

ROBUST DOMAIN GENERALISATION WITH CAUSAL INVARIANT BAYESIAN NEURAL NETWORKS

Anonymous authors

Paper under double-blind review

ABSTRACT

Deep neural networks can obtain impressive performance on various tasks under the assumption that their training domain is identical to their target domain. Performance can drop dramatically when this assumption does not hold. One explanation for this discrepancy is the presence of spurious domain-specific correlations in the training data that the network exploits. Causal mechanisms, in the other hand, can be made invariant under distribution changes as they allow disentangling the factors of distribution underlying the data generation. Yet, learning causal mechanisms to improve out-of-distribution generalisation remains an under-explored area. We propose a Bayesian neural architecture that disentangles the learning of the data distribution from the inference process mechanisms. We show theoretically and experimentally that our model approximates reasoning under causal interventions. We demonstrate the performance of our method, outperforming point estimate-counterparts, on out-of-distribution image recognition tasks where the data distribution acts as strong adversarial confounders.

1 INTRODUCTION

The training of deep neural networks commonly relies on the assumption that the distribution of the training data is representative of the distribution at inference. Despite being widely adopted, this assumption has been heavily criticised as it is often challenged in practice (Langford, 2005; Ben-David et al., 2006; Albuquerque et al., 2019; Arjovsky et al., 2019; Jalaldoust & Bareinboim, 2024). Indeed, despite tremendous progress on many vision tasks over recent years, deep neural networks face limitations and reduced performance on tasks requiring the model to shift from its training distribution at test time (Mao et al., 2022).

Causality theory under the do-calculus framework (Huang & Valtorta, 2006; Pearl, 2009) provides tools to explain these limitations. Neural networks extract correlation patterns but do not possess knowledge on the cause-effect relationships from the data; causal links and potentially spurious correlations are learned alike. However, the two relationships fundamentally differ as correlations can be non-causal: e.g. a correlation between X and Y can be explained by a causal link $X \leftarrow Z \rightarrow Y$. This type of correlation can be specific to the distribution if Y depends on the domain, e.g. *has tusks and a prehensile trunk* \leftarrow *African bush elephant* \rightarrow *is in a savannah environment*. A system only learning the correlation between X and Y may fail to recognise an elephant in a different environment. However, direct causal links make these relationships explicit and are therefore more robust to changes in the distribution (Schölkopf et al., 2021). In particular, the *Independent Causal Mechanisms* principle states that causal relationships are only sparsely connected and altering one should not modify the others (Peters et al., 2017; Schölkopf et al., 2021), allowing some learned mechanisms to be *transported* to new environments. Indeed, changes in the distribution can be modelled via the notion of *transportability* of causal mechanisms (Bareinboim & Pearl, 2013; Jalaldoust & Bareinboim, 2024). As such, transportable causal effects are equivalent to domain-invariant features. These principles motivate the search for factorised neural models composed of independent modules representing either domain-specific or domain-invariant conditional distributions.

Bayesian neural networks (BNNs) (MacKay, 1992; Blundell et al., 2015b; Magris & Iosifidis, 2023) are another line of work attempting to model robust representations of mechanisms by learning the distribution W of a parametric function $y = f_w(x)$, $w \sim W$ instead of a single point estimate w . BNNs are particular suited to express uncertainty in their response and reduce overconfidence. By

054
055
056
057
058
059
060
061
062
063
064
065
066
067
068
069
070
071
072
073
074
075
076
077
078
079
080
081
082
083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099
100
101
102
103
104
105
106
107

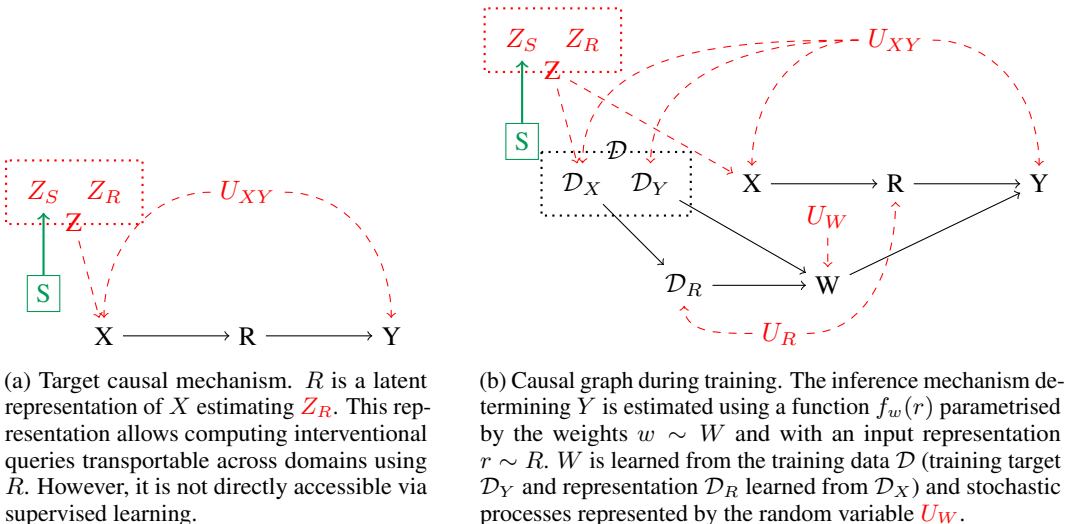


Figure 1: Target causal representation and actual causal graph during training. X is the input image and Y is the output class. Z is the variable representing the factors of distribution generating the input X , it is composed of domain-specific factors Z_S and robust domain-invariant factors Z_R . U_{XY} represents the shared factors of variations between X and Y . The variables in red are unobserved. S is a selection variable determining the current domain. The target causal graph is different from the one used in supervised deep learning as it omits the influence of the datasets $\mathcal{D} = \{\mathcal{D}_X, \mathcal{D}_Y\}$. The datasets depend on the same factors as X and Y and add spurious correlations that are not captured when using the target representation.

differentiating the modelling of uncertainty in the data distribution (aleatoric uncertainty) from the uncertainty in the process distribution (epistemic uncertainty), BNNs can better capture the latter (Kendall & Gal, 2017). However, BNNs are not fundamentally more robust against distribution shifts than their point-estimate counterparts (Izmailov et al., 2021). This problem can be explained by the intractability in the optimisation of the posterior (Magris & Iosifidis, 2023). Motivated by the use of partially stochastic networks (Sharma et al., 2023), we argue that BNNs should also be integrated in a wider framework estimating causal mechanisms.

We propose a Causal-Invariant Bayesian (CIB) neural network taking advantage of the causal structure of the supervised learning process to differentiate domain-specific and domain-invariant mechanisms. Compared to previous work, we integrate a more realistic causal graph that takes into account the Bayesian model update during learning. We use variational inference and partially-stochastic Bayesian neural networks to model the causal paths in an interventional setting. We verify theoretically and experimentally the applicability of our method. In particular, we perform tests on standard out-of-distribution (o.o.d) image recognition tasks. Our contributions are as follows:

- We reformulate the Bayesian inference problem to include causal interventions and the supervised learning mechanism and propose a factorised model explicitly modelling domain-invariant mechanisms in an unsupervised fashion,
- We propose an architecture wrapped around a neural network inspired by this factorised model and test it on challenging tasks requiring domain shifts,
- We show that adding our model can increase the i.i.d and o.o.d performance of the underlying neural network and reduce overfitting,
- Additionally, we find with our model that adding unsupervised contextual information with a Bayesian classifier to a standard ResNet improves its robustness and training stability, even with a small number of context and Bayesian weight samples.

Our code is available at this anonymous repository: <https://anonymous.4open.science/r/cibnn-1B0B>.

2 BACKGROUND AND RELATED WORK

Causal Inference and Transportability Causal inference methods aim to estimate the result of causal queries. Answering these queries require an inference model to have a partial knowledge of the causal mechanisms underlying the studied system. The causal knowledge is typically divided into three layers: *observational* ($I1$), *interventional* ($I2$) and *counterfactual* ($I3$) (Pearl, 2009). Bareinboim et al. (2022) showed that these layers form a non-collapsing hierarchy, i.e. information at level i is needed to answer a question at level i . Richens & Everitt (2024) further showed that an agent must learn a causal model of the world (at level $I2$) to robustly solve a task, e.g. subject to changes of domain. The problem of formally representing causal quantities is addressed by the do-calculus framework, which introduces the $do(\cdot)$ operator corresponding to an intervention on a causal variable (Pearl, 2009). For example, given a treatment T and background covariates X , the probability of getting cured given by C can be expressed with the query $P(C|do(T), X)$. The do-operation implies that the treatment is received in an unbiased manner (i.e. as in a double-blind study). An interventional query as described in the example can be reduced to observations using the rules of do-calculus. We describe them in Appendix A.1. *Transportability* theory furthermore allows representing causal mechanisms across multiple domains (Bareinboim & Pearl, 2013) by differentiating domain-specific and domain-invariant variables in a causal graph. Assuming the domains can be represented with a unified causal graph, domain-specific mechanisms can be modelled by a *selection variable* whose value depends on the current domain. The *S-admissibility* criterion (Jalaldoust & Bareinboim, 2024) states that the value of a causal variable A is invariant when shifting from a domain \mathcal{M}_i to a domain \mathcal{M}_j if it can be *d-separated* from all selection variables S_{ij} (given observations Z), i.e. $A \perp\!\!\!\perp S_{ij}|Z \implies P^i(A|Z) = P^j(A|Z)$. Mao et al. (2022) also integrate causal transportability for vision tasks but do not include the Bayesian learning process. They combine input and contextual information differently and only select context with the same label during training. Conversely, we make the assumption that the context should be diverse and representative of the label distribution. We use label mixup to this end. We also include new regularisation techniques.

Disentanglement A large body of work attempting to disentangle domain-specific and domain-invariant information rely on an autoencoding paradigm and reconstruct the input image as part of their training process (Gabbay et al., 2021; Yang et al., 2021; Zhang et al., 2022). For instance, Zhang et al. (2022) build a latent space regularised to divide domain-invariant *semantic* features from a domain-specific *variation* space. Similarly to our work, the authors add contextual information by feeding different samples to a variation space encoder to improve domain generalisation. However, they mainly rely on a reconstruction loss whereas we argue that this component is not needed for building robust representations. van Steenkiste et al. (2019) argued that having disentangled representations improved systematic generalisation in abstract visual reasoning tasks. The authors also found that a low reconstruction error was not necessary for performance on downstream tasks. Our work differs from the disentanglement literature by detaching from the autoencoding paradigm and adopting a causal Bayesian approach.

Bayesian Neural Networks and Variational Inference Bayesian deep learning methods teach neural networks to simulate Bayesian reasoning when learning new information. This approach allows dealing with uncertainty and has been argued to help mitigate overconfidence and improve robustness in neural networks (Magris & Iosifidis, 2023). Bayesian Neural Networks (BNNs) model the distribution of parametric functions solving a task: the parameters w of a BNN are not directly optimised but sampled from a learned posterior distribution $P(w|\mathcal{D})$. This distribution can be obtained via the *Bayes objective*: $P(w|\mathcal{D}) = \frac{P(\mathcal{D}|w)P(w)}{P(\mathcal{D})}$. However, training BNNs is challenging because the computation of the denominator (called the *evidence*) is often intractable (Izmailov et al., 2021). A standard way to circumvent this issue is to approximate the posterior $P(w|\mathcal{D})$ with a *variational distribution* $q(w)$. This distribution can be estimated by maximising the *Evidence Lower Bound* (ELBO) (Blundell et al., 2015b). This quantity can have many expressions, the one most suitable for our work is shown in Equation 1. It simultaneously maximises the likelihood of the data $\mathcal{D} = \{\mathcal{D}_x, \mathcal{D}_y\}$ given the model w while maintaining the distribution $q(w)$ close to the prior $P(w)$.

$$\text{ELBO}(q) = \mathbb{E}_{q(w)}[\log P(\mathcal{D}_y|\mathcal{D}_x, w)] - \text{KL}(q(w)||P(w)) \quad (1)$$

Izmailov et al. (2021) showed that the posterior distribution of BNNs have a high functional diversity and can outperform their *point-estimate* counterparts on downstream tasks. However, they do not necessarily show a high diversity in the parameters space and strictly following Bayesian posteriors can lead to reduced robustness against distribution shifts. Sharma et al. (2023) further showed that partially-stochastic BNNs outperform networks containing only Bayesian layers. Roy et al. (2024) recently showed that inducing BNNs to assign similar posterior densities to reparametrisations of the same functions (i.e. different weights realising the same function) could improve their ability to fit the data. These findings motivate our work on integrating BNNs as a part of a larger causal neural network.

Mixup Input Mixup (Zhang et al., 2018) and Manifold Mixup (Verma et al., 2019) are regularisation techniques for improving robustness in image classification tasks by interpolating inputs, respectively latent representations. The main principle of mixup consists of combining two inputs (x_i, x_j) and labels (y_i, y_j) together to form a new virtual training example \bar{x} with label \bar{y} : $\bar{x} = \alpha \cdot x_i + (1 - \alpha) \cdot x_j$ and $\bar{y} = \alpha \cdot y_i + (1 - \alpha) \cdot y_j$ (with $\alpha \in [0, 1]$ constant). Manifold Mixup performs a similar operation but interpolates latent representations instead of the inputs. Gendron et al. (2023) argued that mixing could be used to include interventional information into the network’s training. Following these observations, we integrate mixup strategies as part of our training scheme.

3 INTERVENTIONAL BAYESIAN INFERENCE

We propose to learn deep domain-invariant representations that can be leveraged to solve o.o.d tasks. In this section, we show the necessary assumptions needed to learn this representation in a supervised fashion and identify an interventional query that can solve the task. We show that this query can be answered using Bayesian Neural Networks (MacKay, 1992; Blundell et al., 2015b; Magris & Iosifidis, 2023). Section 4 describes our proposed architecture.

Interventional queries The causal graph in Figure 1b shows the causal dependencies between the variables involved in the learning process. The standard inference query $P(Y|x)$ is affected by spurious correlations as they do not distinguish the causal links $X \overset{\text{red}}{\rightsquigarrow} Y$ (i.e. $X \rightarrow \dots \rightarrow Y$) and their common causes $Z \overset{\text{red}}{\leftarrow} U_{XY} \overset{\text{red}}{\rightarrow} Y$. In causality theory, the $do(\cdot)$ operator removes the incoming dependencies of the target variable (Pearl, 2009). Therefore, the interventional query $P(Y|do(x))$ solves the confounding issue (Mao et al., 2022). However, it does not take into account the conditioning on the training data \mathcal{D} . We specify an interventional query that makes the dependency on the data explicit. Note that we simplify the formalism by considering the label Y_{true} and predicted value Y_{pred} with a single variable Y . This is equivalent to having a prediction loss $\mathcal{L}(Y_{\text{pred}}, Y_{\text{true}}) = 0$.

Conditional distributions independent from the domain selection variable S are transportable across domains (Jalaldoust & Bareinboim, 2024). This independence can be obtained using the $do(\cdot)$ operation by removing the outgoing dependencies of S . The conditional distribution $P(y|do(x), do(\mathcal{D}_X), \mathcal{D}_Y)$ is transportable across domains but is not tractable as it requires marginalising over all possible input distributions (via the frontdoor criterion). We give more details in Appendix A.3. Instead, we use a partially transportable query $P(y|do(x), \mathcal{D}_X, \mathcal{D}_Y)$. Equation 2 formulates this query using observations only (proof in Appendix A.2):

$$\begin{aligned}
 & P(y|do(x), \mathcal{D}_X, \mathcal{D}_Y) \\
 &= \underbrace{\int_w P(w|\mathcal{D}_X, \mathcal{D}_Y)}_{\text{marginalisation over W}} \underbrace{\int_r P(r|x, w, \mathcal{D}_X, \mathcal{D}_Y)}_{\text{marginalisation over R}} \underbrace{\int_{x'} P(x'|\mathcal{D}_X, \mathcal{D}_Y)}_{\text{marginalisation over X}} \overbrace{P(y|x', r, w, \mathcal{D}_X, \mathcal{D}_Y)}^{\text{inference}} dx' dr dw
 \end{aligned}
 \tag{2}$$

The interventional query must be decomposed into four components to be represented using observations only. The marginalisation over W is equivalent to an application of the backdoor criterion and ensures that no backdoor paths exist between Y and W . The marginalisations over X and R are applications of the frontdoor criterion. The combination of these mechanisms ensure that no backdoor paths exist between the input X and the output Y that could bias the training.

Transportability We have identified an unbiased interventional query from observational information only. We now study the transportability of each component across domains using the \mathcal{S} -admissibility criterion (Jalaldoust & Bareinboim, 2024): conditional probabilities independent from the domain selection variable \mathcal{S} are transportable. We assume a training source domain \mathcal{M}^s and a target test domain \mathcal{M}^t , it follows graphically that:

$$(W \perp\!\!\!\perp \mathcal{S} | \mathcal{D}_X, \mathcal{D}_Y)_{\mathcal{G}^{\Delta_{st}}} \quad (3)$$

$$(R \perp\!\!\!\perp \mathcal{S} | X, W, \mathcal{D}_X, \mathcal{D}_Y)_{\mathcal{G}^{\Delta_{st}}} \quad (4)$$

$$(X \not\perp\!\!\!\perp \mathcal{S} | \mathcal{D}_X, \mathcal{D}_Y)_{\mathcal{G}^{\Delta_{st}}} \quad (5)$$

$$(Y \not\perp\!\!\!\perp \mathcal{S} | X, R, W, \mathcal{D}_X, \mathcal{D}_Y)_{\mathcal{G}^{\Delta_{st}}} \quad (6)$$

The first two quantities are \mathcal{S} -admissible and can be written as $P^s(w | \mathcal{D}_{X_s}, \mathcal{D}_{Y_s}) = P^t(w | \mathcal{D}_{X_t}, \mathcal{D}_{Y_t}) = P^*(w | \mathcal{D}_X, \mathcal{D}_Y)$ and $P^s(r | x, w, \mathcal{D}_{X_s}, \mathcal{D}_{Y_s}) = P^t(r | x, w, \mathcal{D}_{X_t}, \mathcal{D}_{Y_t}) = P^*(r | x, w, \mathcal{D}_X, \mathcal{D}_Y)$. The last two quantities are not \mathcal{S} -admissible and require to have access to information $\mathcal{D}_{X_t}, \mathcal{D}_{Y_t}$ on the target domain. This is an expected result for $P(x' | \mathcal{D}_X, \mathcal{D}_Y)$. Indeed, the input X directly depends on the distribution of its factors of variations. We must learn a suitable representation of X to use it for inference. Note that this probability is independent of Y and does not require access to labelled data. Under the current SCM problem formulation, $P(y | x', r, w, \mathcal{D}_X, \mathcal{D}_Y)$ is not \mathcal{S} -admissible because of the backdoor path between the training data \mathcal{D} and the inference target Y . However, in out-of-distribution settings, the training and target domains can differ and the backdoor path may not exist. This is a reasonable expectation as the converse of the \mathcal{S} -admissibility criterion does not hold: non \mathcal{S} -admissibility does not implies non-transportability. We discuss the limitations of the current formalism in Section 6.

Tractable Approximation The quantities above are not directly tractable and must be estimated. The marginalisation over W is a standard technique from Bayesian inference (Magris & Iosifidis, 2023). The integral can be approximated using Markov Chain Monte Carlo (MCMC) or Variational Inference (Blundell et al., 2015b; Magris & Iosifidis, 2023). Specifically:

$$\begin{aligned} & P(y | do(x), \mathcal{D}_X, \mathcal{D}_Y) \\ &= \int_w P(w | \mathcal{D}_X, \mathcal{D}_Y) \int_r P(r | x, w, \mathcal{D}_X, \mathcal{D}_Y) \int_{x'} P(x' | \mathcal{D}_X, \mathcal{D}_Y) P(y | x', r, w, \mathcal{D}_X, \mathcal{D}_Y) dx' dr dw \\ &\approx \frac{1}{M} \sum_{j=1}^M \int_r P(r | x, w_j, \mathcal{D}_X, \mathcal{D}_Y) \int_{x'} P(x' | \mathcal{D}_X, \mathcal{D}_Y) P(y | x', r, w_j, \mathcal{D}_X, \mathcal{D}_Y) dx' dr \end{aligned}$$

where $w_j \sim P(w | \mathcal{D}_X, \mathcal{D}_Y)$ and M is the amount of weight samples. The distribution of the weights W can be learned by a Bayesian Neural Network (BNN) (MacKay, 1992; Blundell et al., 2015b; Magris & Iosifidis, 2023). The network does not directly learn the weights w but instead estimate the parameters of the distribution in which the weights belong and sample the weights from this distribution at inference. The learning objective cannot be directly optimised because of the high-dimensional and non-convex posterior distribution $P(w | \mathcal{D})$. Instead, MCMC (Izmailov et al., 2021) and Variational Inference (Blundell et al., 2015b) algorithms are preferred. The marginalisation over R is approximated similarly:

$$\begin{aligned} & \int_r P(r | x, w_j, \mathcal{D}_X, \mathcal{D}_Y) \int_{x'} P(x' | \mathcal{D}_X, \mathcal{D}_Y) P(y | x', r, w_j, \mathcal{D}_X, \mathcal{D}_Y) dx' dr \\ &\approx \frac{1}{L} \sum_{k=1}^L \int_{x'} P(x' | \mathcal{D}_X, \mathcal{D}_Y) P(y | x', r_k, w_j, \mathcal{D}_X, \mathcal{D}_Y) dx' \end{aligned}$$

where $r_k \sim P(r | x, w_j, \mathcal{D}_X, \mathcal{D}_Y)$ and L is the amount of latent representation samples. Variational Autoencoders (VAEs) (Kingma & Welling, 2014) are commonly used to estimate this distribution. In

270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323

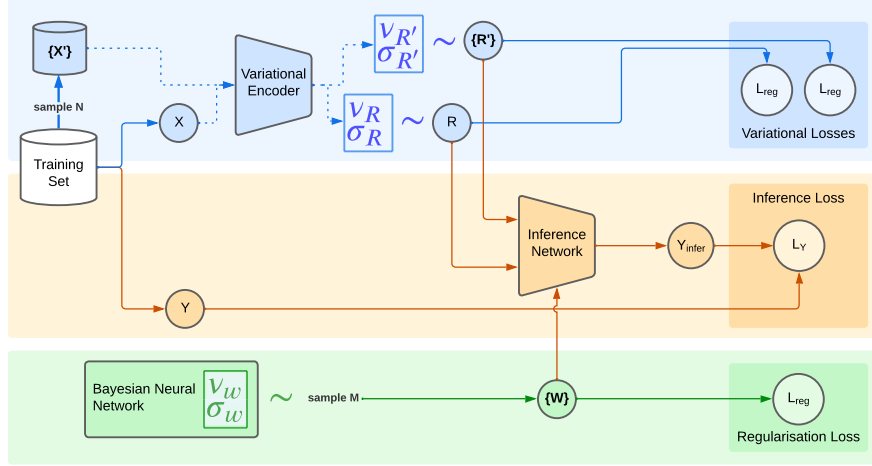


Figure 2: Architecture of the Causal-Invariant Bayesian (CIB) neural network. At the top (in blue), a variational encoder generates the respective parameters of the distributions of the intermediate representations R and $\{R'_i\}_{i=1}^N$ of the input X and the contextual information $\{X'_i\}_{i=1}^N$. R and $\{R'_i\}_{i=1}^N$ are then provided to the inference module (in orange) that retrieves Y . This procedure aims to disentangle the learning of the representation R from the learning of the inference mechanisms and force the inference module to only learn the latter. The weights W of the inference module are sampled from a distribution learned using variational inference (in green). The weight sampling and the variational encoding are regularised using an ELBO loss.

practice, we consider a single representation R per input and context image, i.e. $L = 1$. Finally, the last marginalisation can be obtained by sampling input images from the current data distribution:

$$\int_{x'} P(x' | \mathcal{D}_X, \mathcal{D}_Y) P(y | x', r, w_j, \mathcal{D}_X, \mathcal{D}_Y) dx' \approx \frac{1}{N} \sum_{k=1}^N P(x_i | \mathcal{D}_X, \mathcal{D}_Y) P(y | x_i, r_k, w_j, \mathcal{D}_X, \mathcal{D}_Y)$$

where N is the amount of input image samples.

4 CAUSAL-INVARIANT BAYESIAN NEURAL NETWORK

We now propose our Causal-Invariant Bayesian (CIB) architecture estimating the posterior probability defined above. The architecture is summarised in Figure 2. It can be decomposed into three components:

- A variational encoder (in blue) learning domain-invariant information R by jointly optimising the inference loss with an ELBO loss,
- An inference module (in orange) combining domain-specific and domain-invariant information to solve the task,
- A weight distribution W (in green) learned by optimising the inference objective with an ELBO loss and from which are sampled the parameters of the inference module.

Variational Encoder The variational encoder is used to learn the intermediate domain-invariant representations R and $\{R'_i\}_{i=1}^N$. The network weights are learned by optimising an estimate lower bound (ELBO) loss. Existing work attempting to disentangle domain-specific and domain-invariant features typically use a Variational Autoencoder (VAE) and add a reconstruction loss (Gabbay et al., 2021; Yang et al., 2021; Zhang et al., 2022). The input and the contextual information are fed to two different encoders, learning to specialise to invariant features and domain-specific features, respectively. However, we find that a single encoder can yield better performance. By jointly giving

the input and the context to the encoder, we teach it to discard irrelevant domain-specific features at once. We use a pre-trained ResNet-18 (He et al., 2016) without the final classification layer as the encoder. For comparison, we perform additional experiments with the Causal Transportability architecture that integrates a VAE and a reconstruction loss (Mao et al., 2022). Following the methodology of partially-stochastic neural networks (Sharma et al., 2023), we use a point-estimate network for the encoder and only use a Bayesian network for the inference network. This choice allows to maintain the expressivity on the level of uncertainty inherent to Bayesian neural networks while having a limited impact on the optimisation efficiency.

Bayesian Inference Network For the inference model, we use a dense network with stochastic layers for the input and output and a single point-estimate hidden layer. The inference module takes batches of N inputs. Each input is a sum of the input representation R and one context representation R'_i . We select the mean of the results with respect to N to estimate the marginalisation over the context. We use a reparametrisation trick (Blundell et al., 2015a) to maintain the full differentiability of the network. Following Elfwing et al. (2018), we change the activation layers from ReLU to SiLU.

Context Regularisation We add label-mixing regularisation to complement the contextual information during training, following the procedure used in Mixup (Zhang et al., 2018) and Manifold Mixup (Verma et al., 2019). We mix the true label \mathbf{y}_t with the labels $\mathbf{Y}_{\mathbf{x}_i}$ of the N context images x_i :

$$\mathbf{y} = \alpha \cdot \mathbf{y}_t + (1 - \alpha) \cdot \frac{1}{N} \sum_{i=1}^N \mathbf{y}_{\mathbf{x}_i} \quad (7)$$

Weight Function Regularisation We further regularise the weights sampling of the inference module to enhance the diversity of the weights while aligning it with functional diversity. Following the work of Roy et al. (2024), we add a regularisation loss inducing the network to assign the same posterior distribution to weights realising an identical function. For a batch of inputs \mathbf{x} and a set of weight samples $\mathcal{W} = \{w_i\}_{i=1}^M$, we define our *weight function regularisation* loss in Equation 8. $\text{Comb}(\mathcal{W}, 2)$ denotes the set of 2-combinations from \mathcal{W} .

$$\mathcal{L}_{weight_func} = \sum_{(w_i, w_j) \in \text{Comb}(\mathcal{W}, 2)} \|f_{w_i}(\mathbf{x}) - f_{w_j}(\mathbf{x})\|^2 \quad (8)$$

Loss function The final loss function comprises the mixed labels loss, the weight function regularisation and the regularisation terms of the ELBO losses. The mixed label loss is a cross entropy loss between the predicted distribution and the true mixed distribution as described above. We add the weight function regularisation loss described in the previous paragraph, weighted by a hyperparameter β . Then, KL divergence losses are added for regularising the variational parameters for the input representation (ν_R, σ_R) , the context representation $(\nu_{R'}, \sigma_{R'})$ and the Bayesian weights (ν_W, σ_W) . These additional losses are weighted by their respective hyperparameters γ , μ and ϵ . The complete loss is shown in Equation 9.

$$\mathcal{L} = \text{CrossEntropy}(\mathbf{y}_{\text{pred}}, \mathbf{y}) + \beta \cdot \mathcal{L}_{weight_func} + \gamma \cdot \text{KL}(\mathcal{N}(\nu_R, \sigma_R) \| \mathcal{N}(0, 1)) + \mu \cdot \text{KL}(\mathcal{N}(\nu_{R'}, \sigma_{R'}) \| \mathcal{N}(0, 1)) + \epsilon \cdot \text{KL}(\mathcal{N}(\nu_W, \sigma_W) \| \mathcal{N}(0, 1)) \quad (9)$$

5 EXPERIMENTS

5.1 DATASETS

We perform experiments on image recognition tasks. All datasets contain train, validation and test splits in-distribution (i.i.d) and out-of-distribution (o.o.d). We first conduct image recognition on the image recognition CIFAR10 dataset (Krizhevsky et al., 2009). We build the o.o.d sets by performing random translations of the input images. We also use the OFFICEHOME dataset (Venkateswara et al., 2017) which contains four subsets of images containing objects in various configurations (real world, product sheet, art, clipart). We train on one configuration and evaluate o.o.d in another configuration.

Table 1: Accuracy on CIFAR-10. The CIFAR-10 column shows the results on the i.i.d test set while the columns on the right show the results with an increasing level of perturbation of the test set. The mean and standard deviation across three runs are shown. Our model outperforms the baselines and the ablated models. Particularly, mixup information has a significant impact on the final accuracy.

	CIFAR-10	CIFAR-10-o.o.d-perturb.		
		0.1	0.2	0.4
ResNet-18	0.739 \pm 0.005	0.599 \pm 0.017	0.471 \pm 0.022	0.284 \pm 0.010
ResNet-18-CT	0.682 \pm 0.004	0.507 \pm 0.010	0.397 \pm 0.006	0.244 \pm 0.010
-w/-pretrainedVAE	0.641 \pm 0.004	0.464 \pm 0.017	0.358 \pm 0.015	0.229 \pm 0.010
CIBResNet-18 (ours)	0.763 \pm 0.004	0.646 \pm 0.006	0.510 \pm 0.012	<i>0.301 \pm 0.006</i>
-w/o-mixup	0.761 \pm 0.004	0.625 \pm 0.013	0.494 \pm 0.009	0.297 \pm 0.004
-w/o-weight-func	0.762 \pm 0.009	0.635 \pm 0.010	0.505 \pm 0.009	0.305 \pm 0.007

5.2 EXPERIMENTAL SETUP

We compare our CIBResNet-18 against ResNet-18 and ResNet-18-CT in i.i.d and o.o.d settings. The second baseline is the causal-Transportability (CT) algorithm of Mao et al. (2022). Following their methodology, we use a VAE for the $P(R|X)$ encoder. Their proposed inference network is a three-layers convolution network followed by a two-layers dense classifier. For a fair comparison, we use a modified ResNet-18 (the same size as our model) that takes input and contextual information. More details are given in Appendix B. We conduct our experiments on a single GPU Nvidia A100 with the AdamW optimiser (Loshchilov & Hutter, 2019). We initialise the Bayesian classifier parameters from the Normal distribution $\mathcal{N}(0, 1)$. On CIFAR-10, we train our model and the baselines for 20 epochs. We find that all models usually do not show improvements after 10 epochs. On OFFICEHOME, we train the models for 80 epochs. The training sets contain less samples than CIFAR-10 and more epochs are required for convergence. ResNet-18 stops improving after 40 epochs while CIBResNet-18 keep improving until around 60 epochs. We use hyperparameter grid search on the validation set of CIFAR-10 to set the hyperparameters. We use the default learning rate of 0.005 for the baselines and 0.01 for CIBResNet-18. We use a batch size of 64. By default, we use $N = 16$, $M = 16$, $\alpha = 0.4$, $\beta = 0.01$, $\gamma = 10^{-6}$, $\nu = 10^{-6}$ and $\epsilon = 10^{-6}$. The remaining hyperparameters are set to their default values. When evaluating the models in o.o.d settings, we use the batch statistics to build the mean and variance in the batch normalisation layers instead of the values learned on the training set that follow a different distribution.

5.3 VISUAL RECOGNITION

We compare our model against ResNet-18 and ResNet-18-CT on CIFAR10. The results are shown in Table 1. Curiously, the base ResNet-18 outperforms the ResNet-18-CT models. This result reinforces our hypothesis that a reconstruction loss is not necessary and can even be adversarial in some situations. We further study how our model behaves on a more challenging dataset and compare it with ResNet-18. Figure 3 shows the results on OFFICEHOME. Our model systematically outperforms the baselines. Ablation studies in Table 1 demonstrate the importance of adding contextual mixup information to the label during training to help the model incorporate the context images during learning. We further compare the training of our model with the baseline on CIFAR-10 and the first domain of OFFICEHOME in Figure 4. We observe that the validation accuracy of the baselines first decreases to reach a plateau but then increases again, highlighting an overfitting to the training data. However, this behaviour is not observed with CIBResNet-18, which steadily decreases across the entire training, demonstrating higher stability. We can also note that our proposed model yields a lower standard deviation in the accuracy.

5.4 IMPACT OF SAMPLING ON POSTERIOR

We perform additional experiments to investigate the impact of the number of samples. We vary the amounts N and M of context images and inference weights, respectively. Figure 5 shows the evolution of the accuracy of our model when varying N and M . The main contributing factor for performance is the number of context samples: using four context samples instead of one significantly

432
433
434
435
436
437
438
439
440
441
442
443

	RealWorld	Product	Art	Clipart		RealWorld	Product	Art	Clipart
RealWorld	0.253 (±0.019)	0.124 (±0.01)	0.091 (±0.016)	0.084 (±0.006)	RealWorld	0.303 (±0.009)	0.191 (±0.016)	0.121 (±0.009)	0.123 (±0.004)
Product	0.034 (±0.013)	0.509 (±0.061)	0.023 (±0.005)	0.023 (±0.009)	Product	0.052 (±0.003)	0.574 (±0.015)	0.036 (±0.002)	0.042 (±0.007)
Art	0.06 (±0.001)	0.019 (±0.003)	0.125 (±0.017)	0.035 (±0.001)	Art	0.076 (±0.019)	0.033 (±0.015)	0.159 (±0.023)	0.034 (±0.009)
Clipart	0.064 (±0.01)	0.03 (±0.007)	0.038 (±0.006)	0.462 (±0.025)	Clipart	0.083 (±0.009)	0.055 (±0.01)	0.046 (±0.012)	0.517 (±0.009)

Figure 3: Domain transfer results on the OFFICEHOME dataset. Each row represents the category of the training subset and each column represents the category of the test subset. Accuracy with a random guess is 0.015. In the right figure, a cell is shown in green if its value is higher than the baseline on the left. The mean and standard deviation across three runs are shown. Our proposed model (on the right) systematically outperforms the baseline (on the left). Only the model trained on the Art domain shows little to no improvement. As this domain demonstrates the lowest i.i.d and o.o.d accuracy, we explain it by the lack of exploitable training data.

451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466




Figure 4: Evolution of the validation loss during training. The mean and standard deviation across three runs are shown. After a period of improvement, the ResNet-18 and ResNet-18-CT baselines overfit as the training progresses. This behaviour is not observed with CIBResNet-18, which also demonstrates a lower standard deviation.

472
473
474
475
476

improves accuracy. We hypothesize that a lower number is not sufficient to be representative of the distribution and has an adversarial effect instead. Figure 6 shows the evolution of the accuracy during training for several values of N and M . We observe that increasing the number of weight samples M improves sample efficiency and reduces the amount of training steps required to converge.

6 LIMITATIONS

479
480
481
482
483
484
485

This study focuses on developing a theory of interventional queries for o.o.d tasks and implementing an architecture suitable for answering such queries. We do not focus on the optimisation and the efficiency of such a system. Bayesian neural networks are notoriously difficult to optimise (Magris & Iosifidis, 2023). As discussed in the experimental section, the sampling step can hurt the efficiency of the model. As argued by Hooker (2021), the available hardware has a great impact on the success of a machine learning method. The current hardware landscape is well suited for point estimate neural networks but not for Bayesian and variational methods requiring many sampling steps. We argue that developing better and usable Bayesian causal models will require, not only improving the theory, but




Figure 5: Accuracy heatmap of CIBResNet-18 on the CIFAR-10 test set as a function of the number of weight and context samples. The amount of weight samples has a negligible effect on performance. Only increasing the context size improves accuracy.




Figure 6: Evolution of the validation accuracy of CIBResNet-18 on CIFAR-10 during training with varying context images N and weight samples M . Increasing the weight samples reduces the amount of training steps required for learning.

more importantly developing new optimisation strategies that can bridge the efficiency gap with their point estimate counterparts.

We perform our analysis under the SCM framework (Pearl, 2009). The underlying causal model is represented by a DAG that does not allow feedback loops. As the training of a neural network is an iterative process, a representation that allows cycles could better represent the inner causal mechanisms behind the learning process and help improve generalisation. This is a challenging task (Bongers et al., 2021) that we will tackle in our future work.

7 CONCLUSION

We propose a Causal-Invariant Bayesian neural network architecture to improve domain generalisation on visual recognition tasks. We build a theoretical analysis for domain generalisation with deep learning based on causality theory and taking into account the parameter learning process. We show experimentally that following causal principles can improve robustness to distribution shifts and reduce overfitting. These problems are particularly prominent on strong reasoning tasks that require neural systems to use abstract domain-invariant mechanisms and ignore spurious domain-specific information (Chollet, 2019; Zhang et al., 2021). In our future work, we will extend our framework to solve these challenging reasoning tasks. This investigation also provides elements to motivate the development of a causality-based theory of deep learning following the Independent Causal Mechanisms principle (Peters et al., 2017; Schölkopf et al., 2021). Such theory could improve robustness and generalisation in neural networks and provide additional benefits such as improved interpretability and modularity.

REFERENCES

- Isabela Albuquerque, João Monteiro, Mohammad Darvishi, Tiago H Falk, and Ioannis Mitliagkas. Generalizing to unseen domains via distribution matching. *arXiv preprint arXiv:1911.00804*, 2019.
- Martín Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *CoRR*, abs/1907.02893, 2019. URL <http://arxiv.org/abs/1907.02893>.
- Elias Bareinboim and Judea Pearl. A general algorithm for deciding transportability of experimental results. *CoRR*, abs/1312.7485, 2013. URL <http://arxiv.org/abs/1312.7485>.

- 540 Elias Bareinboim, Juan D. Correa, Duligur Ibeling, and Thomas Icard. On pearl’s hierarchy and the
541 foundations of causal inference. In Hector Geffner, Rina Dechter, and Joseph Y. Halpern (eds.),
542 *Probabilistic and Causal Inference: The Works of Judea Pearl*, volume 36 of *ACM Books*, pp.
543 507–556. ACM, 2022. doi: 10.1145/3501714.3501743. URL [https://doi.org/10.1145/
544 3501714.3501743](https://doi.org/10.1145/3501714.3501743).
- 545 Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for
546 domain adaptation. In Bernhard Schölkopf, John C. Platt, and Thomas Hofmann (eds.), *Advances
547 in Neural Information Processing Systems 19, Proceedings of the Twentieth Annual Conference
548 on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December
549 4-7, 2006*, pp. 137–144. MIT Press, 2006. URL [https://proceedings.neurips.cc/
550 paper/2006/hash/b1b0432ceafb0ce714426e9114852ac7-Abstract.html](https://proceedings.neurips.cc/paper/2006/hash/b1b0432ceafb0ce714426e9114852ac7-Abstract.html).
- 551 Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty
552 in neural networks. *CoRR*, abs/1505.05424, 2015a. URL [http://arxiv.org/abs/1505.
553 05424](http://arxiv.org/abs/1505.05424).
- 554 Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in
555 neural network. In *International conference on machine learning*, pp. 1613–1622. PMLR, 2015b.
- 556 Stephan Bongers, Patrick Forré, Jonas Peters, and Joris M. Mooij. Foundations of structural causal
557 models with cycles and latent variables. *The Annals of Statistics*, 49(5):2885 – 2915, 2021. doi:
558 10.1214/21-AOS2064. URL <https://doi.org/10.1214/21-AOS2064>.
- 559 François Chollet. On the measure of intelligence. *CoRR*, abs/1911.01547, 2019. URL [http:
560 //arxiv.org/abs/1911.01547](http://arxiv.org/abs/1911.01547).
- 561 Stefan Elfving, Eiji Uchibe, and Kenji Doya. Sigmoid-weighted linear units for neural network
562 function approximation in reinforcement learning. *Neural Networks*, 107:3–11, 2018. doi:
563 10.1016/J.NEUNET.2017.12.012. URL [https://doi.org/10.1016/j.neunet.2017.
564 12.012](https://doi.org/10.1016/j.neunet.2017.12.012).
- 564 Aviv Gabbay, Niv Cohen, and Yedid Hoshen. An image is worth more than a thousand words: Towards
565 disentanglement in the wild. In Marc’ Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy
566 Liang, and Jennifer Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems
567 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December
568 6-14, 2021, virtual*, pp. 9216–9228, 2021. URL [https://proceedings.neurips.cc/
569 paper/2021/hash/4cef5b5e6ff1b3445db4c013f1d452e0-Abstract.html](https://proceedings.neurips.cc/paper/2021/hash/4cef5b5e6ff1b3445db4c013f1d452e0-Abstract.html).
- 570 Gaël Gendron, Michael Witbrock, and Gillian Dobbie. A survey of methods, challenges and
571 perspectives in causality. *CoRR*, abs/2302.00293, 2023. doi: 10.48550/ARXIV.2302.00293. URL
572 <https://doi.org/10.48550/arXiv.2302.00293>.
- 573 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
574 recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR
575 2016, Las Vegas, NV, USA, June 27-30, 2016*, pp. 770–778. IEEE Computer Society, 2016. doi:
576 10.1109/CVPR.2016.90. URL <https://doi.org/10.1109/CVPR.2016.90>.
- 577 Sara Hooker. The hardware lottery. *Commun. ACM*, 64(12):58–65, 2021. doi: 10.1145/3467017.
578 URL <https://doi.org/10.1145/3467017>.
- 579 Yimin Huang and Marco Valtorta. Pearl’s calculus of intervention is complete. In
580 *UAI ’06, Proceedings of the 22nd Conference in Uncertainty in Artificial Intelligence, Cambridge, MA, USA, July 13-16, 2006*. AUAI Press, 2006. URL
581 [https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=
582 1&smnu=2&article_id=1301&proceeding_id=22](https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=1&smnu=2&article_id=1301&proceeding_id=22).
- 582 Pavel Izmailov, Sharad Vikram, Matthew D. Hoffman, and Andrew Gordon Wilson. What are bayesian
583 neural network posteriors really like? In Marina Meila and Tong Zhang (eds.), *Proceedings of the
584 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*,
585 volume 139 of *Proceedings of Machine Learning Research*, pp. 4629–4640. PMLR, 2021. URL
586 <http://proceedings.mlr.press/v139/izmailov21a.html>.

- 594 Kasra Jalaldoust and Elias Bareinboim. Transportable representations for domain generaliza-
595 tion. In Michael J. Wooldridge, Jennifer G. Dy, and Sriraam Natarajan (eds.), *Thirty-Eighth*
596 *AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innova-*
597 *tive Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational*
598 *Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada,*
599 pp. 12790–12800. AAAI Press, 2024. doi: 10.1609/AAAI.V38I11.29175. URL <https://doi.org/10.1609/aaai.v38i11.29175>.
600
- 601 Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer
602 vision? In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus,
603 S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Sys-*
604 *tems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017,*
605 *Long Beach, CA, USA*, pp. 5574–5584, 2017. URL [https://proceedings.neurips.cc/](https://proceedings.neurips.cc/paper/2017/hash/2650d6089a6d640c5e85b2b88265dc2b-Abstract.html)
606 [paper/2017/hash/2650d6089a6d640c5e85b2b88265dc2b-Abstract.html](https://proceedings.neurips.cc/paper/2017/hash/2650d6089a6d640c5e85b2b88265dc2b-Abstract.html).
607
- 608 Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and
609 Yann LeCun (eds.), *2nd International Conference on Learning Representations, ICLR 2014,*
610 *Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. URL <http://arxiv.org/abs/1312.6114>.
611
- 612 Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
613
- 614 John Langford. Tutorial on practical prediction theory for classification. *J. Mach. Learn. Res.*, 6:
615 273–306, 2005. URL <http://jmlr.org/papers/v6/langford05a.html>.
616
- 617 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *7th International*
618 *Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019.*
619 OpenReview.net, 2019. URL <https://openreview.net/forum?id=Bkg6RiCqY7>.
620
- 621 David JC MacKay. A practical bayesian framework for backpropagation networks. *Neural computa-*
622 *tion*, 4(3):448–472, 1992.
- 623 Martin Magris and Alexandros Iosifidis. Bayesian learning for neural networks: an algorithmic
624 survey. *Artif. Intell. Rev.*, 56(10):11773–11823, 2023. doi: 10.1007/S10462-023-10443-1. URL
625 <https://doi.org/10.1007/s10462-023-10443-1>.
626
- 627 Chengzhi Mao, Kevin Xia, James Wang, Hao Wang, Junfeng Yang, Elias Bareinboim, and Carl
628 Vondrick. Causal transportability for visual recognition. In *IEEE/CVF Conference on Computer*
629 *Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pp.
630 7511–7521. IEEE, 2022. doi: 10.1109/CVPR52688.2022.00737. URL [https://doi.org/](https://doi.org/10.1109/CVPR52688.2022.00737)
631 [10.1109/CVPR52688.2022.00737](https://doi.org/10.1109/CVPR52688.2022.00737).
- 632 Judea Pearl. *Causality*. Cambridge university press, 2009.
633
- 634 Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations*
635 *and learning algorithms*. The MIT Press, 2017.
636
- 637 Jonathan Richens and Tom Everitt. Robust agents learn causal world models. In *The Twelfth Inter-*
638 *national Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024.*
639 OpenReview.net, 2024. URL <https://openreview.net/forum?id=pOoKI3ouv1>.
- 640 Hritik Roy, Marco Miani, Carl Henrik Ek, Philipp Hennig, Marvin Pförtner, Lukas Tatzel, and
641 Søren Hauberg. Reparameterization invariance in approximate bayesian inference. *CoRR*,
642 abs/2406.03334, 2024. doi: 10.48550/ARXIV.2406.03334. URL [https://doi.org/10.](https://doi.org/10.48550/arXiv.2406.03334)
643 [48550/arXiv.2406.03334](https://doi.org/10.48550/arXiv.2406.03334).
644
- 645 Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner,
646 Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proc. IEEE*, 109(5):
647 612–634, 2021. doi: 10.1109/JPROC.2021.3058954. URL [https://doi.org/10.1109/](https://doi.org/10.1109/JPROC.2021.3058954)
[JPROC.2021.3058954](https://doi.org/10.1109/JPROC.2021.3058954).

- 648 Mrinank Sharma, Sebastian Farquhar, Eric T. Nalisnick, and Tom Rainforth. Do bayesian neural
649 networks need to be fully stochastic? In Francisco J. R. Ruiz, Jennifer G. Dy, and Jan-Willem
650 van de Meent (eds.), *International Conference on Artificial Intelligence and Statistics, 25-27 April*
651 *2023, Palau de Congressos, Valencia, Spain*, volume 206 of *Proceedings of Machine Learning Re-*
652 *search*, pp. 7694–7722. PMLR, 2023. URL [https://proceedings.mlr.press/v206/](https://proceedings.mlr.press/v206/sharma23a.html)
653 [sharma23a.html](https://proceedings.mlr.press/v206/sharma23a.html).
- 654 Sjoerd van Steenkiste, Francesco Locatello, Jürgen Schmidhuber, and Olivier Bachem. Are dis-
655 entangled representations helpful for abstract visual reasoning? In Hanna M. Wallach, Hugo
656 Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett (eds.),
657 *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information*
658 *Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp.
659 14222–14235, 2019. URL [https://proceedings.neurips.cc/paper/2019/hash/](https://proceedings.neurips.cc/paper/2019/hash/bc3c4a6331a8a9950945a1aa8c95ab8a-Abstract.html)
660 [bc3c4a6331a8a9950945a1aa8c95ab8a-Abstract.html](https://proceedings.neurips.cc/paper/2019/hash/bc3c4a6331a8a9950945a1aa8c95ab8a-Abstract.html).
- 661 Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep
662 hashing network for unsupervised domain adaptation. In *2017 IEEE Conference on Computer*
663 *Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pp. 5385–5394.
664 IEEE Computer Society, 2017. doi: 10.1109/CVPR.2017.572. URL [https://doi.org/10.](https://doi.org/10.1109/CVPR.2017.572)
665 [1109/CVPR.2017.572](https://doi.org/10.1109/CVPR.2017.572).
- 666 Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz,
667 and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. In
668 Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International*
669 *Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*,
670 volume 97 of *Proceedings of Machine Learning Research*, pp. 6438–6447. PMLR, 2019. URL
671 <http://proceedings.mlr.press/v97/verma19a.html>.
- 672 Mengyue Yang, Furui Liu, Zhitang Chen, Xinwei Shen, Jianye Hao, and Jun Wang. Causalvae:
673 Disentangled representation learning via neural structural causal models. In *IEEE Conference*
674 *on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pp.
675 9593–9602. Computer Vision Foundation / IEEE, 2021. doi: 10.1109/CVPR46437.2021.00947.
676 URL [https://openaccess.thecvf.com/content/CVPR2021/html/Yang_](https://openaccess.thecvf.com/content/CVPR2021/html/Yang_CausalVAE_Disentangled_Representation_Learning_via_Neural_Structural_Causal_Models_CVPR_2021_paper.html)
677 [CausalVAE_Disentangled_Representation_Learning_via_Neural_](https://openaccess.thecvf.com/content/CVPR2021/html/Yang_CausalVAE_Disentangled_Representation_Learning_via_Neural_Structural_Causal_Models_CVPR_2021_paper.html)
678 [Structural_Causal_Models_CVPR_2021_paper.html](https://openaccess.thecvf.com/content/CVPR2021/html/Yang_CausalVAE_Disentangled_Representation_Learning_via_Neural_Structural_Causal_Models_CVPR_2021_paper.html).
- 679 Chi Zhang, Baoxiong Jia, Mark Edmonds, Song-Chun Zhu, and Yixin Zhu. ACRE: abstract causal
680 reasoning beyond covariation. In *IEEE Conference on Computer Vision and Pattern Recognition,*
681 *CVPR 2021, virtual, June 19-25, 2021*, pp. 10643–10653. Computer Vision Foundation / IEEE,
682 2021. doi: 10.1109/CVPR46437.2021.01050. URL [https://openaccess.thecvf.](https://openaccess.thecvf.com/content/CVPR2021/html/Zhang_ACRE_Abstract_Causal_REasoning_Beyond_Covariation_CVPR_2021_paper.html)
683 [com/content/CVPR2021/html/Zhang_ACRE_Abstract_Causal_REasoning_](https://openaccess.thecvf.com/content/CVPR2021/html/Zhang_ACRE_Abstract_Causal_REasoning_Beyond_Covariation_CVPR_2021_paper.html)
684 [Beyond_Covariation_CVPR_2021_paper.html](https://openaccess.thecvf.com/content/CVPR2021/html/Zhang_ACRE_Abstract_Causal_REasoning_Beyond_Covariation_CVPR_2021_paper.html).
- 685 Hanlin Zhang, Yi-Fan Zhang, Weiyang Liu, Adrian Weller, Bernhard Schölkopf, and Eric P. Xing.
686 Towards principled disentanglement for domain generalization. In *IEEE/CVF Conference on*
687 *Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*,
688 pp. 8014–8024. IEEE, 2022. doi: 10.1109/CVPR52688.2022.00786. URL [https://doi.org/](https://doi.org/10.1109/CVPR52688.2022.00786)
689 [10.1109/CVPR52688.2022.00786](https://doi.org/10.1109/CVPR52688.2022.00786).
- 690 Hongyi Zhang, Moustapha Cissé, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical
691 risk minimization. In *6th International Conference on Learning Representations, ICLR 2018,*
692 *Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net,
693 2018. URL <https://openreview.net/forum?id=r1Ddpl-Rb>.
- 694
695
696
697
698
699
700
701

A PROOFS FOR THE INTERVENTIONAL BAYESIAN INFERENCE

A.1 RULES OF DO-CALCULUS

We first describe the three rules of do-calculus (Pearl, 2009) as they are required for the proofs in Sections A.2 and A.3. The first rule states that an observation z can be ignored if it does not affect the outcome y of the query given the current causal graph. The second rule states that an intervention on a variable $do(z)$ can be considered an observation z if there are no backdoor paths linking it to the outcome y , i.e. if the variables Z and Y do not share any common ancestors not captured by the observation z . The third rule states that an intervention $do(z)$ can be ignored if there are no direct paths between Z and Y or backdoor paths between Y and observed descendants W of Z .

The three rules are formally defined as follows. For each rule, the equality holds if the independence test between the variable of interest Z and the outcome Y is verified in the given causal graph. The graph $\mathcal{G}_{\bar{X}}$ corresponds to the initial graph with the incoming links to X removed, corresponding to an intervention on X . The graph $\mathcal{G}_{\bar{X}\bar{Z}}$ furthermore removes the outgoing links of Z . The term $\overline{Z(W)}$ removes the incoming links of Z if Z is not an ancestor of the observed variable W .

- **Rule 1** Deletion of observation

$$P(y|do(x), z, w) = P(y|do(x), w) \text{ if } (Y \perp\!\!\!\perp Z|X, W)_{\mathcal{G}_{\bar{X}}} \quad (10)$$

- **Rule 2** Reduction of intervention to observation

$$P(y|do(x), do(z), w) = P(y|do(x), z, w) \text{ if } (Y \perp\!\!\!\perp Z|X, W)_{\mathcal{G}_{\bar{X}\bar{Z}}} \quad (11)$$

- **Rule 3** Deletion of intervention

$$P(y|do(x), do(z), w) = P(y|do(x), w) \text{ if } (Y \perp\!\!\!\perp Z|X, W)_{\mathcal{G}_{\bar{X}, \overline{Z(W)}}} \quad (12)$$

A.2 IDENTIFIABILITY OF $P(y|do(x), \mathcal{D}_X, \mathcal{D}_Y)$

Proof. Proof of Equation 2. We start by marginalising on W to block the backdoor path between R and Y ($R \leftarrow U_R \rightarrow \mathcal{D}_R \rightarrow W \rightarrow Y$) and use the rules of do-calculus to simplify the quantities. This step ensures that we can later apply the frontdoor criterion with R and block the backdoor paths between X and Y .

$$\begin{aligned} & P(y|do(x), \mathcal{D}_X, \mathcal{D}_Y) \\ &= \int_w P(y|do(x), \mathcal{D}_X, \mathcal{D}_Y, w) P(w|do(x), \mathcal{D}_X, \mathcal{D}_Y) dw && \text{Marginalisation over } W \\ &= \int_w \boxed{P(y|do(x), \mathcal{D}_X, \mathcal{D}_Y, w)} P(w|\mathcal{D}_X, \mathcal{D}_Y) dw && \text{Rule 3: } (W \perp\!\!\!\perp X|\mathcal{D}_X, \mathcal{D}_Y)_{\mathcal{G}_{\bar{X}}} \end{aligned}$$

We now focus on the quantity in the red box. We must use the frontdoor criterion to block all backdoor paths between X and Y : $X \leftarrow U_{XY} \rightarrow Y$, $X \leftarrow Z \rightarrow \mathcal{D}_X \leftarrow U_{XY} \rightarrow Y$. We marginalise on R .

$$\begin{aligned} & \boxed{P(y|do(x), \mathcal{D}_X, \mathcal{D}_Y, w)} \\ &= \int_r P(y|do(x), \mathcal{D}_X, \mathcal{D}_Y, w, r) P(r|do(x), \mathcal{D}_X, \mathcal{D}_Y, w) dr && \text{Marginalisation over } R \\ &= \int_r \boxed{P(y|do(x), \mathcal{D}_X, \mathcal{D}_Y, w, r)} P(r|x, \mathcal{D}_X, \mathcal{D}_Y, w) dr && \text{Rule 2: } (R \perp\!\!\!\perp X|\mathcal{D}_X, \mathcal{D}_Y, W)_{\mathcal{G}_{\bar{X}}} \end{aligned}$$

Again, we focus on the quantity in the blue box. We continue to apply the frontdoor criterion and marginalise on X after switching the intervened variable using do-calculus rules.

$$\begin{aligned}
& P(y|do(x), \mathcal{D}_X, \mathcal{D}_Y, w, r) \\
&= P(y|do(x), do(r), \mathcal{D}_X, \mathcal{D}_Y, w) && \text{Rule 2: } (Y \perp\!\!\!\perp R|X, \mathcal{D}_X, \mathcal{D}_Y, W)_{\mathcal{G}_{\overline{XR}}} \\
&= P(y|do(r), \mathcal{D}_X, \mathcal{D}_Y, w) && \text{Rule 3: } (Y \perp\!\!\!\perp X|R, \mathcal{D}_X, \mathcal{D}_Y, W)_{\mathcal{G}_{\overline{XR}}} \\
&= \int_{x'} P(y|do(r), \mathcal{D}_X, \mathcal{D}_Y, w, x') P(x'|do(r), \mathcal{D}_X, \mathcal{D}_Y, w) dx' && \text{Marginalisation over } X \\
&= \int_{x'} P(y|r, \mathcal{D}_X, \mathcal{D}_Y, w, x') P(x'|do(r), \mathcal{D}_X, \mathcal{D}_Y, w) dx' && \text{Rule 2: } (Y \perp\!\!\!\perp R|X, \mathcal{D}_X, \mathcal{D}_Y, W)_{\mathcal{G}_{\overline{R}}} \\
&= \int_{x'} P(y|r, \mathcal{D}_X, \mathcal{D}_Y, w, x') P(x'|\mathcal{D}_X, \mathcal{D}_Y, w) dx' && \text{Rule 3: } (X \perp\!\!\!\perp R|\mathcal{D}_X, \mathcal{D}_Y, W)_{\mathcal{G}_{\overline{R}}} \\
&= \int_{x'} P(y|r, \mathcal{D}_X, \mathcal{D}_Y, w, x') P(x'|\mathcal{D}_X, \mathcal{D}_Y) dx' && \text{Rule 1: } (X \perp\!\!\!\perp W|\mathcal{D}_X, \mathcal{D}_Y)_{\mathcal{G}}
\end{aligned}$$

We put together all the quantities and obtain Equation 2:

$$\begin{aligned}
& P(y|do(x), \mathcal{D}_X, \mathcal{D}_Y) \\
&= \int_w P(w|\mathcal{D}_X, \mathcal{D}_Y) \int_r P(r|x, w, \mathcal{D}_X, \mathcal{D}_Y) \int_{x'} P(x'|\mathcal{D}_X, \mathcal{D}_Y) P(y|x', r, w, \mathcal{D}_X, \mathcal{D}_Y) dx' dr dw
\end{aligned}$$

□

A.3 INTRACTABILITY OF $P(y|do(x), do(\mathcal{D}_X), \mathcal{D}_Y)$

Proof. Proof of the intractability of $P(y|do(x), do(\mathcal{D}_X), \mathcal{D}_Y)$. We show that estimating this query from observational data requires to marginalise on \mathcal{D}_X . We take as a postulate that this is an intractable operation because it implies to access all possible input domains. We also note that estimating this operation via sampling can be achieved but goes against the purpose of domain generalisation as we aim to adapt to new domains from limited information.

From the causal graph in Figure 1b, we observe two causal paths linking \mathcal{D}_X to Y :

1. A direct path: $\mathcal{D}_X \rightarrow \mathcal{D}_R \rightarrow W \rightarrow Y$
2. A backdoor path: $\mathcal{D}_X \leftarrow Z \rightarrow X \rightarrow R \rightarrow Y$
3. A backdoor path: $\mathcal{D}_X \leftarrow U_{XY} \rightarrow Y$

The third path is blocked by the do operation on X but the first two paths realise the frontdoor criterion as shown in the simplified causal graph in Figure 7. The expression can therefore be simplified as follows:

$$\begin{aligned}
& P(y|do(x), do(\mathcal{D}_X), \mathcal{D}_Y) \\
&= \int_a P(a|\mathcal{D}_X, do(x), \mathcal{D}_Y) \int_{\mathcal{D}'_X} P(y|do(x), \mathcal{D}'_X, \mathcal{D}_Y) P(\mathcal{D}'_X|do(x), \mathcal{D}_Y) d\mathcal{D}'_X \quad \text{frontdoor criterion}
\end{aligned}$$

Answering this query from observations requires marginalising on \mathcal{D}_X , as highlighted in blue.

We can see graphically that this path cannot be blocked by any other variable than \mathcal{D}_X because the variable U_{XY} is not observed, therefore marginalising on \mathcal{D}_X is necessary.

□




Figure 7: Simplified causal graph linking the input domain \mathcal{D}_X to the output label Y . The paths can be simplified as a backdoor path via U_{XY} and a direct path through an abstract variable A .

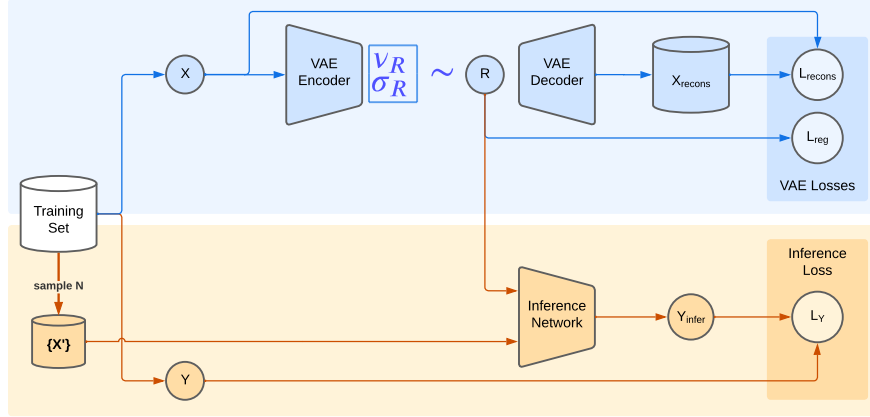


Figure 8: Illustration of the Causal Transportability baseline model.

B DETAILS ON THE CAUSAL TRANSPORTABILITY MODEL

We implement the Causal Transportability algorithm defined by Mao et al. (2022) as a baseline. The model is illustrated in Figure 8. The algorithm realises the following quantity:

$$P(y|do(x)) = P(r|x) \sum_{k=1}^N P(y|r, x_i) P(x_i) \quad (13)$$

The sampling strategy is the same as the one used in our work. The authors use a small network composed of three convolution layers followed by two fully connected layers to realise the quantity $P(y|r, x_i)$. For a fair comparison with our model, we replace it with a modified ResNet-18. We alter the first layer is modified to take a concatenation of an input image x with the representation r . The rest of the network is left untouched. The quantity $P(r|x)$ can be obtained via several methods. We use a VAE in our implementation, as the main method described in the original paper. We use the model in two settings. First, we jointly train the VAE and the inference network. Second, we compare this baseline against a second model where we train the VAE separately on a reconstruction task and then include the weights to the full model. We allow the training of the VAE weights at the second stage as we observed that freezing the weights leads to decreased performance.