

# Meta Learning for Causal Direction

Jean-François Ton,<sup>1</sup> Dino Sejdinovic,<sup>1</sup> Kenji Fukumizu<sup>2</sup>

<sup>1</sup> University of Oxford

<sup>2</sup> The Institute of Statistical Mathematics

ton@stats.ox.ac.uk, dino.sejdinovic@stats.ox.ac.uk, fukumizu@ism.ac.jp

## Abstract

The inaccessibility of controlled randomized trials due to inherent constraints in many fields of science has been a fundamental issue in causal inference. In this paper, we focus on distinguishing the cause from effect in the bivariate setting under limited observational data. Based on recent developments in meta learning as well as in causal inference, we introduce a novel generative model that allows distinguishing cause and effect in the small data setting. Using a learnt task variable that contains distributional information of each dataset (task), we propose an end-to-end algorithm that makes use of similar training datasets at test time. We demonstrate our method on various synthetic as well as real-world data and show that it is able to maintain high accuracy in detecting directions across varying dataset sizes.

## Introduction

Discovering causal links between variables has been a long standing problem in many areas of science. Ideally all experiments are in a randomized controlled environment where each variable can be accounted for separately. However in most cases, this is impossible due to physical, financial or ethical reasons. The problem of determining causal direction becomes even more apparent, when trying to understand the true generating process of data. Although modern machine learning models are able to achieve impressive performances in capturing complex nonlinear relationships among variables, many of them do not take into account causal structure, which might lead to generalization errors when faced with different data than seen during training.

Hence, in recent years, researchers have focused on inferring causal relations from observational data and have developed many different algorithms. One of the major approaches is constraint-based methods, which analyze the conditional independence among variables to determine the causal graph up to a Markov equivalence class under certain assumptions (Neuberg 2003); in addition to early example of the PC algorithm (Spirtes et al. 2000), there are also nonlinear methods for capturing independence (Sun et al. 2007; Sun, Janzing, and Schölkopf 2007; Zhang et al. 2011).

Another category is score-based methods which use search algorithms to find the best causal graph with respect

to such a score as BIC (Chickering 2002). These methods are, however, often unable to determine the correct structure and can be computationally very expensive. There are also hybrid methods which mitigate such difficulty (Tsamardinos, Brown, and Aliferis 2006).

A new line of research has taken specific interest in the bivariate case, i.e., the cause-effect inference, where one decides between causal hypotheses “ $X \rightarrow Y$ ” and “ $Y \rightarrow X$ ” (Hoyer et al. 2009; Goudet et al. 2017; Mitrovic, Sejdinovic, and Teh 2018; Wu and Fukumizu 2020). In this setting, methods that exploit the inherent asymmetries between cause and effect are the most prominent. The data is analysed under the *Functional Causal Model* (FCM, (Pearl 2009)) formalism, following respective model assumptions. Due to the intrinsic asymmetry of the problem, several statistics have been proposed to infer the direction of the cause-effect pairs (Shimizu et al. 2006; Hoyer et al. 2009; Mooij et al. 2016; Mitrovic, Sejdinovic, and Teh 2018).

Most relevant prior work to this paper is the framework of Causal Generative Neural Networks (CGNN, (Goudet et al. 2017)). When applied to cause-effect inference, CGNN learns a generative model using a neural network for each direction and compares their fitting to determine the causal directionality. Many advanced methods including CGNN, however, assume an access to a large dataset to make use of strong learning models such as neural networks. This in turn may lead significantly degraded performances in small data settings encountered in practical problems (Look and Riedelbauch 2018). We demonstrate this phenomenon in our experimental section. In addition, training a model for every dataset may cause significantly slow inference time by the computational burden of neural networks.

In this paper, we revisit the problem of cause-effect inference from the viewpoint of empirical learning. Specifically, we are interested in learning from many examples of cause-effect datasets together with their true causal directions. Our purpose is to develop a method for using this empirical knowledge on causal directions effectively, when making cause-effect inference on a new unseen and purely observational dataset.

Learning-based cause-effect inference has been already explored in the literature. For instance, Randomized Causation Coefficient (RCC) (Lopez-Paz, Muandet, and Recht 2015) and Neural Causation Coefficient (NCC) (Lopez-

Paz et al. 2017) make a learning-based binary classifier for causal directions. RCC and NCC, however, require to synthesize vast amounts of problem specific data pairs up to 10000 (Lopez-Paz et al. 2017). Other examples, such as NonSENS (Monti, Zhang, and Hyvarinen 2019) and Causal Mosaic (Wu and Fukumizu 2020), aim to recover a FCM using nonlinear independent component analysis. An important assumption for these methods is availability of multiple datasets sharing the same causal mechanism and the same exponential family of latent variables. We then need to assume or select datasets to satisfy this.

Different from these works, aiming at alleviating the problem of small data, we consider methods of meta learning by introducing a dataset-feature extractor. The feature represents the distributional information of each dataset, aiming to encode similar causal mechanisms of datasets into similar features. For this purpose, we employ two approaches: the formalism of kernel mean embeddings (Muandet et al. 2017) and DeepSets (Zaheer et al. 2017).

We propose a neural network-based generative model that trains jointly on all the training datasets. The model has an encoder-decoder architecture, which has been employed successfully by meta learning frameworks (Garnelo et al. 2018b,a); the encoder gives the dataset-features, and the decoder realizes a generative model or FCM, which is accompanied with the Feature-wise Linear Modulation layers (FiLM (Perez et al. 2018)) to adapt the generator to the dataset at hand. With this meta-learning architecture, the proposed method is able to determine the cause-effect direction efficiently for new unseen, possibly small, and purely observational datasets.

The contributions of this work can be summarized as follows:

- We introduce a new meta learning algorithm that can leverage similar datasets for unseen causal pairs in causal direction discovery.
- We exploit structural asymmetries with an adaptive generative model, thus avoiding the need to retrain at test time.
- We propose an end-to-end algorithm that assumes no *a priori* assumptions on the causal mechanism between cause and effect.
- High performance on small dataset sizes can be achieved by virtue of meta learning.

## Meta Learning for Detecting Causal Direction

We first give a brief summary of FCM and explain the proposed method including its building blocks.

### Functional Causal Model

Functional Causal Models (FCM) have been widely used when conducting causal inference. Formally, a FCM on a random vector  $X = (X_1, \dots, X_d)$  is a triplet  $C = (\mathcal{G}, f, \mathcal{E})$ , where  $\mathcal{G}$  is the causal graph and  $f, \mathcal{E}$  such that:

$$X_i = f_i(X_{Pa(i;\mathcal{G})}, Z_i), \quad Z_i \sim \mathcal{E}, \text{ for } i = 1, \dots, d, \quad (1)$$

where  $X_i$  are the observed variables,  $Z_i$  are the independent hidden variables,  $Pa(i;\mathcal{G})$  being the parents of  $X_i$  and

$f_i$  being the mechanism linking the cause and the effect. Under this formulation, there is clear asymmetry between cause and effect, given that cause is used to infer the effect, and hence numerous work has been done exploiting this fact.

For inferring causal direction  $X \rightarrow Y$  for bivariate  $(X, Y)$ , we can consider only the FCM  $Y = f(X, Z)$ , where  $Z$  and  $X$  are independent. Among other inference methods, CGNN (Goudet et al. 2017) uses neural networks to train the mechanism  $f$  for a dataset. More precisely, given dataset  $\mathcal{D} = \{(x_j, y_j)\}_{j=1}^m$ , we generate  $Z_j$  by the standard normal distribution  $N(0, 1)$ , and train neural network  $f$  so that the distribution of  $\{(X_j, Y_j)\}_j$  be close to that of  $\{(X_j, f(X_j, Z_j))\}_j$ . The difference of the distributions is measured by the Maximum Mean Discrepancy (MMD, (Gretton et al. 2012)). CGNN learns two models  $\hat{Y} = f_y(X, Z)$  and  $\hat{X} = f_x(Y, Z)$ , and chooses a better fit to determine the direction.

Unlike CGNN, which trains networks  $f_y$  and  $f_x$  for each dataset, our method considers a single neural network working for all the datasets in a cause-effect database  $\{\mathcal{D}_i\}_{i=1}^N$  where  $\mathcal{D}_i = \{(X_j^i, Y_j^i)\}_{j=1}^{m_i}$  is a dataset in the database. We assume that the causal direction is known for all  $\mathcal{D}_i$  during training. More specifically, given  $X^i \rightarrow Y^i$  is the true causal direction, we wish to create a single suitable model  $F(X, Z)$  based on neural networks so that the distribution of  $\mathcal{D}_i$  is approximately the same as that of  $\{(X_j^i, \hat{Y}_j^i)\}$  for *any*  $i$ , where  $\hat{Y}_j^i = F(X_j^i, Z_j^i)$ .

This approach involves obvious difficulty, since a wide-variety of cause-effect relations must be learnt by a single network. Naïvely training a single model jointly over all the different dataset does not yield desired performance as we demonstrate in our ablation study in the Appendix. In order to achieve successful training, we introduce two novel and crucial components:

1. **Dataset-feature:** This feature  $C$  represents the causal mechanism of each dataset as the distributional information on the dataset  $\mathcal{D}$ . The feature will be used to adapt our single network efficiently at test time.
2. **FiLM layers:** To adapt the base neural network to each dataset using the dataset-features  $C$ , the FiLM layers enable us to adapt the weights of our network to a given new dataset  $\mathcal{D}$  quickly at test time.

Together, with these two additional apparatus, we are able to train our model across datasets and therefore harness information from all the datasets together, instead of treating them independently as it is done for example in CGNN.

In the next section we describe methods to capture distributional information of dataset  $\mathcal{D}_i$ , by leveraging, the well studied area of conditional mean embeddings (CME) (Song, Fukumizu, and Gretton 2013) as well as DeepSets (Zaheer et al. 2017). For a briefly give an high level overview of meta learning see Appendix.

### Dataset Features via DeepSets

DeepSets (Zaheer et al. 2017) have been used as task embedding  $C_i$  in previous meta learning literature (Garnelo et al. 2018a,b; Xu et al. 2019). Using a neural network  $\phi_{x,y}$ , the

task embedding is defined by

$$C_i = \frac{1}{m_i} \sum_{j=1}^m \phi_{x,y}([x_j, y_j]). \quad (2)$$

DeepSets is a simple flexible approach to encoding sets into vectors, which is also permutation invariant. The latter is important as we do not want the embeddings to change solely based on the order of the elements in the dataset. Zaheer et al. (2017) show that DeepSets is a universal approximating for any set function. Hence this aggregation method allows us to have a good representation of the dataset.

However, given that we use a concatenation in DeepSets we do not encode the conditional distribution information but rather the joint. Therefore, in this paper, we in addition also consider conditional mean embeddings as dataset-features.

### Dataset features via Conditional Mean Embeddings (CME)

Kernel mean embeddings of distributions provide a powerful framework for representing probability distributions (Song, Fukumizu, and Gretton 2013; Muandet et al. 2017). Given sets  $\mathcal{X}$  and  $\mathcal{Y}$ , with a distribution  $P$  over the random variables  $(X, Y)$  taking values in  $\mathcal{X} \times \mathcal{Y}$ , the conditional mean embedding (CME) of the conditional density  $p(y|x)$ , is defined as:

$$\mu_{Y|X=x} := \mathbb{E}_{Y|X=x}[\phi_y(Y)] = \int_{\mathcal{Y}} \phi_y(y)p(y|x)dy. \quad (3)$$

where  $\phi_y$  is the feature map associated to the reproducing kernel Hilbert space (RKHS) of  $Y$ ,  $\mathcal{H}_Y$ . Intuitively, the equation above allows us to represent a probability distribution  $p(y|x)$  in a function space such as a RKHS, by taking the expectation under  $p(y|x)$  of the features  $\phi_y(y) \in \mathcal{H}_Y$ . Hence, for each value of the conditioning variable  $x$ , we obtain  $\mu_{Y|X=x} \in \mathcal{H}_Y$ .

Following (Song, Fukumizu, and Gretton 2013), the CME can be associated with the operator  $\mathcal{C}_{Y|X} : \mathcal{H}_X \rightarrow \mathcal{H}_Y$ , known as the conditional mean embedding operator (CMEO), which satisfies

$$\mu_{Y|X=x} = \mathcal{C}_{Y|X}\phi_x(x) \quad (4)$$

where  $\mathcal{C}_{Y|X} := \mathcal{C}_{YX}\mathcal{C}_{XX}^{-1}$  with  $\mathcal{C}_{YX} := \mathbb{E}_{Y,X}[\phi_y(Y) \otimes \phi_x(X)]$  and  $\mathcal{C}_{XX} := \mathbb{E}_{X,X}[\phi_x(X) \otimes \phi_x(X)]$ . The operator inverse should be understood as a regularized inverse, unless the rigorous inverse exists.

As a result, the finite sample estimator of  $\mathcal{C}_{Y|X}$  based on the dataset  $\{(x_j, y_j)\}_{j=1}^n$  can be written as

$$\widehat{\mathcal{C}}_{Y|X} = \Phi_y(K + \lambda I)^{-1}\Phi_x^T \quad (5)$$

where  $\Phi_y := (\phi_y(y_1), \dots, \phi_y(y_n))$  and  $\Phi_x := (\phi_x(x_1), \dots, \phi_x(x_n))$  are the feature matrices,  $K := \Phi_x^T \Phi_x$  is the kernel matrix with entries  $K_{i,j} = k_x(x_i, x_j) := \langle \phi_x(x_i), \phi_x(x_j) \rangle$ , and  $\lambda > 0$  is a regularization parameter.

Hence  $\widehat{\mu}_{Y|X=x} = \widehat{\mathcal{C}}_{Y|X}\phi_x(x)$  simplifies to a weighted sum of the feature maps of the observed points  $y_i$ :

$$\widehat{\mu}_{Y|X=x} = \sum_{i=1}^n \beta_i(x)\phi_y(y_i) = \Phi_y\beta(x), \quad (6)$$

$$\beta(x) = (\beta_1(x), \dots, \beta_n(x))^T = (K + \lambda I)^{-1}K_{:x}, \quad (7)$$

where  $K_{:x} = (k_x(x, x_1), \dots, k_x(x, x_n))^T$ .

In fact, when using finite-dimensional feature maps, the conditional mean embedding operator is simply a solution to a vector-valued ridge regression problem (regressing  $\phi_y(y)$  to  $\phi_x(x)$ ), which allows computation scaling linearly in the number of observations  $n$ .

The Woodbury matrix identity allows us to have computations of either order  $\mathcal{O}(n^3)$  or  $\mathcal{O}(d^3) + \mathcal{O}(d^2n)$ , where  $d$  is the dimension of the feature map  $\phi_x$ . In our case, given that we are in the meta learning setting, the dataset size  $n$  is usually rather small and hence the CME can be efficiently computed.

The CME is a canonical way for capturing conditional densities and thus the mechanism in a functional causal model. Therefore the CME also encodes the causal direction. We give further motivations on why we use CMEO as dataset-features in the next few sections.

### A Feature-wise Linear Modulation (FiLM)

We propose an architecture that is able to adapt the network  $f(X, Z)$ , i.e. a FCM, with the dataset-feature by using the Feature-wise Linear Modulation (FiLM) (Perez et al. 2018).

The FiLM layers are known to allow network adaptation to new environments quickly without adding further model parameters. They have been shown to work effectively in various tasks of computer vision (Perez et al. 2018) and regression (Requeima et al. 2019). In essence, the FiLM layers work as follows: given a conditioning variable  $C$  (this may be the label for image classification) and  $l_a$  being the  $a^{\text{th}}$  layer of a network, the FiLM layer  $FL_a$ , constructed by a neural network, adapts  $l_a$  to  $l_a^{FL}$  by

$$(\beta_a, \gamma_a) = FL_a(C), \quad l_a^{FL} = \beta_a + \gamma_a \circ l_a, \quad (8)$$

where  $\circ$  is the element-wise multiplication. Intuitively, the FiLM layer learns shift and scale parameters, conditioned on  $C$ , for any given layer  $l_a$ . We shall use the CME or DeepSets for  $C$ .

### Proposed Method

We propose a new meta learning algorithm, *meta-CGNN*, which works for cause-effect inference given cause-effect training database  $\mathbb{D} = \{\mathcal{D}_i\}_{i=1}^N$  with known directionality, where  $\mathcal{D}_i = \{(x_j^i, y_j^i)\}_{j=1}^{m_i}$ . Without loss of generality, we assume  $X^i \rightarrow Y^i$  for any dataset  $i$ . The FCM is trained on this database  $\mathbb{D}$ , and is used to infer the causal direction for an unseen dataset  $\mathcal{D}_{test}$ .

**Overview:** We use the popular encoder-decoder based architecture for meta learning, which is similar to the Neural Process (Garnelo et al. 2018b). The encoder first maps the

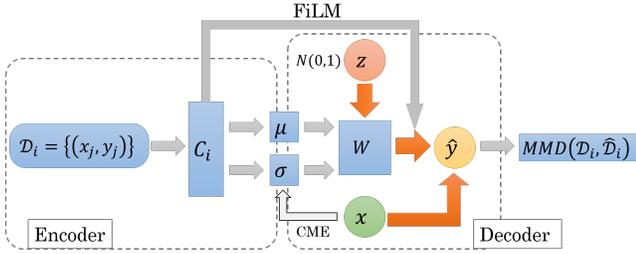


Figure 1: Proposed meta-CGNN Algorithm for only one dataset  $\mathcal{D}_i$  in the mini-batch

dataset  $\mathcal{D}_i$  into a dataset-feature  $C_i$ , which is given as an input to two further neural networks, (1) FiLM network and (2) amortization network. The FiLM network, operates as described in the above, by producing shift and scale parameters that allow us to adapt the decoder  $D_{FL}$  accordingly. The amortization network outputs  $(\mu(C), \sigma(C))$ , with which we use to modulate the latent random variable  $Z_j^i \sim N(0, 1)$  to  $W_j^i := \mu(C_i) + \sigma(C_i)Z_j^i$ .

The decoder network,  $D_{FL}$ , then maps  $(X, W)$  to  $\hat{Y}$  so that the distribution of  $\{(X_j^i, \hat{Y}_j^i)\}_j$  is close to  $\{(X_j^i, Y_j^i)\}_j$ . By using an encoder network, which trains across datasets jointly, we are able to share distributional information between datasets (tasks) and apply it to a new unseen task. See Figure 1 and Algorithm 1 (Appendix) for a detailed breakdown of *meta-CGNN*.

The overall functional causal model in the proposed meta-CGNN is thus

$$\hat{y}_j = F((x_j, z_j); C) \text{ , where } z_j \sim N(0, 1), \quad z_j \perp\!\!\!\perp x_j. \quad (9)$$

**Encoder:** For representing dataset-specific distributional information or mechanism for each task, we consider both the CME and DeepSets approach.

As argued by Mitrovic, Sejdinovic, and Teh (2018), CMEs hold critical information for causal-effect inference by representing the mechanism in the FCM. The claim is that the Kolmogorov complexity (Grunwald et al. 2008) of the mechanism tends to be larger in the anti-causal direction than it is in the causal direction. Hence, we expect CME to capture relevant distributional information, and inform adapting our generative model to the task at hand.

More concretely, we use conditional mean embeddings as follows. We first compute the CMEO  $\mathcal{C}_{Y|X}$  from  $\mathcal{D}_i$  as described in (Eq 5), and use it to obtain the CME for each datapoint using:

$$C_{i,j} = \mathcal{C}_{Y|X} \phi_x(x_j). \quad (10)$$

where,  $\phi_x$  is the feature vector such that  $k(x_i, x_k) = \langle \phi_x(x_j), \phi_x(x_k) \rangle$ ,  $k$  being the RBF kernel. To obtain a finite dimensional representation efficiently, we will use Random Fourier Features (Rahimi and Recht 2008) to approximate the CME. Throughout the paper we will use 100 features as we work on problems with at most 1500 datapoints. According to (Li et al. 2019), we need approximately  $\sqrt{N}$ ,

where  $N$  is the number of datapoints. For DeepSets the dataset feature  $C_i$  is simply defined in (Eq 2).

**Decoder:** For the generative part of our model, the dataset-feature  $C_i$  or  $C_{i,j}$  gives modulation through the FiLM and amortization network. Both FiLM and amortization networks take as input  $C_i$  (DeepSets) or  $C_{i,j}$  (CME). The FiLM layer is able to adapt the weights of the decoder network depending on the distributional feature of a dataset. This is crucial for a single network to learn FCMs of all the datasets. Naïvely training a single network over multiple datasets resulted in poor performance (see ablation study in Appendix). The amortization network works on  $z \sim N(0, 1)$  similarly to FiLM. It can be interpreted as the adaption of latent Gaussian distribution;  $p(w|C_i)$  with  $W^i = \mu(C_i) + \sigma(C_i)Z$  regarded as a new latent variable for the dataset  $\mathcal{D}_i$ .

Together with the FiLM layer we construct a decoder which generates data from the conditional distribution, by firstly sampling from  $p(w|\mathcal{D}_i)$  and concatenate  $w$  with  $x$  before pushing it through the decoder  $D_{FL}$ . This novel architecture, allows us to model and more importantly, **sample** from the conditional distribution of unseen task quickly and efficiently. In the next section, we will describe how we will make use of the samples to train our networks.

$$\hat{y}_j = D_{FL}((x_j, w_j); C) \text{ , where } w_j \sim N(\mu(C), \sigma(C)) \quad (11)$$

**Training:** The objective function of training is similar to (Goudet et al. 2017); by sampling  $\{\hat{y}_j^i\}_{j=1}^{m_i}$  from (11) and estimating the Maximum Mean Discrepancy (MMD) (Gretton et al. 2012) between the sampled data  $\hat{\mathcal{D}}_i = \{(x_j^i, \hat{y}_j^i)\}_{j=1}^{m_i}$  and the original data  $\mathcal{D}_i = \{(x_j^i, y_j^i)\}_{j=1}^{m_i}$ . MMD is a popular metric to measure the distance between two distributions. It has been widely used in two-sample tests (Gretton et al. 2012) as well as in applications of Generative Adversarial Network (GAN) (Li et al. 2017) etc.

More formally, the MMD estimator between two datasets  $\mathcal{U} = \{u_i\}_{i=1}^m$  and  $\mathcal{V} = \{v_i\}_{i=1}^n$  is, given by

$$\begin{aligned} \widehat{\text{MMD}}^2(\mathcal{U}, \mathcal{V}) &= \frac{1}{m(m-1)} \sum_{i=1}^m \sum_{j \neq i}^m k(u_i, u_j) \\ &+ \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j \neq i}^n k(v_i, v_j) \\ &- \frac{2}{mn} \sum_{i=1}^m \sum_{j=1}^n k(u_i, v_j). \end{aligned}$$

We use Gaussian kernel, which is characteristic (Sriperumbudur, Fukumizu, and Lanckriet 2011). With  $k$  being the Gaussian kernel, the above expression is differentiable and thus can be optimized as already demonstrated in various works such as (Goudet et al. 2017; Li et al. 2017). A drawback of using MMD as a loss function however is that it scales quadratically in the sample size. We can again use

Random Fourier Features (Rahimi and Recht 2008; Lopez-Paz, Muandet, and Recht 2015), which give us linear-time estimators of MMD.

Using stochastic mini-batches of  $q$  datasets  $\{\mathcal{D}_i\}_{i=1}^q$ , the objective function to minimize is

$$LOSS = \sum_{i=1}^q \widehat{\text{MMD}}^2(\widehat{\mathcal{D}}_i, \mathcal{D}_i) \quad (12)$$

This joint training, similar to that used in other encoder-decoder architectures for meta learning (Garnelo et al. 2018a,b; Xu et al. 2019), allows us to utilize the information of all the available training datasets in a single generative model. This is in stark contrast to CGNN (Goudet et al. 2017), which trains a separate generative model for each dataset. As noted in (Goudet et al. 2017) training these neural networks separately can be very costly and one of the major drawbacks of the CGNN. Hence *meta-CGNN* aims to alleviate this constraint by training over datasets jointly. This considerably speeds up the model inference time as *meta-CGNN* does not need to retrain the network for a new task.

**Inference of Causal Direction:** After training, when we wish to infer the causal direction for a new dataset  $\mathcal{D}_{test} = \{(x_j, y_j)\}_{j=1}^m$ , we feed both of  $\mathcal{D}^{xy} = \{(x_j, y_j)\}_{j=1}^m$  and  $\mathcal{D}^{yx} = \{(y_j, x_j)\}_{j=1}^m$  into the trained model and estimate the MMDs between the generated samples and the true ones, i.e.,  $\mathcal{M}_{xy} := \widehat{\text{MMD}}(\mathcal{D}^{xy}, \widehat{\mathcal{D}}^{xy})$  and  $\mathcal{M}_{yx} := \widehat{\text{MMD}}(\mathcal{D}^{yx}, \widehat{\mathcal{D}}^{yx})$ . If  $\mathcal{M}_{xy} < \mathcal{M}_{yx}$ , we deduce that  $X \rightarrow Y$  as it agrees better with a postulated FCM than  $Y \rightarrow X$ , and similarly  $Y \rightarrow X$  if  $\mathcal{M}_{xy} > \mathcal{M}_{yx}$ . Intuitively, this means that we choose the direction of whose samples that match the ground truth best.

## Related Work

There have been already some works that exploit the asymmetry by taking a closer look at decomposing the joint distribution  $P(X, Y)$  into either  $P(Y|X)P(X)$  or  $P(X|Y)P(Y)$ . Relevant to this work is the approach using the asymmetry in terms of functional causal models (FCMs). Some of the previous methods make strong assumptions on the model; LiNGAM (Shimizu et al. 2006) considers linear non-Gaussian model for finding causal structure. For nonlinear relations, Hoyer et al. (2009) discusses nonlinear additive noise models, and Zhang and Hyvarinen (2009) invertible interactions between the covariates and the noise models. There are other methods to consider the nonlinear models such as Gaussian process regression (Stegle et al. 2010).

Information theory gives an alternative view on asymmetry using Kolmogorov complexity, following the postulate that the mechanism in the causal direction should be less complex than the one in the anti-causal direction. Several papers have proposed to approximate or use certain proxies for intractable Kolmogorov complexity (Janzing and Scholkopf 2010; Lemeire and Dirkx 2006; Daniusis et al. 2010; Mitrovic, Sejdinovic, and Teh 2018).

Method	Inference Time
ANM	<1sec CPU
CDS	<1sec CPU
RECI	<1sec CPU
ICGI	<1sec CPU
RCC	<1sec CPU
CGNN	24 mins GPU
meta-CGNN (DeepSets)	<1min GPU
meta-CGNN (CME)	<1min GPU

Table 1: Time needed during inference i.e. to determine the causal direction for a new dataset at test time. *meta-CGNN* is much faster than CGNN, which is crucial, given that CGNN needs to be trained from scratch for every new dataset.

The method that comes closest to ours is CGNN (Goudet et al. 2017). However, our meta-CGNN method differs from CGNN in a multitude of aspects.

1. Our method employs meta learning, while CGNN only considers one dataset at a time. Hence CGNN is not able leverage similarity between datasets. A naïve way of training the CGNN jointly over datasets as in (Eq 12) was analysed in our ablation study and performed poorly.
2. CGNN averages over 32-64 separately trained generative networks per direction, which is computationally very expensive, i.e. training up to 128 models per dataset. In contrast, we merely train the model with four random initializations and average over the resulting MMDs for each dataset, thereby achieving similar results to CGNN for larger datasets and significantly better for smaller ones.
3. CGNN needs to be trained separately for every new dataset, which in practice can be slow as well as computationally expensive. *meta-CGNN* does not need to be trained for a new dataset and can give a causal direction through a simple forward pass at the test time. This is a crucial difference which allows much faster inference once a new dataset is presented to the model.

Lastly, there have recently been other applications of meta learning to causal inference. MetaCI (Sharma et al. 2019) uses a MAML-style (Finn, Abbeel, and Levine 2017) (optimization based) learner but mainly deals with counterfactual causality rather than causal directionality.

Bengio et al. (2020) also consider a meta learning method using FCM, based on the principle that the model assuming true causal direction can be “adapted faster” than the one assuming the anti-causal. However, their method is designed for different settings. Firstly, the test distribution comes from a perturbation of the training distribution, which is modeled by a known parametric family. The choice of the model is not trivial for continuous domains in real-world data. Secondly, for training of neural networks, they assume to have access to a large training dataset around 3000 for a single mechanism, which is different from our setting of small data.

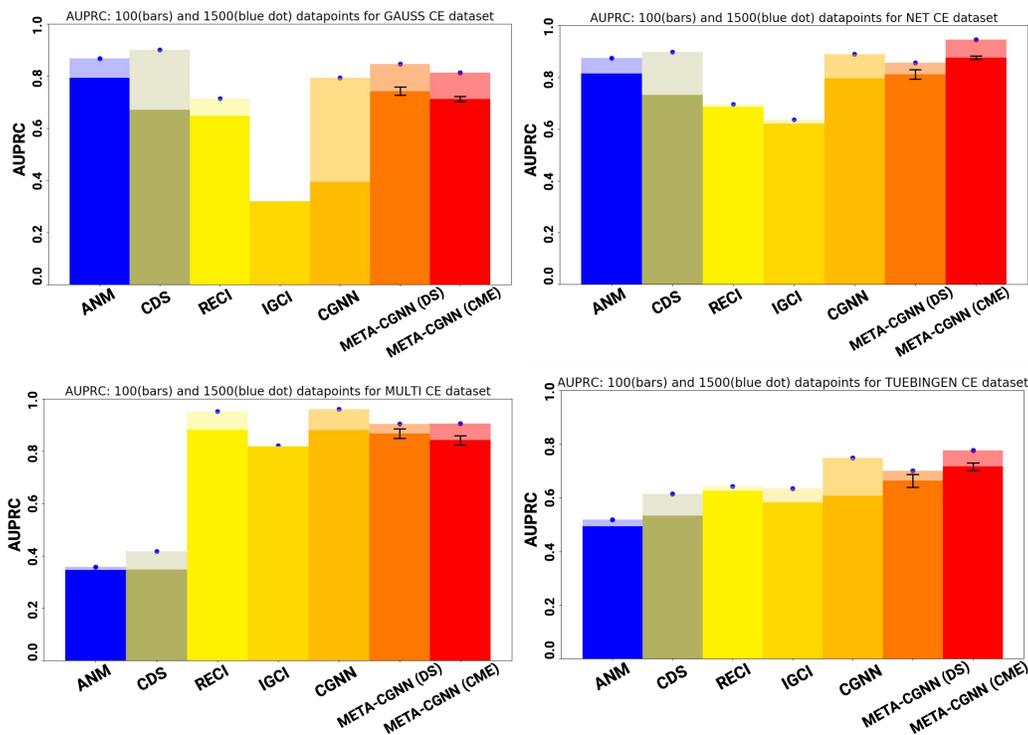


Figure 2: AUPRC for Net, Multi, Gauss and Tuebingen dataset. The thin colored bars with blue dots represent the AUPRC with 1500 datapoints, whereas the thicker barplots are with 100 datapoints. The proposed *meta-CGNN* shows among the best results in all the cases. Note that the other methods show significant degradation for some datasets or small data size.

## Experiments

### Synthetic Datasets

For the synthetic experiments we use three different types of datasets taken from Goudet et al. (2017), each of which exhibits distinct cause-effect (CE) mechanism. The *CE-Net* contains 300 cause-effect pairs with random distributions for the causes and random Neural Networks as mechanisms to create the effects. The *CE-Gauss* also contains 300 data pairs, where the cause is sampled from a random mixture of Gaussians and the mechanism is drawn from a Gaussian process (Mooij et al. 2016). Lastly, *CE-Multi* datasets take the cause from Gaussian distributions and the mechanisms are built using random linear and polynomial functions. They also include multiplicative and additive noise before or after the causal mechanism, making the task harder. In addition, to confirm the advantage of meta learning, we use two different data size regimes: 1500 and 100 datapoints.

We measure the performance of distinguishing the causal direction using the Area Under the Precision Recall Curve (AUPRC), as in (Goudet et al. 2017) (accuracy is presented in the Appendix). With AUPRC, we are able to take into account the confidence of an algorithm, thus allowing models not to commit to a prediction if not certain.

For *meta-CGNN*, we use 100 datasets for training and the remaining 200 for testing. We average the MMD of each dataset over 4 independent runs in order to get our final prediction. At testing time, we only need to do a simple forward

pass through our model, as our model has been trained to adapt to new dataset quickly and efficiently. Hence it takes **<1 minute** for each new dataset at inference time. This is contrary to CGNN that needs to be trained on each new dataset separately. Note that CGNN (Goudet et al. 2017) significantly benefits from averaging their model over multiple runs i.e. around 32-64 different runs, which takes about **24 minutes** per dataset. This would be infeasible without high-performance computing environments as we have hundreds of datasets to do inference on. Hence we have restricted ourselves to averaging CGNN over 12 runs in the comparisons, which is still computationally very heavy.

Regarding architectures, we use 2 hidden layers with ReLU activation function for FiLM, amortization and encoder network. For the decoder we use a 1 hidden layer with ReLU. As noted in Goudet et al. (2017), the number of hidden nodes in the decoder is very important; too small networks are not able to realize the mapping properly, while too big networks tend to overfit so that both directions have low MMD values. Hence for our experiments, for simplicity, we solely cross-validated over the number of decoder nodes [5, 40] (Goudet et al. 2017) by leaving out a few datasets at training aside for validation. Around 40 nodes for the 1-hidden layer was optimal for (Goudet et al. 2017).

In addition, following (Goudet et al. 2017), for our loss function, we use a sum over Gaussian kernel MMDs with bandwidth  $\eta \in \{0.005, 0.05, 0.25, 0.5, 1, 5, 50\}$  together with an Adam optimizer (Kingma and Ba 2014) and a

fixed learning rate 0.01. For the mini-batch size we fix  $q$  to 10. These were parameters we fixed in the beginning and seemed to work well in our experiments. We compare *meta-CGNN* against several competing ones that have open source codes; the methods are:

1. Additive noise model (ANM) (Mooij et al. 2016) with Gaussian process regression and HSIC test.
2. Information Geometric Causal Inference (IGCI) (Daniusis et al. 2010) with entropy estimator and Gaussian reference measure
3. Conditional Distribution Similarity statistic (CDS) (Fonollosa 2019), which analyses the variance of the conditional distributions
4. Regression Error based Causal Inference (RECI) (Blöbaum et al. 2019), which analyses the residual of each direction using a polynomial fit
5. Randomized Causation Coefficient (Lopez-Paz, Muandet, and Recht 2015), which builds creates a synthetic classification problem and use CME as the feature.
6. Causal Generative Neural Networks (CGNN) (Goudet et al. 2017)

We use the implementation by (Kalainathan and Goudet 2020) which provides a GitHub repository toolbox for the above mentioned methods. In order to keep the comparisons fair, we use the same 200 data pairs as the one that our *meta-CGNN* is tested on. Finally, in order to demonstrate the effect of each building block in our model, we have also conducted an ablation study, highlighting the importance of the task embeddings and FiLM layer (see Appendix).

### Tuebingen Cause-Effect dataset

As a real-world example, we use the popular Tuebingen benchmark (Mooij et al. 2016), from which we take 99 bivariate datasets. We use a similar setup as in the synthetic experiments, and the only difference is that we employ 5-fold cross-validation for training and testing, i.e. We train on 4 of the folds and test on the last one. We repeat this procedure such that each fold has been the test fold at one point, while the remaining were acting as training. That way we obtain a prediction on each of the 99 datasets. We repeat this 3 times with different random splits and report the results in Figure 2, where we also check the performance on decreasing size of samples in a dataset.

### Results

Figure 2 illustrates how both our *meta-CGNN* algorithms are the only ones that can retain high AUPRC across different datasets, as well as dataset sizes, i.e. 1500 and 100 datapoints. **NOTE:** *meta-CGNN* performs well in all the small dataset settings, while remaining computationally more efficient at inference time, in contrast to CGNN, which has significantly worse result in the small data for *CE-Gauss* and *CE-Tueb* dataset and requires training the whole model at test time. Another important point to notice is that similarly to CGNN, *meta-CGNN* does not have any assumptions on the data and hence can be used in a variety

of datasets, while retaining good performance. Algorithms such as ANM (Mooij et al. 2016), for example, seems to do reasonably well on the *CE-Net* and *CE-Gauss* dataset, but completely fails in the *CE-Multi* and, more importantly, on the real-world Tuebingen dataset. This occurs mainly because of the strict assumptions which ANM imposes on the FCM. Similarly, RECI which performs well on the *CE-Multi*, but not on the *CE-Net*, *CE-Gauss* and the Tuebingen dataset. Similarly for the remaining methods.

Our proposed method is amongst the top performing ones and is consistently doing well for both 1500 and 100 datapoints settings. These results illustrate that *meta-CGNN* is able to retain high performance, even when faced with small data, by leveraging the meta-learning setting.

Lastly, note that *meta-CGNN* does not need to be trained for each test dataset, but instead only needs a forward pass through the generative model to determine the causal direction, which makes it vastly more computationally efficient than CGNN, while attaining higher performance (Table 1).

### Conclusion

We introduced a novel meta learning algorithm for cause-effect inference which performs well even in the small data regime, in sharp contrast to the existing methods. By leveraging a dataset-feature extractor that can be learned during training, we are able to efficiently adapt our model at test time to new previously unseen datasets by using amortization and FiLM layers. We also demonstrate the utility of using conditional mean embeddings, as they allow us to capture the distribution and adapt the model at test time.

In addition, we extended Causal Generative Neural Network (CGNN) (Goudet et al. 2017) by learning a single generative network, readily adaptable for new datasets, vastly alleviating the computation burden of CGNNs. In particular, instead of having to train multiple models on each dataset separately, our proposed methods are able to achieve similar or better performance than existing methods with simple forward passes through our generative network at test time.

Recently there has been an increase in interest in causality topics in reinforcement learning (Zhu and Chen 2020; Buesing et al. 2019; Dasgupta et al. 2019), where the proposed methods may also be applicable. Assuming that we have a skeleton graph of the relevant quantities in the RL model, *meta-CGNN* could be used to efficiently infer causal direction of unoriented edges.

### Acknowledgements

We would like to thank Pengzhou (Abel) Wu for interesting discussions. JFT is supported by the EPSRC and MRC through the OxWaSP CDT programme EP/L016710/1. KF is supported by SPS KAKENHI 18K19793 and JST CREST JPMJCR2015.

### References

Bengio, Y.; Deleu, T.; Rahaman, N.; Ke, R.; Lachapelle, S.; Bilaniuk, O.; Goyal, A.; and Pal, C. 2020. A meta-transfer objective for learning to disentangle causal mechanisms.

- In *International Conference on Learning Representations*. URL <https://openreview.net/forum?id=ryxWlgBFPS>.
- Blöbaum, P.; Janzing, D.; Washio, T.; Shimizu, S.; and Schölkopf, B. 2019. Analysis of cause-effect inference by comparing regression errors. *PeerJ Computer Science* 5: e169.
- Buesing, L.; Weber, T.; Zwols, Y.; Racaniere, S.; Guez, A.; Lespiau, J.-B.; and Heess, N. 2019. Woulda, coulda, shoulda: Counterfactually-guided policy search. In *International Conference on Learning Representations*. URL <https://openreview.net/forum?id=BJG0voC9YQ>.
- Chickering, D. M. 2002. Optimal structure identification with greedy search. *Journal of machine learning research* 3(Nov): 507–554.
- Daniusis, P.; Janzing, D.; Mooij, J.; Zscheischler, J.; Steudel, B.; Zhang, K.; and Schölkopf, B. 2010. Inferring deterministic causal relations. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence, UAI'10*, 143–150.
- Dasgupta, I.; Wang, J.; Chiappa, S.; Mitrovic, J.; Ortega, P.; Raposo, D.; Hughes, E.; Battaglia, P.; Botvinick, M.; and Kurth-Nelson, Z. 2019. Causal reasoning from meta-reinforcement learning. *arXiv preprint arXiv:1901.08162*.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 1126–1135. JMLR. org.
- Fonollosa, J. A. 2019. Conditional distribution variability measures for causality detection. In *Cause Effect Pairs in Machine Learning*, 339–347. Springer.
- Garnelo, M.; Rosenbaum, D.; Maddison, C. J.; Ramalho, T.; Saxton, D.; Shanahan, M.; Teh, Y. W.; Rezende, D. J.; and Eslami, S. 2018a. Conditional neural processes. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, 1704–1713. PMLR. URL <http://proceedings.mlr.press/v80/garnelo18a.html>.
- Garnelo, M.; Schwarz, J.; Rosenbaum, D.; Viola, F.; Rezende, D. J.; Eslami, S.; and Teh, Y. W. 2018b. Neural processes. *arXiv preprint arXiv:1807.01622*.
- Goudet, O.; Kalainathan, D.; Caillou, P.; Guyon, I.; Lopez-Paz, D.; and Sebag, M. 2017. Causal generative neural networks. *arXiv preprint arXiv:1711.08936*.
- Gretton, A.; Borgwardt, K. M.; Rasch, M. J.; Schölkopf, B.; and Smola, A. 2012. A kernel two-sample test. *Journal of Machine Learning Research* 13(Mar): 723–773.
- Grunwald, P. D.; et al. 2008. Algorithmic information theory. *arXiv preprint arXiv:0809.2754*.
- Hoyer, P. O.; Janzing, D.; Mooij, J. M.; Peters, J.; and Schölkopf, B. 2009. Nonlinear causal discovery with additive noise models. In *Advances in neural information processing systems*, 689–696.
- Janzing, D.; and Schölkopf, B. 2010. Causal inference using the algorithmic Markov condition. *IEEE Transactions on Information Theory* 56(10): 5168–5194.
- Kalainathan, D.; and Goudet, O. 2020. Causal Discovery Toolbox: Uncover causal relationships in Python. *Journal of Machine Learning Research* 21(37): 1–5. URL <http://jmlr.org/papers/v21/19-187.html>.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Lemeire, J.; and Dirx, E. 2006. Causal models as minimal descriptions of multivariate systems. URL [http://parallel.vub.ac.be/~jan/papers/JanLemeire\\_CausalModelsAsMinimalDescriptionsOfMultivariateSystems\\_October2006.pdf](http://parallel.vub.ac.be/~jan/papers/JanLemeire_CausalModelsAsMinimalDescriptionsOfMultivariateSystems_October2006.pdf).
- Li, C.-L.; Chang, W.-C.; Cheng, Y.; Yang, Y.; and Póczos, B. 2017. MMD GAN: Towards deeper understanding of moment matching network. In *Advances in Neural Information Processing Systems*, 2203–2213.
- Li, Z.; Ton, J.-F.; Oglic, D.; and Sejdinovic, D. 2019. Towards A Unified Analysis of Random Fourier Features. In *International Conference on Machine Learning (ICML)*, PMLR 97:3905–3914. URL <http://proceedings.mlr.press/v97/li19k.html>.
- Look, A.; and Riedelbauch, S. 2018. Learning with Little Data: Evaluation of Deep Learning Algorithms. <https://openreview.net/pdf?id=rylU8oRctX>.
- Lopez-Paz, D.; Muandet, K.; and Recht, B. 2015. The randomized causation coefficient. *The Journal of Machine Learning Research* 16(1): 2901–2907.
- Lopez-Paz, D.; Nishihara, R.; Chintala, S.; Schölkopf, B.; and Bottou, L. 2017. Discovering causal signals in images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6979–6987.
- Mitrovic, J.; Sejdinovic, D.; and Teh, Y. W. 2018. Causal inference via kernel deviance measures. In *Advances in Neural Information Processing Systems*, 6986–6994.
- Monti, R. P.; Zhang, K.; and Hyvarinen, A. 2019. Causal discovery with general non-linear relationships using non-linear ICA. In *35th Conference on Uncertainty in Artificial Intelligence (UAI 2019)*.
- Mooij, J. M.; Peters, J.; Janzing, D.; Zscheischler, J.; and Schölkopf, B. 2016. Distinguishing cause from effect using observational data: methods and benchmarks. *The Journal of Machine Learning Research* 17(1): 1103–1204.
- Muandet, K.; Fukumizu, K.; Sriperumbudur, B.; and Schölkopf, B. 2017. Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends® in Machine Learning* 10(1-2): 1–141. doi:10.1561/22000000060. URL <http://dx.doi.org/10.1561/22000000060>.
- Neuberg, L. G. 2003. Causality: models, reasoning, and inference, by judea pearl, cambridge university press, 2000. *Econometric Theory* 19(4): 675–685.
- Pearl, J. 2009. Causality. <http://bayes.cs.ucla.edu/BOOK-2K/>.
- Perez, E.; Strub, F.; De Vries, H.; Dumoulin, V.; and Courville, A. 2018. Film: Visual reasoning with a general conditioning layer. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

- Rahimi, A.; and Recht, B. 2008. Random features for large-scale kernel machines. In *Advances in neural information processing systems*, 1177–1184.
- Requeima, J.; Gordon, J.; Bronskill, J.; Nowozin, S.; and Turner, R. E. 2019. Fast and flexible multi-task classification using conditional neural adaptive processes. In *Advances in Neural Information Processing Systems*, 7957–7968.
- Sharma, A.; Gupta, G.; Prasad, R.; Chatterjee, A.; Vig, L.; and Shroff, G. 2019. MetaCI: Meta-Learning for Causal Inference in a Heterogeneous Population. *arXiv preprint arXiv:1912.03960*.
- Shimizu, S.; Hoyer, P. O.; Hyvärinen, A.; and Kerminen, A. 2006. A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research* 7(Oct): 2003–2030.
- Song, L.; Fukumizu, K.; and Gretton, A. 2013. Kernel embeddings of conditional distributions: A unified kernel framework for nonparametric inference in graphical models. *IEEE Signal Processing Magazine* 30(4): 98–111.
- Spirtes, P.; Glymour, C. N.; Scheines, R.; and Heckerman, D. 2000. *Causation, prediction, and search*.
- Sriperumbudur, B. K.; Fukumizu, K.; and Lanckriet, G. R. 2011. Universality, characteristic kernels and RKHS embedding of measures. *Journal of Machine Learning Research* 12(70): 2389–2410. URL <http://jmlr.org/papers/v12/sriperumbudur11a.html>.
- Stegle, O.; Janzing, D.; Zhang, K.; Mooij, J. M.; and Schölkopf, B. 2010. Probabilistic latent variable models for distinguishing between cause and effect. In *Advances in neural information processing systems*, 1687–1695.
- Sun, X.; Janzing, D.; and Schölkopf, B. 2007. Distinguishing between cause and effect via kernel-based complexity measures for conditional distributions. In *15th European Symposium on Artificial Neural Networks (ESANN 2007)*, 441–446. D-Side Publications.
- Sun, X.; Janzing, D.; Schölkopf, B.; and Fukumizu, K. 2007. A Kernel-Based Causal Learning Algorithm. In *Proceedings of the 24th International Conference on Machine Learning, ICML '07*, 855–862. New York, NY, USA: Association for Computing Machinery. ISBN 9781595937933. doi: 10.1145/1273496.1273604. URL <https://doi.org/10.1145/1273496.1273604>.
- Tsamardinos, I.; Brown, L. E.; and Aliferis, C. F. 2006. The max-min hill-climbing Bayesian network structure learning algorithm. *Machine learning* 65(1): 31–78.
- Wu, P.; and Fukumizu, K. 2020. Causal Mosaic: Cause-Effect Inference via Nonlinear ICA and Ensemble Method. In *Proceedings of Artificial Intelligence and Statistics*.
- Xu, J.; Ton, J.-F.; Kim, H.; Kosiorek, A. R.; and Teh, Y. W. 2019. MetaFun: Meta-Learning with Iterative Functional Updates. *arXiv preprint arXiv:1912.02738*.
- Zaheer, M.; Kottur, S.; Ravanbakhsh, S.; Póczos, B.; Salakhutdinov, R. R.; and Smola, A. J. 2017. Deep sets. In *Advances in neural information processing systems*, 3391–3401.
- Zhang, K.; and Hyvärinen, A. 2009. On the identifiability of the post-nonlinear causal model. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, UAI '09*, 647–655.
- Zhang, K.; Peters, J.; Janzing, D.; and Schölkopf, B. 2011. Kernel-based conditional independence test and application in causal discovery. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence, UAI '11*, 804–813.
- Zhu, S.; and Chen, Z. 2020. Causal discovery with reinforcement learning. In *International Conference on Learning Representations*. URL <https://openreview.net/forum?id=S1g2skStPB>.