

Mitigating attribute amplification in counterfactual image generation

Tian Xia, Mélanie Roschewitz, Fabio De Sousa Ribeiro, Charles Jones, and Ben Glocker

Department of Computing, Imperial College London, London, UK
`t.xia@imperial.ac.uk`

Abstract. Causal generative modelling is gaining interest in medical imaging due to its ability to answer interventional and counterfactual queries. Most work focuses on generating counterfactual images that look plausible, using auxiliary classifiers to enforce *effectiveness* of simulated interventions. We investigate pitfalls in this approach, discovering the issue of *attribute amplification*, where unrelated attributes are spuriously affected during interventions, leading to biases across protected characteristics and disease status. We show that attribute amplification is caused by the use of hard labels in the counterfactual training process and propose *soft counterfactual fine-tuning* to mitigate this issue. Our method substantially reduces the amplification effect while maintaining effectiveness of generated images, demonstrated on a large chest X-ray dataset. Our work makes an important advancement towards more faithful and unbiased causal modelling in medical imaging. Code available at <https://github.com/biomed-mira/attribute-amplification>.

1 Introduction

Scientific investigation has always been driven by causal questions such as: “*What is the effect of treatment X on disease Y ?*”. In medical imaging, we may ask “*What would this patient’s image look like if they had no disease?*”. These causal questions cannot be answered with statistical tools alone, but require a mathematical framework that allows such questions to be answered from data. Causal models describe (assumed) causal mechanisms of a system [18], in which the causal relationships between variables are directed from cause to effect, and change in a cause would result in change in its effect, but not the other way around. Causal models allow us to analyse interactions of variables within our environment (interventions) and hypothetical alternative worlds (counterfactuals). The ability to reason about cause-and-effect relationships has gained significant interest [22, 21, 34]. Efforts have been made to combine causality and deep learning models [2, 22], but few works have attempted to satisfy all three rungs of Pearl’s causal ladder [18]: (i) association, (ii) intervention and (iii) counterfactual. Notable works include Deep Structural Causal Models (DSCMs) [17, 3] and Neural Causal Models (NCMs) [29, 30]. Our work builds upon DSCMs used for producing high-fidelity counterfactual images for real-world data [3]. Ribeiro

et al. [3] proposed to train the generative causal model using a hierarchical variational auto-encoder (HVAE) conditioned on the assumed causal parents. Solely relying on standard likelihood-based training, however, was found to lead to suboptimal axiomatic *effectiveness* [16], where conditioning on intervened upon parents may be ignored by the forward model post-abduction. To encourage the model to obey (counterfactual) conditioning, the authors proposed an additional training step called *counterfactual fine-tuning*, whereby pretrained parent predictors are used to help fine-tune the HVAE such that conditioning on counterfactual parents results in semantically meaningful changes.

We analyse the quality of counterfactuals generated by DSCMs under the counterfactual fine-tuning setup. We discover that while DSCMs produce plausible counterfactual images, these may exhibit amplified attributes that were not intervened upon. For example, when intervening on biological sex for a healthy male patient, a generated female counterfactual may appear healthier than the real image. We term this phenomenon *attribute amplification*. Amplification of unrelated attributes needs to be addressed as it is in conflict with the assumed causal graph and can cause distribution shift and introduce harmful bias. Attribute amplification can lead to the encoding of spurious correlations in the generated images between protected characteristics and disease. We find that attribute amplification occurs in the counterfactual fine-tuning step. To mitigate this issue, we propose a simple yet effective strategy, *soft counterfactual fine-tuning*, where we replace hard labels with inferred *soft* labels during counterfactual fine-tuning. We show through a series of experiments that soft counterfactual fine-tuning effectively mitigates attribute amplification in counterfactual inference models, whilst retaining the ability of high-fidelity image synthesis. Some visual examples of generated CFs are presented in Fig 1.

2 Counterfactual image generation

Various works have used generative models such as VAEs [11], GANs [7], normalizing flows [19] and diffusion models [23,8,25] for causal effect estimation [14,12,26], causal discovery [31,20,5], and other tasks modelling conditional [27,15,24,4] and interventional distributions [13,10,29,33]. Very few works [30,17] satisfied all three rungs of Pear’s ladder of causation [18,1]. Our work builds upon DSCMs first introduced by Pawlowski et al. [17], and recently improved by Ribeiro et al. [3], leveraging hierarchical variational autoencoders (HVAE) to generate high-quality high-resolution images. In the following, we summarise the main components of DSCMs as introduced by Ribeiro et al., for more details refer to [3].

Structural Causal Models [18] are defined by a triplet $\langle U, A, F \rangle$, where $U = \{u_i\}_{i=1}^K$ are a set of *exogenous variables*, $A = \{a_i\}_{i=1}^K$ a set of *endogenous variables*, and $F = \{f_i\}_{i=1}^K$ a set of functions such that $a_k := f_k(\mathbf{pa}_k, u_k)$, where $\mathbf{pa}_k \subseteq A \setminus a_k$ are called direct causes or *parents* of a_k . SCMs allow us to perform interventions by substituting one or multiple parents, denoted by the *do*-operator. The estimation of counterfactuals follows three steps: (i) *abduction*:

infer exogenous noise given observed data; (ii) *action*: perform an intervention $do(a_k := c)$; (iii) *prediction*: infer counterfactuals via the modified model.

Deep Structural Causal Models were first introduced in [17] and recently improved in [3] for high-resolution image generation. Let $\mathbf{x}$ be an image and $\{a_1, \dots, a_{K-1}\} \supseteq \mathbf{pa}_{\mathbf{x}}$ *ancestors* of $\mathbf{x}$. For each low-dimensional attribute, a conditional normalizing flow is used as its (invertible) mechanism $a_k = f_k(u_k; \mathbf{pa}_k)$, such that abduction is tractable and explicit. For high-dimensional structured variables like images, the mechanism is implemented by a HVAE. To generate a counterfactual, we first approximately infer the exogenous noise $\mathbf{z} \sim q_\phi(\mathbf{z} | \mathbf{x}, \mathbf{pa}_{\mathbf{x}})$, where q_ϕ is the encoder of the HVAE. Similarly, we can explicitly infer the exogenous noises for attributes via: $u_k = f^{-1}(a_k; \mathbf{pa}_k)$. We can then perform an intervention $do(a_i := c)$ by e.g. setting a_i to some target c . Note, we may intervene on multiple a_i at the same time. We compute counterfactuals of attributes using abducted noise u_k , and thus obtain $\widetilde{\mathbf{pa}}_{\mathbf{x}}$, the counterfactuals of parents of $\mathbf{x}$. Finally, we generate the counterfactual image $\widetilde{\mathbf{x}} = g_\theta(\mathbf{z}, \widetilde{\mathbf{pa}}_{\mathbf{x}})$. Ribeiro et al [3] noticed that using only likelihood training for the HVAE, some parents may be ignored during inference, effectively leading to ‘ignored counterfactual conditioning’, i.e. $\widetilde{\mathbf{x}}$ does not obey $\widetilde{\mathbf{pa}}_{\mathbf{x}}$. To mitigate this issue the authors introduced *counterfactual fine-tuning* using hard labels for parents. *Hard-CFT* leverages a pre-trained predictor $q_\psi(\widetilde{\mathbf{pa}}_{\mathbf{x}} | \widetilde{\mathbf{x}})$, and optimises the pre-trained HVAE weights $\{\theta, \phi\}$ to maximise $\log q_\psi(\widetilde{\mathbf{pa}}_{\mathbf{x}} | \widetilde{\mathbf{x}})$ with ψ fixed. This encourages the DSCM to produce effective counterfactuals which obey counterfactual conditioning.

3 The attribute amplification problem

Datasets and experimental setup. We use chest X-ray images of the MIMIC-CXR dataset [9] for our experiments. We follow the dataset splits and filtering of [3,6], and focus on the disease label of pleural effusion. The resulting dataset contained subjects with and without pleural effusion, consisting of 62,336 images for training, 9,968 for validation and 30,535 for testing. All images were resized to 224×224 resolution. Following [3], we consider four attributes: (self-reported) race, biological sex, disease status, and age for the underlying causal graph (see Fig. A1). Note, that disease, sex, and race are assumed to be independent.

Evaluating counterfactual effectiveness. The ability to generate realistic counterfactual images can be measured via *effectiveness* [3,16], which assesses whether counterfactuals (CFs) can fool attribute predictors. The attribute predictors were trained on real data with ground truth labels in a supervised manner. Here, we measure the predictive performance in terms of AUC. Results are summarised in Table 1, where we report (i) the baseline test performance of each attribute predictor on real test images and (ii) the performance when testing on different types of CFs. Ideally, we would find similar performance across real and counterfactual images. Without counterfactual fine-tuning, CFs yield much lower AUC on intervened attributes, highlighting the importance of CFT.

Table 1: Quantitative evaluation of *effectiveness* of CFs of different types. For *intervened* and *unintervened* attributes, AUC % (test on CFs) is reported (with change in brackets) compared to the baseline (test on real images, first row).

INTERVENTION	CF TYPE	EFFECTIVENESS / AMPLIFICATION		
		RACE	SEX	DISEASE
None	Real images	93.8	99.6	94.2
<i>do</i> (race)	No CFT	68.9 (↓ 24.9)	99.6 (· 0.0)	94.2 (· 0.0)
	Hard-CFT [3]	99.0 (↑ 5.2)	99.7 (↑ 0.1)	96.4 (↑ 2.2)
	Soft-CFT	98.7 (↑ 4.9)	99.5 (↓ 0.1)	94.3 (↑ 0.1)
<i>do</i> (sex)	No CFT	92.4 (↓ 1.4)	92.8 (↓ 6.8)	94.0 (↓ 0.2)
	Hard-CFT [3]	97.1 (↑ 3.3)	99.8 (↑ 0.2)	97.8 (↑ 3.6)
	Soft-CFT	93.3 (↓ 0.5)	99.7 (↑ 0.1)	94.5 (↑ 0.3)
<i>do</i> (disease)	No CFT	93.7 (↓ 0.1)	99.2 (↓ 0.4)	70.6 (↓ 23.6)
	Hard-CFT [3]	97.3 (↑ 3.5)	99.7 (↑ 0.1)	97.9 (↑ 3.7)
	Soft-CFT	93.8 (· 0.0)	99.6 (· 0.0)	98.1 (↑ 3.9)

The attribute amplification problem. From the results in Table 1, we identify the key problem of *attribute amplification*, where attributes that should remain unchanged during interventions are amplified in the generated images, indicated by an increase in AUC beyond the baseline for the attribute predictor. When using Hard-CFT as proposed in [3], counterfactual images yield consistently higher AUC compared to real images for *unintervened* attributes. These attributes seem more strongly encoded in the generated images than in the real images, leading to undesired dataset shift and potential bias. For example, if x is *healthy* and *male*, when $do(sex := female)$, $\tilde{x}$ becomes *healthier* than x , although in the assumed causal graph *disease* is not a descendant of *sex* and thus $do(sex)$ should not affect *disease*. We observe an increase in AUC for the disease attribute by 3.6% after intervening on sex. Similar, the AUC for race increases by 3.5% after intervening on disease. These undesirable side effects question the faithfulness of generated counterfactual images when using the Hard-CFT strategy.

Why is attribute amplification a problem? The effect on *unintervened* attributes violates the assumed causal graph, affecting the causal relationships between variables. This violation is illustrated in Fig. A1. Attribute amplification not only affects the quality of counterfactuals but may introduce harmful, spurious correlations between protected characteristics and disease status. Note these spurious correlations are caused by the counterfactual training process, which is different from the spurious correlations discussed in shortcut learning [28,32] that mainly come from data issues. This could have negative consequences when using generated images in downstream applications such as counterfactual data augmentation or counterfactual explainability.

Algorithm 1 Soft counterfactual fine-tuning (Soft-CFT)**Input:** training data $\{\mathbf{x}, \mathbf{pa}_{\mathbf{x}}\}$; frozen predictor q_{ψ} ; DSCM q_{ϕ}, g_{θ} .**Compute counterfactuals:**

1. Compute $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{pa}}_{\mathbf{x}}$ using q_{ϕ} and g_{θ} upon intervention $\mathbf{do}(\cdot)$.

Fine-tune DSCM:

1. Identify attributes in $\tilde{\mathbf{pa}}_{\mathbf{x}}$ that are not intervened on or non-descendants of intervened attributes, denoted as $\mathbf{A}^*$.
2. For attributes $a_k \in \mathbf{A}^*$ use soft labels as targets. With l a classification loss function: $\mathcal{L}_{soft} := l(q_{\psi}(a_k | \mathbf{x}), q_{\psi}(a_k | \tilde{\mathbf{x}}))$.
4. For $a_k \notin \mathbf{A}^*$ use hard labels as targets i.e. $\mathcal{L}_{interv} := l(a_k, q_{\psi}(a_k | \tilde{\mathbf{x}}))$.
5. Fine-tune q_{ϕ}, g_{θ} by minimising $\mathcal{L}_{soft} + \mathcal{L}_{interv}$.

Why does Hard-CFT cause attribute amplification? In Hard-CFT, $\tilde{\mathbf{x}}$ is encouraged to maximise the performance of a pretrained attribute predictor $q_{\psi}(\tilde{a}_k | \tilde{\mathbf{x}})$, $\forall \tilde{a}_k \in \tilde{\mathbf{pa}}_{\mathbf{x}}$, regardless of whether $\tilde{a}_k$ should actually be affected by the intervention. Suppose we have sample $\mathbf{x}$ with mild pleural effusion ($d = 1$), and that the attribute predictor outputs a disease probability of $q_{\psi}(d | \mathbf{x}) = 0.65$ for the real image. When we intervene on race, as both race and disease are parents of $\mathbf{x}$, Hard-CFT will not only maximise the probability of the race predictor but will also maximise the predicted disease probability $q_{\psi}(\tilde{d} | \tilde{\mathbf{x}})$ by optimizing the HVAE weights. As a consequence, the model will amplify pleural effusion related features in the generated race counterfactual image.

4 Mitigating attribute amplification with soft labels

We discovered that attribute amplification occurs during counterfactual fine-tuning, which, however, is necessary to increase effectiveness on intervened attributes. Amplification of unintervened attributes is caused by the use of hard labels in the CFT step, making all attributes more extreme than they may have been in the original image. Ideally, we only change attributes that are affected by the intervention, and leave other attributes unchanged. To this end, we propose soft counterfactual fine-tuning (Soft-CFT) to mitigate attribute amplification. In Soft-CFT, we treat attributes differently during fine-tuning depending on whether they are being intervened on or not. For intervened attributes, we fine-tune the model with the hard labels to ensure the model obeys conditioning. For unintervened attributes, we use inferred soft labels. We first pass the real image through the trained attribute predictors to obtain the predicted probability for each non-intervened attribute. These predicted probabilities then form the targets during counterfactual fine-tuning. This encourages the network to not change the encoding of unintervened attributes, and the resulting counterfactuals should have the same predicted probability on the unintervened attributes as the original real images. The Soft-CFT process is summarised in Algorithm 1.




Fig. 1: Generated CFs with (a) Hard-CFT and (b) Soft-CFT. First rows show original image $\mathbf{x}$ and CFs $\tilde{\mathbf{x}}$; second rows show direct effect of CFs, i.e. $\tilde{\mathbf{x}} - \mathbf{x}$.

4.1 Effect of Soft-CFT on effectiveness and attribute amplification

Table 1 compares effectiveness and amplification of Soft-CFT and Hard-CFT. With Soft-CFT, AUCs for unintervened attributes are closer to the baseline, substantially reducing the amplification effect while preserving effectiveness on intervened attributes. When intervening on sex, disease amplification is reduced to 0.3% compared to 3.6% for Hard-CFT. Amplification on race is completely removed (0.0) compared to 3.5% when intervening on disease. Visual examples in Fig. 1 show that both Hard-CFT and Soft-CFT generate plausible CFs. However, with Hard-CFT, the generated images exhibit more global and stronger direct effects. With Soft-CFT, images show more targeted, localised changes aligned with the intervened variable. See Fig. A2 for additional visual examples.

4.2 Utility of counterfactual images in downstream applications

We evaluate the utility of CFs by training multiple attribute predictors on the generated counterfactuals. We train attribute predictors on each set of counterfactuals (race, sex, disease), and evaluate on real images from the test set. To make a fair comparison, we randomly generate one CF per subject per attribute, such that we train on the same number of CFs as real training samples. Note, this is the opposite in Table 1 where attribute predictors are trained on real images, and evaluated on counterfactuals. To compare the performance of predictors trained on counterfactuals only, we compare to the baseline where predictors are trained and tested on real data only. Without CFT, AUCs are generally low on intervened attributes due to the low effectiveness (see Table 1). We find consistent and substantial improvements for Soft-CFT over Hard-CFT. Predictors trained on counterfactuals generated with our Soft-CFT strategy are much closer to baseline performance on both intervened and unintervened attributes for all three sets of race, sex, and disease counterfactuals. Hard-CFT, on the other hand, appears to introduce undesired distribution shift. Note the decrease in performance on sex counterfactuals even for the sex predictor, despite their high effectiveness observed in Table 1. Here, attribute amplification

Table 2: Performance of predictors trained **only** on CFs with **interven**ed and **un-interven**ed attributes, evaluated on real images. Comparing different CF training strategies. AUC % performance is compared to a baseline trained on real images.

TRAINING SET	CF TYPE	ATTRIBUTE PREDICTION TASK		
		RACE	SEX	DISEASE
Original	Real images	93.8	99.6	94.2
Race CFs	No CFT	56.9 (↓ 36.9)	99.5 (↓ 0.1)	93.9 (↓ 0.3)
	Hard-CFT	81.6 (↓ 12.2)	98.3 (↓ 1.3)	92.2 (↓ 2.0)
	Soft-CFT	81.2 (↓ 12.6)	99.6 (↓ 0.0)	93.8 (↓ 0.4)
Sex CFs	No CFT	82.1 (↓ 11.7)	82.2 (↓ 17.4)	92.3 (↓ 1.9)
	Hard-CFT	75.4 (↓ 18.4)	79.6 (↓ 20.0)	89.8 (↓ 4.4)
	Soft-CFT	92.4 (↓ 1.4)	90.0 (↓ 9.6)	94.0 (↓ 0.2)
Disease CFs	No CFT	79.7 (↓ 14.1)	98.1 (↓ 1.5)	30.4 (↓ 63.8)
	Hard-CFT	77.6 (↓ 16.2)	95.2 (↓ 4.4)	88.8 (↓ 5.4)
	Soft-CFT	91.8 (↓ 2.0)	99.5 (↓ 0.1)	90.0 (↓ 4.2)

of Hard-CFT results in images being too extreme. While they are correctly classified with predictors trained on real images, these images are not suitable for training predictors evaluated on real images. The results in Table 2 together with Table 1 indicate that Soft-CFT both improves effectiveness and reduces attribute amplification yielding more faithful counterfactual images which is important for their utility in downstream applications.

4.3 Assessing distribution shift of counterfactual images

To gain further insights into the faithfulness of counterfactual images under different training strategies, we inspect the latent space of image embeddings extracted from attribute predictors trained on real data. Here, we use a publicly available pre-trained multi-task model from [6]. This model was trained on the same splits of MIMIC-CXR as used in our work. The model was previously used to visualise subgroup differences across race, sex and disease. Applying principal component analysis (PCA) to the high-dimensional image embeddings showed that each independent attribute is encoded in a different PCA mode, with the first mode encoding disease, the second mode encoding sex, and the third mode encoding race differences. We leverage this feature inspection approach to compare image embeddings from real and counterfactual images. Ideally, we should be unable to see differences in the data distribution between real and corresponding generated images, if the counterfactuals are faithful.

We randomly select 1000 samples from the test set, generate counterfactuals and inspect the distributions in the corresponding PCA modes. Fig. 2 focuses on distribution shift in the race attribute under interventions on disease and sex.




Fig. 2: Marginal distribution of image embeddings across PCA mode 3 of a multi-task model. This mode encodes changes in the race attribute. We plot distributions for subgroups of real data alongside distributions of counterfactual images when intervening on disease and sex attributes. When training with Hard-CFT (left) there is a clear distribution shift between real (blue) and counterfactual images (red). This shift is removed for both interventions when using our proposed Soft-CFT (right). These results suggest that Soft-CFT successfully mitigates attribute amplification and generates more faithful counterfactual images.

We observe that with Hard-CFT, there is a clear distribution shift between the real images from the subgroup ‘White’ and the counterfactual subgroups ‘White *do*(disease)’ and ‘White *do*(sex)’. In contrast, the images generated with Soft-CFT align much better with real data distribution, suggesting that Soft-CFT mitigates attribute amplification and yields more faithful counterfactual images that are more similar to real data. Additional plots inspecting other combinations of intervened and unintervened attributes across different subgroups are provided in Fig. A3, confirming that Soft-CFT improves counterfactual image generation.

5 Conclusion

In this paper, we discover the important issue of *attribute amplification* in counterfactual image generation with DSCMs when using the previously proposed counterfactual fine-tuning step with hard labels [3]. Attribute amplification violates the assumed causal graph, introducing distribution shift, and resulting in potentially harmful spurious correlations. For example, we observed how interventions on independent attributes such as race and sex can cause distribution shift in disease status. An intervened image may appear healthier or more diseased than the original real images. Such undesired effects could be harmful when counterfactual images are used for data augmentation in downstream applications. We mitigate this with Soft-CFT, using inferred probabilities for unintervened attributes. Our experiments on a large chest X-ray dataset with results on attribute prediction performance, together with a detailed feature inspection to

analyse distribution shift across subgroups, suggest that Soft-CFT successfully mitigates attribute amplification with the ability to generate effective and more faithful counterfactual images. However, we still observe discrepancies between the attribute prediction performance when training on real vs counterfactual images. Closing this gap remains an open problem for future work. Visual assessment of CFs by clinical expert will be helpful in future work.

The proposed method is not only applicable to DSCMs but may improve other causal generative models that use classifier guidance during training. Thus, our work makes an important advancement towards more faithful and hopefully unbiased causal modelling in medical imaging.

Acknowledgement. This work was supported via the UKRI AI programme and EPSRC as part of the AI Hub for Causality in Healthcare AI (grant number EP/Y028856/1). T.X., F.R. and B.G. received funding from the ERC under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 757173). B.G. also received support from the Royal Academy of Engineering as part of his Kheiron/RAEng Research Chair in Safe Deployment of Medical Imaging AI. M.R. is funded through an Imperial College London President’s PhD Scholarship. C.J. is supported by Microsoft Research, EPSRC and The Alan Turing Institute through a Microsoft PhD Scholarship and Turing PhD enrichment award.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bareinboim, E., Correa, J.D., Ibeling, D., Icard, T.: On Pearl’s Hierarchy and the Foundations of Causal Inference, p. 507–556. Association for Computing Machinery, New York, NY, USA, 1 edn. (2022), <https://doi.org/10.1145/3501714.3501743>
2. Bengio, Y., Courville, A., Vincent, P.: Representation learning: a review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence* **35**(8), 1798–1828 (August 2013). <https://doi.org/10.1109/tpami.2013.50>
3. De Sousa Ribeiro, F., Xia, T., Monteiro, M., Pawlowski, N., Glocker, B.: High fidelity image counterfactuals with probabilistic causal models. *ICML* (2023)
4. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems* **34**, 8780–8794 (2021)
5. Geffner, T., Antoran, J., Foster, A., Gong, W., Ma, C., Kiciman, E., Sharma, A., Lamb, A., Kukla, M., Pawlowski, N., et al.: Deep end-to-end causal inference. *arXiv preprint arXiv:2202.02195* (2022)
6. Glocker, B., Jones, C., Bernhardt, M., Winzeck, S.: Algorithmic encoding of protected characteristics in chest x-ray disease detection models. *Ebiomedicine* **89** (2023)
7. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. *Communications of the ACM* **63**(11), 139–144 (2020)

8. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* **33**, 6840–6851 (2020)
9. Johnson, A.E., Pollard, T.J., Berkowitz, S.J., Greenbaum, N.R., Lungren, M.P., Deng, C.y., Mark, R.G., Horng, S.: Mimic-cxr, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific data* **6**(1), 1–8 (2019)
10. Ke, N.R., Bilaniuk, O., Goyal, A., Bauer, S., Larochelle, H., Schölkopf, B., Mozer, M.C., Pal, C., Bengio, Y.: Learning neural causal models from unknown interventions. *arXiv preprint arXiv:1910.01075* (2019)
11. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013)
12. Kocaoglu, M., Snyder, C., Dimakis, A.G., Vishwanath, S.: Causalgan: Learning causal implicit generative models with adversarial training. *arXiv preprint arXiv:1709.02023* (2017)
13. Kocaoglu, M., Snyder, C., Dimakis, A.G., Vishwanath, S.: CausalGAN: Learning causal implicit generative models with adversarial training. In: *International Conference on Learning Representations* (2018), <https://openreview.net/forum?id=BJE-4xWOW>
14. Louizos, C., Shalit, U., Mooij, J.M., Sontag, D., Zemel, R., Welling, M.: Causal effect inference with deep latent-variable models. *Advances in neural information processing systems* **30** (2017)
15. Mirza, M., Osindero, S.: Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (2014)
16. Monteiro, M., Ribeiro, F.D.S., Pawlowski, N., Castro, D.C., Glocker, B.: Measuring axiomatic soundness of counterfactual image models. In: *The Eleventh International Conference on Learning Representations* (2023), <https://openreview.net/forum?id=1Z0UQQvwI3q>
17. Pawlowski, N., Coelho de Castro, D., Glocker, B.: Deep structural causal models for tractable counterfactual inference. *Advances in Neural Information Processing Systems* **33**, 857–869 (2020)
18. Pearl, J.: *Causality*. Cambridge university press (2009)
19. Rezende, D., Mohamed, S.: Variational inference with normalizing flows. In: *International conference on machine learning*. pp. 1530–1538. PMLR (2015)
20. Sanchez, P., Liu, X., O’Neil, A.Q., Tsaftaris, S.A.: Diffusion models for causal discovery via topological ordering. *arXiv preprint arXiv:2210.06201* (2022)
21. Schölkopf, B.: *Causality for Machine Learning*, p. 765–804. Association for Computing Machinery, New York, NY, USA, 1 edn. (2022), <https://doi.org/10.1145/3501714.3501755>
22. Schölkopf, B., Locatello, F., Bauer, S., Ke, N.R., Kalchbrenner, N., Goyal, A., Bengio, Y.: Toward causal representation learning. *Proceedings of the IEEE* **109**(5), 612–634 (2021)
23. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: *International Conference on Machine Learning*. pp. 2256–2265. PMLR (2015)
24. Sohn, K., Lee, H., Yan, X.: Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems* **28** (2015)
25. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: *International Conference on Learning Representations* (2021), <https://openreview.net/forum?id=PXTIG12RRHS>

26. Tran, D., Blei, D.M.: Implicit causal models for genome-wide association studies. arXiv preprint arXiv:1710.10742 (2017)
27. Trippe, B.L., Turner, R.E.: Conditional density estimation with bayesian normalising flows. arXiv preprint arXiv:1802.04908 (2018)
28. Weng, N., Pegios, P., Feragen, A., Petersen, E., Bigdeli, S.: Fast diffusion-based counterfactuals for shortcut removal and generation. arXiv preprint arXiv:2312.14223 (2023)
29. Xia, K., Lee, K.Z., Bengio, Y., Bareinboim, E.: The causal-neural connection: Expressiveness, learnability, and inference. *Advances in Neural Information Processing Systems* **34**, 10823–10836 (2021)
30. Xia, K.M., Pan, Y., Bareinboim, E.: Neural causal models for counterfactual identification and estimation. In: *The Eleventh International Conference on Learning Representations* (2023), <https://openreview.net/forum?id=vouQcZS8KfW>
31. Yang, M., Liu, F., Chen, Z., Shen, X., Hao, J., Wang, J.: Causalvae: Disentangled representation learning via neural structural causal models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9593–9602 (2021)
32. Yang, W., Kirichenko, P., Goldblum, M., Wilson, A.G.: Chroma-vae: Mitigating shortcut learning with generative classifiers. *Advances in Neural Information Processing Systems* **35**, 20351–20365 (2022)
33. Zečević, M., Dhimi, D.S., Veličković, P., Kersting, K.: Relating graph neural networks to structural causal models. arXiv preprint arXiv:2109.04173 (2021)
34. Zhang, J., Jennings, J., Zhang, C., Ma, C.: Towards causal foundation model: on duality between causal inference and attention. arXiv preprint arXiv:2310.00809 (2023)