitly modelling causal mechanisms, reducing spurious correlations and tackling statistical and societal biases. The resulting models offer better interpretability due to counterfactual explanations and could yield novel understanding through causal discovery.

However, causal modelling relies on strong assumptions and cannot always unambiguously determine the true causal structure of observational data. It therefore is necessary to carefully consider and communicate the assumptions being made by the analyst. In this light, our methodology is susceptible to being used to wrongly claim the discovery of causal structures due to careless application or intentional misuse. Particularly, the use of ‘black-box’ components as causal mechanisms may exacerbate concerns about identifiability, already present even for simple linear models. Whereas deep causal models can be useful for deriving insights from data, we must be cautious about their use in consequential decision-making, such as in informing policies or in the context of healthcare.

Acknowledgements

We thank Thanos Vlontzos for helpful comments on a draft of this paper. This research received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 757173, project MIRA, ERC-2017-STG). NP is supported by a Microsoft Research PhD Scholarship. DC and NP are also supported by the EPSRC Centre for Doctoral Training in High Performance Embedded and Distributed Systems (HiPEDS, grant ref EP/L016796/1). We gratefully acknowledge the support of NVIDIA with the donation of one Titan X GPU. The UK Biobank data is accessed under Application Number 12579.

References

- [1] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press, Cambridge, MA, 2017.
- [2] Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2nd edition, 2009.
- [3] H. Wold. Causality and econometrics. *Econometrica*, 22(2):162–177, 1954. URL <http://www.jstor.org/stable/1907540>.
- [4] Sander Greenland, Judea Pearl, and James M. Robins. Causal diagrams for epidemiologic research. *Epidemiology*, 10(1):37–48, 1999. URL <http://www.jstor.org/stable/3702180>.
- [5] Bernhard Schölkopf. Causality for machine learning. *arXiv preprint arXiv:1911.10500*, 2019.
- [6] Mengyue Yang, Furui Liu, Zhitang Chen, Xinwei Shen, Jianye Hao, and Jun Wang. CausalVAE: Structured causal disentanglement in variational autoencoder. *arXiv preprint arXiv:2004.08697*, 2020.
- [7] Giambattista Parascandolo, Niki Kilbertus, Mateo Rojas-Carulla, and Bernhard Schölkopf. Learning independent causal mechanisms. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *PMLR*, pages 4036–4044. PMLR, 2018.

- [8] Olivier Goudet, Diviyani Kalainathan, Philippe Caillou, Isabelle Guyon, David Lopez-Paz, and Michèle Sebag. Learning functional causal models with generative neural networks. In Hugo Jair Escalante, Sergio Escalera, Isabelle Guyon, Xavier Baró, Yağmur Güçlütürk, Umut Güçlü, and Marcel van Gerven, editors, *Explainable and Interpretable Models in Computer Vision and Machine Learning*, pages 39–80. Springer International Publishing, Cham, 2018. doi: 10.1007/978-3-319-98131-4_3.
- [9] Yoshua Bengio, Tristan Deleu, Nasim Rahaman, Nan Rosemary Ke, Sebastien Lachapelle, Olexa Bilaniuk, Anirudh Goyal, and Christopher Pal. A meta-transfer objective for learning to disentangle causal mechanisms. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=ryxWigBFPS>.
- [10] Álvaro Parafita Martínez and Jordi Vitrià Marca. Explaining visual models by causal attribution. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 4167–4175. IEEE, 2019.
- [11] Sumedha Singla, Brian Pollack, Junxiang Chen, and Kayhan Batmanghelich. Explanation by progressive exaggeration. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=H1xFWgrFPS>.
- [12] Divyansh Kaushik, Eduard Hovy, and Zachary C. Lipton. Learning the difference that makes a difference with counterfactually-augmented data. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SkIgs0NFvr>.
- [13] Jason Hartford, Greg Lewis, Kevin Leyton-Brown, and Matt Taddy. Deep IV: A flexible approach for counterfactual prediction. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *PMLR*, pages 1414–1423. PMLR, 06–11 Aug 2017.
- [14] Murat Kocaoglu, Christopher Snyder, Alexandros G. Dimakis, and Sriram Vishwanath. Causal-GAN: Learning causal implicit generative models with adversarial training. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=BJE-4xWOW>.
- [15] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2979–2989, 2017.
- [16] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations*, 2014. URL <http://arxiv.org/abs/1312.6199>.
- [17] Matt J. Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. Counterfactual fairness. In *Advances in Neural Information Processing Systems 30 (NIPS 2017)*, pages 4066–4076, 2017.
- [18] Adarsh Subbaswamy, Peter Schulam, and Suchi Saria. Preventing failures due to dataset shift: Learning predictive models that transport. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics (AISTATS 2019)*, volume 89 of *PMLR*, pages 3118–3127. PMLR, 2019.
- [19] Judea Pearl. The seven tools of causal inference, with reflections on machine learning. *Communications of the ACM*, 62(3):54–60, feb 2019. doi: 10.1145/3241036.
- [20] Esteban G. Tabak and Cristina V. Turner. A family of nonparametric density estimation algorithms. *Communications on Pure and Applied Mathematics*, 66(2):145–164, feb 2013. doi: 10.1002/cpa.21423.
- [21] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *PMLR*, pages 1530–1538. PMLR, 2015.
- [22] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *arXiv preprint arXiv:1912.02762*, 2019.

- [23] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *PMLR*, pages 1278–1286. PMLR, 2014.
- [24] Diederik P. Kingma and Max Welling. Auto-encoding variational Bayes. In *International Conference on Learning Representations*, 2014. URL <http://arxiv.org/abs/1312.6114>.
- [25] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with Gumbel-softmax. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=rkE3y85ee>.
- [26] Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. The Concrete distribution: A continuous relaxation of discrete random variables. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=S1jE5L5gl>.
- [27] Brian L. Trippe and Richard E. Turner. Conditional density estimation with Bayesian normalising flows. In *NIPS 2017 Workshop on Bayesian Deep Learning*, 2017. URL <http://arxiv.org/abs/1802.04908>.
- [28] You Lu and Bert Huang. Structured output learning with conditional generative flows. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020)*. AAAI, 2020. URL <http://arxiv.org/abs/1905.13288>. To appear.
- [29] Christina Winkler, Daniel Worrall, Emiel Hoogeboom, and Max Welling. Learning likelihoods with conditional normalizing flows. *arXiv preprint arXiv:1912.00042*, 2019.
- [30] Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=BJtNZAFgg>.
- [31] Vincent Dumoulin, Ishmael Belghazi, Ben Poole, Olivier Mastropietro, Alex Lamb, Martin Arjovsky, and Aaron Courville. Adversarially learned inference. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=B1Elr4cgg>.
- [32] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [33] Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. Autoencoding beyond pixels using a learned similarity metric. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *PMLR*, pages 1558–1566. PMLR, 2016.
- [34] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27 (NIPS 2014)*, pages 2672–2680, 2014.
- [35] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems 28 (NIPS 2015)*, pages 3483–3491, 2015.
- [36] Matthew J. Johnson, David K. Duvenaud, Alex Wiltschko, Ryan P. Adams, and Sandeep R. Datta. Composing graphical models with neural networks for structured representations and fast inference. In *Advances in Neural Information Processing Systems 29 (NIPS 2016)*, pages 2946–2954, 2016.
- [37] Wu Lin, Mohammad Emtiyaz Khan, and Nicolas Hubacher. Variational message passing with structured inference networks. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=HyH9lbZAW>.
- [38] Chongxuan Li, Max Welling, Jun Zhu, and Bo Zhang. Graphical generative adversarial networks. In *Advances in Neural Information Processing Systems 31 (NeurIPS 2018)*, pages 6069–6080, 2018.

- [39] Tejas D. Kulkarni, William F. Whitney, Pushmeet Kohli, and Josh Tenenbaum. Deep convolutional inverse graphics network. In *Advances in Neural Information Processing Systems 28 (NIPS 2015)*, pages 2539–2547, 2015.
- [40] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems 29 (NIPS 2016)*, pages 2172–2180, 2016.
- [41] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. β -VAE: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=Sy2fzU9gl>.
- [42] Ashis Pati and Alexander Lerch. Attribute-based regularization of VAE latent spaces. *arXiv preprint arXiv:2004.05485*, 2020.
- [43] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2414–2423, 2016.
- [44] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1125–1134, 2017.
- [45] G. Antipov, M. Baccouche, and J. Dugelay. Face aging with conditional generative adversarial networks. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 2089–2093, 2017.
- [46] Tian Xia, Agisilaos Chartsias, Chengjia Wang, and Sotirios A. Tsafaris. Learning to synthesise the ageing brain without longitudinal data. *arXiv preprint arXiv:1912.02620*, 2019.
- [47] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.
- [48] Daniel C. Castro, Jeremy Tan, Bernhard Kainz, Ender Konukoglu, and Ben Glocker. Morpho-MNIST: Quantitative assessment and diagnostics for representation learning. *Journal of Machine Learning Research*, 20(178), 2019.
- [49] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, pages 8024–8035, 2019.
- [50] Eli Bingham, Jonathan P. Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul Szerlip, Paul Horsfall, and Noah D. Goodman. Pyro: Deep universal probabilistic programming. *Journal of Machine Learning Research*, 20(28), 2019.
- [51] Cathie Sudlow, John Gallacher, Naomi Allen, Valerie Beral, Paul Burton, John Danesh, Paul Downey, Paul Elliott, Jane Green, Martin Landray, et al. UK Biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS Medicine*, 12(3), 2015.
- [52] Ruth Peters. Ageing and the brain. *Postgraduate Medical Journal*, 82(964):84–88, 2006.
- [53] Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows. In *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, pages 7511–7522, 2019.
- [54] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp, 2016. URL <https://arxiv.org/abs/1605.08803>.

- [55] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [56] Fidel Alfaro-Almagro, Mark Jenkinson, Neal K. Bangerter, Jesper L. R. Andersson, Ludovica Griffanti, Gwenaëlle Douaud, Stamatios N. Sotiropoulos, Saad Jbabdi, Moises Hernandez-Fernandez, Emmanuel Vallee, et al. Image processing and quality control for the first 10,000 brain imaging datasets from UK Biobank. *NeuroImage*, 166:400–424, 2018.
- [57] Michael Oberst and David Sontag. Counterfactual off-policy evaluation with Gumbel-Max structural causal models. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *PMLR*, pages 4881–4890. PMLR, 2019.
- [58] Chris J. Maddison and Daniel Tarlow. Gumbel machinery, jan 2017. URL <https://cmaddis.github.io/gumbel-machinery>.

A Synthetic Morpho-MNIST Experiment

A.1 Data Generation

We use the original MNIST dataset [47] together with the morphometric measurements introduced with Morpho-MNIST [48] to add functionality to measure intensity as well as set the intensity and thickness to a given value.

We implement MeasureIntensity by following the processing steps proposed by Castro et al. [48], and measure the intensity i of an image as the median intensity of pixels within the extracted binary mask. Once the intensity is measured, the entire image is rescaled to match the target intensity, with values clamped between 0 and 255 (images are assumed to be in unsigned 8-bit format).

Originally, Morpho-MNIST only proposed relative thinning and thickening operations. We expand those operations to absolute values by calculating the amount of dilation or erosion based on the ratio between target thickness and measured thickness.

Finally, we follow Eq. (7) to modify each image within the MNIST dataset and randomly split the original training set into a training and validation set. We show random samples from the resulting test set in Fig. A.1.




Figure A.1: Random exemplars from the synthetically generated Morpho-MNIST test dataset

A.2 Experimental Setup

We use (conditional) normalising flows for all variables apart from the images, which we model using (conditional) deep encoder-decoder architectures. The flows consist of components that constrain the support of the output distribution (where applicable) and components relevant for fitting the distribution. We use unit Gaussians as base distributions for all exogenous noise distributions $P(\epsilon)$ and, if available, we use the implementations in PyTorch [49] or Pyro [50] for all transformations. Otherwise, we adapt the available implementations, referring to [53] for details. We indicate with θ the modules with learnable parameters.

We model the mechanisms of the thickness t and intensity i as

$$t := f_T(\epsilon_T) = (\exp \circ \text{AffineNormalisation} \circ \text{Spline}_\theta)(\epsilon_T), \quad (\text{A.1})$$

$$i := f_I(\epsilon_I; t) = (\text{AffineNormalisation} \circ \text{sigmoid} \circ \text{ConditionalAffine}_\theta(\hat{t}))(\epsilon_I). \quad (\text{A.2})$$

In the independent model, where i is not conditioned on t , we use instead

$$i := f_I(\epsilon_I) = (\text{AffineNormalisation} \circ \text{sigmoid} \circ \text{Spline}_\theta \circ \text{Affine}_\theta)(\epsilon_I). \quad (\text{A.3})$$

We found that including normalisation layers help learning dynamics⁵ and therefore include flows to perform commonly used normalisation transformations. For doubly bounded variable y we learn the flows in unconstrained space and then constrain them by a sigmoid transform and rescale to the original range using fixed affine transformations with bias $\min(Y)$ and scale $[\max(Y) - \min(Y)]$. We constrain singly bounded values by applying an exponential transform to the unbounded values and using an affine normalisation equivalent to a whitening operation in unbounded log-space. We denote those fixed normalisation transforms as $\text{AffineNormalisation}$ and use a hat to refer to the unconstrained, normalised values (e.g. $\widehat{\mathbf{p}}\mathbf{a}_k$). The Spline_θ transformation refers to first-order neural spline flows [53], Affine_θ is an element-wise affine transformation, and sigmoid refers to the logistic function. $\text{ConditionalAffine}_\theta(\cdot)$ is a regular affine transform whose transformation parameters are predicted by a context neural network taking $\cdot$ as input. In the case of $f_I(\epsilon_I; t)$, the context network is represented by a simple linear transform. Further, we model x using a low-level flow:

$$h_X(u_X; \mathbf{p}\mathbf{a}_X) = [\text{Preprocessing} \circ \text{ConditionalAffine}_\theta(\widehat{\mathbf{p}}\mathbf{a}_X)](u_X), \quad (\text{A.4})$$

where the ConditionalAffine transform practically reparametrises the noise distribution into another Gaussian distribution and Preprocessing describes a fixed preprocessing transformation. We follow the same preprocessing as used with RealNVP [54]. The context network for the conditional affine transformation is the high-level mechanism $g_X(z_X; \mathbf{p}\mathbf{a}_X)$ and is implemented as a decoder network that outputs the bias for of the affine transformation, while the log-variance is fixed to $\log \sigma^2 = -5$. We implement the decoder network as a CNN:

$$\begin{aligned} g_X(z_X; \mathbf{p}\mathbf{a}_X) = & (\text{Conv}_\theta(1; 1; 1; 0) \circ \text{ConvTranspose}_\theta(1; 4; 2; 1) \circ \text{ReLU} \circ \text{BN}_\theta \\ & \circ \text{ConvTranspose}_\theta(64; 4; 2; 1) \circ \text{Reshape}(64, 7, 7) \\ & \circ \text{ReLU} \circ \text{BN}_\theta \circ \text{Linear}_\theta(1024) \\ & \circ \text{ReLU} \circ \text{BN}_\theta \circ \text{Linear}_\theta(1024))([z_X, \widehat{\mathbf{p}}\mathbf{a}_X]), \end{aligned} \quad (\text{A.5})$$

where the operators describe neural network layers as follows: BN is batch normalisation; ReLU the ReLU activation function; $\text{Conv}(c; k; s; p)$ and $\text{ConvTranspose}(c; k; s; p)$ are a convolution or transposed convolution using a kernel with size k , a stride of s , a padding of p and outputting c channels; $\text{Linear}(h)$ is a linear layer with h output neurons; and $\text{Reshape}(\cdot)$ reshapes its inputs into the given shape $\cdot$. Lastly, $[z_X, \mathbf{p}\mathbf{a}_X]$ denotes the concatenation of z_X and $\mathbf{p}\mathbf{a}_X$, and $z_X \in \mathbb{R}^{16}$.

Equivalently, we implement the the encoder function as a simple CNN that outputs mean and log-variance of a independent Gaussian:

$$\begin{aligned} e_X(x; \mathbf{p}\mathbf{a}_X) = & ([\text{Linear}_\theta(16), \text{Linear}_\theta(16)] \circ [\text{LeakyReLU}(0.1), \widehat{\mathbf{p}}\mathbf{a}_X] \\ & \circ \text{BN}_\theta \circ \text{Linear}_\theta(100) \circ \text{Reshape}(128 \cdot 7 \cdot 7) \\ & \circ \text{LeakyReLU}(0.1) \circ \text{BN}_\theta \circ \text{Conv}_\theta(128; 4; 2, 1) \\ & \circ \text{LeakyReLU}(0.1) \circ \text{BN}_\theta \circ \text{Conv}_\theta(64; 4; 2, 1))(x), \end{aligned} \quad (\text{A.6})$$

where $\text{LeakyReLU}(\ell)$ is the leaky ReLU activation function with a leakiness of ℓ .

We use Adam [55] for optimisation with batch size of 256 and a learning rate of 10^{-4} for the encoder-decoder and 0.005 for the covariate flows. We set the number of particles (MC samples) for estimating the ELBO to 4. We use 32 MC samples for estimating reconstruction and counterfactuals. We train all models for 1000 epochs and report the results of the model with the best validation loss.

A.3 Additional Results

Here we further illustrate the associative, interventional, and counterfactual capabilities of the trained independent, conditional, and full models. (Continued on the next page.)

⁵We observed that not normalising the inputs can lead to the deep models prioritising learning the dependence on the variable with largest magnitude. This phenomenon should be investigated further.

A.3.1 Association




Figure A.2: Random samples generated by the independent, conditional and full model. Note how all models appear to have the same unconditional generation capacity.




Figure A.3: Conditional samples generated by the independent, conditional, and full model. The high-level noise, z_X , is shared for all samples from each model, ensuring the same ‘style’ of the generated digit. The independent model generates images independent of the thickness and intensity values, resulting in identical samples. For the conditional and full models, thickness and intensity change consistently along each column and row, respectively.




Figure A.4: Reconstructions. These are computed as Monte Carlo averages approximating $\mathbb{E}_{Q(z_X|e_X(x;\mathbf{p}_X))}[g_X(z_X;\mathbf{p}_X)]$, where e_X and g_X are the image encoder and decoder networks. All models seem capable of producing faithful reconstructions.




Figure A.5: Comparison of the target covariates and the corresponding values measured from the generated images. The leftmost column refers to the accuracy of the SetThickness and SetIntensity transforms used in generating the synthetic dataset, and the remaining three columns describe the fidelity of samples generated by each of the learned models. While images sampled from the independent model are trivially inconsistent with the sampled covariates, the conditional and full models show comparable conditioning performance.

A.3.2 Intervention




Figure A.6: Difference between conditioning and intervening, based on the trained full model. The joint density $p(t, i)$ is shown as contours in the background, for reference, and the ‘violin’ shapes represent the density of one variable when conditioning or intervening on three different values of the other variable. Since t causes i , notice how $p(t|i)$ (left) is markedly different from $p(t|do(i))$ (middle), which collapses to $p(t)$. On the other hand, $p(i|do(t))$ and $p(i|t)$ (right) are identical.

A.3.3 Counterfactual




Figure A.7: Original samples and counterfactuals from the full model. The first column shows the original image and true values of the non-imaging data. The even rows show the difference maps between the original image and the corresponding counterfactual image. We observe that all counterfactuals preserve the digits’ identity and style. Our model even generates sensible counterfactual images (with some artefacts) in very low-density regions, e.g. ‘0’ with $do(i = 64)$ (thick but dim), and very far from the original, e.g. ‘2’ with $do(t = 5.0)$.

B Brain Modelling

B.1 Data Generation

The original three-dimensional (3D) T1-weighted brain MRI scans have been pre-processed by the data providers of the UK Biobank Imaging study using the FSL neuroimaging toolkit [56]. The pre-processing involves skull removal, bias field correction, and automatic segmentation of brain structures. In addition, we have rigidly registered all scans to the standard MNI atlas space using an in-house image registration tool, which enabled us to extract anatomically corresponding mid-axial 2D slices that were used for the experiments presented in this paper. The 2D slices were normalised in intensity by mapping the minimum and maximum values inside the brain mask to the range $[0, 255]$. Background pixels outside the brain were set to zero. Age and biological sex for each subject were retrieved from the UK Biobank database along with the pre-computed brain and ventricle volumes. These volumes are derived from the 3D segmentation maps obtained with FSL, and although these

are image-derived measurements, they may serve as reasonable proxies of the true measurements within our (simplified yet plausible) causal model of the physical manifestation of the brain anatomy.




Figure B.1: Random exemplars from the test set of the adopted UK Biobank dataset

B.2 Experimental Setup

The setup for the brain imaging experiment closely follows the MNIST example as described in Appendix A.2. We randomly split the available 13,750 brain images into train, validation and test sets with the respective ratios 70%, 15% and 15%. During training, we randomly crop the brain slices from their original size of $233 \text{ px} \times 197 \text{ px}$ to $192 \text{ px} \times 192 \text{ px}$ and use center crops during validation and testing. The cropped images are downsampled by a factor of 3 to a size of $64 \text{ px} \times 64 \text{ px}$.

We use the same low-level mechanism for the image x as with MNIST images but change the encoder and decoder functions to a deeper architecture with 5 scales consisting of 3 blocks of $(\text{LeakyReLU}(0.1) \circ \text{BN}_\theta \circ \text{Conv}_\theta)$ each as well as a linear layer that converts to and from the latent space with 100 dimensions. We directly learn the binary probability of the sex s and use the following invertible transforms to model the age a , brain volume b , and ventricle volume v as

$$a := f_A(\epsilon_A) = (\exp \circ \text{AffineNormalisation} \circ \text{Spline}_\theta)(\epsilon_A), \quad (\text{B.1})$$

$$b := f_B(\epsilon_B; s, a) = (\exp \circ \text{AffineNormalisation} \circ \text{ConditionalAffine}_\theta([s, \hat{a}]))(\epsilon_B), \quad (\text{B.2})$$

$$v := f_V(\epsilon_V; a, b) = (\exp \circ \text{AffineNormalisation} \circ \text{ConditionalAffine}_\theta([\hat{b}, \hat{a}]))(\epsilon_V), \quad (\text{B.3})$$

where the context networks are implemented as a fully-connected network with 8 and 16 hidden units, and a LeakyReLU(0.1) nonlinearity.

B.3 Additional Results

Likewise, we present more detailed analyses of the model trained on UK Biobank brain images and covariates, in terms of modelling the observational distribution and computing various counterfactual queries. (Continued on the next page.)

B.3.1 Association




Figure B.2: Random samples from the model trained on the UK Biobank dataset




Figure B.3: Conditional samples from the model trained on the UK Biobank dataset. Images in each 3×3 block share the same the high-level noise vector, z_X . Each row consistently changes the brain size, whereas each column changes the ventricle volume.






Figure B.4: Original samples and reconstructions from the model trained on the UK Biobank dataset



(a) Age vs. brain volume: $p(a, b | s)$. Here we see differences in head size across biological sexes (reflected in brain volume), as well as a downward trend in brain volume as age progresses.



(b) Age vs. ventricle volume: $p(a, v | b \in \cdot)$. As expected from the literature [52], we observe a consistent increase in ventricle volume with age, in addition to a proportionality relationship with the overall brain volume.

Figure B.5: Densities for the true data (KDE) and for the learned model. The overall trends and interactions present in the true data distribution seem faithfully captured by the model.

B.3.2 Counterfactual




Figure B.6: Original samples and counterfactuals from the model trained on the UK Biobank dataset. The first column shows the original image and true values of the non-imaging data. The even rows show the difference maps between the original image and the corresponding counterfactual image.

C Discrete counterfactuals

As mentioned in the main text, the DSCM framework supports not only low- and high-dimensional continuous data, but also discrete variables. In particular, discrete mechanisms with a Gumbel–max parametrisation have been shown to lead to counterfactuals satisfying desirable properties [57]. For example, they are invariant to category permutations and are stable, such that increasing the odds only of the observed outcome cannot produce a different counterfactual outcome. More computational details and properties of the Gumbel distribution are found in Maddison and Tarlow [58].

Consider a discrete random variable over K categories, y , with a conditional likelihood described by logits λ , assumed to be a function g_Y of its parents, $\mathbf{pa}_Y$:

$$P(y = k | \mathbf{pa}_Y) = \frac{e^{\lambda_k}}{\sum_{l=1}^K e^{\lambda_l}}, \quad \lambda = g_Y(\mathbf{pa}_Y). \quad (\text{C.1})$$

Under the Gumbel–max parametrisation, the mechanism generating y can be described as

$$y := f_Y(\epsilon_Y; \mathbf{pa}_Y) = \arg \max_{1 \leq l \leq K} (\epsilon_Y^l + \lambda_l), \quad \epsilon_Y^l \sim \text{Gumbel}(0, 1). \quad (\text{C.2})$$

Samples from the $\text{Gumbel}(0, 1)$ distribution can be generated by computing $-\log(-\log U)$, where $U \sim \text{Unif}(0, 1)$.

The Gumbel distribution has certain special properties [58] that enable tractable abduction. Given that we observed $y = k$, samples can be generated from the exact posterior $P(\epsilon_Y | y = k, \mathbf{pa}_Y)$:

$$\begin{aligned} \epsilon_Y^k &= G_k + \log \sum_l e^{\lambda_l} - \lambda_k, & G_k &\sim \text{Gumbel}(0, 1), \\ \epsilon_Y^l &= -\log(e^{-G_l - \lambda_l} + e^{-\epsilon_Y^k - \lambda_k}) - \lambda_l, & G_l &\sim \text{Gumbel}(0, 1), \quad \forall l \neq k. \end{aligned} \quad (\text{C.3})$$

Finally, given an upstream counterfactual intervention such that $\tilde{\lambda} = \tilde{g}_Y(\tilde{\mathbf{pa}}_Y)$, the counterfactual outcome for y can be determined simply as

$$y = f_Y(\epsilon_Y; \tilde{\mathbf{pa}}_Y) = \arg \max_{1 \leq l \leq K} (\epsilon_Y^l + \tilde{\lambda}_l). \quad (\text{C.4})$$

Note that this entire derivation applies to a truly discrete variable, without the need for continuous relaxations as commonly used in deep generative models [25, 26], as the likelihood is given in closed form and no gradients of expectations are necessary.