

Adversarial Orthogonal Regression: Two non-Linear Regressions for Causal Inference

M. Reza Heydari,¹ Saber Salehkaleybar,¹ Kun Zhang²

¹Department of Electrical Engineering, Sharif University of Technology, heydari_mr@ee.sharif.edu, saleh@sharif.edu

²Department of Philosophy, Carnegie Mellon University, kunz1@cmu.edu

Abstract

We propose two nonlinear regression methods, named Adversarial Orthogonal Regression (AdOR) for additive noise models and Adversarial Orthogonal Structural Equation Model (AdOSE) for the general case of structural equation models. Both methods try to make the residual of regression independent from regressors, while putting no assumption on noise distribution. In both methods, two adversarial networks are trained simultaneously where a regression network outputs predictions and a loss network that estimates mutual information (in AdOR) and KL-divergence (in AdOSE). These methods can be formulated as a minimax two-player game; at equilibrium, AdOR finds a deterministic map between inputs and output and estimates mutual information between residual and inputs, while AdOSE estimates a conditional probability distribution of output given inputs. The proposed methods can be used as subroutines to address several learning problems in causality, such as causal direction determination (or more generally, causal structure learning) and causal model estimation. Synthetic and real-world experiments demonstrate that the proposed methods have remarkable performance with respect to previous solutions.

1 Introduction

Identifying cause-effect relationships between variables in complex high dimensional networks has been studied in many fields such as neuroscience (Jazayeri and Afraz 2017; Shadlen et al. 1996), computational genomics (Marbach et al. 2012; Haury et al. 2012), economics (Zellner 1988), and social networks (Ver Steeg and Galstyan 2012; 2013). For instance, in genomics, it is known that each cell of living creatures consists of a huge number of genes that produce proteins in a procedure called “gene expression,” in which they can inhibit or promote each others’ activities. These cause-effect relationships can be represented by a causal graph in which each variable is depicted by a node, and a directed edge that shows the direct causal effect from the “parent” node to the “child” node. It is commonly assumed that there is no directed cycle in the causal graph, i.e., it is a Directed Acyclic Graph (DAG). The goal is to recover the causal graph from the data sampled from variables. In the literature, learning causal graphs has been studied extensively in two main settings: random variables and time series.

In the setting of random variables, Shimizu et al.(2006) proposed LiNGAM algorithm which can identify the causal graph in linear model under the assumption of non-Gaussianity of exogenous noises in the system. Hoyer et al.(2009) proposed a method to reveal the direction of causality in additive noise model where the effect is a function of direct causes plus some exogenous noise. The basic idea of their method is the following: for a given candidate DAG, one solves a regression problem for each node, modeling it as a (possibly nonlinear) function of its parents. Then, a statistical independence test is performed to assess whether all residuals are jointly independent. If that is the case, the candidate DAG is accepted, otherwise it is rejected. Peters, Janzing, and Schölkopf(2013) extended this idea for time series in additive noise models. All these methods require nonparametric nonlinear regression such that it ensures the residual is independent of regressors.

In the setting of time series, much efforts exerted to define statistical definition of causality such as Granger causality (Granger 1969; 1963). Marko(1973) defined an information theoretic measure called Directed Information (DI), which is a statistical criterion to detect the existence of direct causal effect between any pair of time series. Based on DI and inspired by G-causality, Quinn, Kiyavash, and Coleman(2015) proved that minimal generative model, i.e., a graph with minimum number of edges that does not miss the full dynamics, can be discovered by causally conditioned DI. Experiments showed that the proposed criterion can be used to reconstruct efficiently the causal graphs with linear relationships.

Causally conditioned DI and the other information theoretic measures for causality in time series typically utilize “differential entropy,” (Peters, Janzing, and Schölkopf 2017) which is an extension of Shannon entropy for continuous random variables. Since differential entropy is defined based on the probability distribution, numerous works have been done for entropy estimation of general distributions using only observational data. In this regard, Hausser and Strimmer(2009) used a naive binning method to estimate the value of joint distribution in each bin and then adjusted these values by a shrinkage factor based on James-Stein estimator (James and Stein 1992). In (Liu, Aviyente,

and Al-khassaweneh 2009; Darbellay and Tichavsky 2000; Miller 2003), the joint distribution is estimated by partitioning the domain in such a way that more accurate values are achieved in the regions where the density of sampled data is high. However, the proposed methods are sophisticated and need huge computational cost in high dimension. Recently, Quinn, Kiyavash, and Coleman(2015) used a regression based method for estimating DI. In order to check whether a variable Y is the parent of variable X , two regressions are performed: one by considering the Y in the regressors, and another without it. Then, DI can be obtained by differing the entropy of residuals in two regressions. The Y is considered as a parent of X if DI is non-zero. The above procedure works correctly only if the obtained residuals are independent of regressors in both regressions.

According to what mentioned above, several causal learning algorithms in the setting of random variables (such as the one in (Hoyer et al. 2009)) or time series (such as TiMINO algorithm in (Peters, Janzing, and Schölkopf 2013) or DI estimator in (Quinn, Kiyavash, and Coleman 2015)), require a subroutine that can perform non-linear regression such that the residual becomes independent of the regressors as much as possible. However, common regression methods are confined to minimize Mean Squared Error (MSE) loss (Wang and Bovik 2009; Kay 1993). Thus, in these common methods, the residuals and regressors become only uncorrelated. While these methods are fully efficient in linear Gaussian case, they might not be statistically efficient in nonlinear or non-Gaussian scenarios. To resolve this issue, Mooij et al.(2009) proposed a novel regression method which minimizes the dependence between residuals and regressors that is measured by Hilbert-Schmidt Independence Criterion (HSIC). In the proposed method, it is needed to carefully tune the kernel parameter in HSIC.

Contributions: In this paper, we propose two nonlinear regression methods, named Adversarial Orthogonal Regression (AdOR) and Adversarial Orthogonal Structural Equation Model (AdOSE). AdOR assumes that the noise is modeled as an additive term while AdOSE relaxes this assumption. The models are “Adversarial”, in the sense that in both methods, two neural networks compete with each other, the regression network and the loss network. In AdOR, the loss network estimates the mutual information between regressors and residuals, and in AdOSE, it acts as a Kullback-Leibler (KL)-divergence estimator between correct responses and predicts (which are the output of regression network). As discussed above, independence of residuals and regressors is vital in inferring the correct causal relationships. Thus, AdOR tries to make the residual independent of regressors, and AdOSE achieves this target by independently generating noise. The proposed methods can be used as subroutines to address several learning problems in causality, such as determining causal direction, causal structure learning, or causal model estimation. Experiment show that the proposed methods have remarkable performance in estimating the true non-linear function with respect to previous solutions. While our main contribution is in causal inference, the proposed methods might also be useful in the other regression tasks.

The rest of the paper is organized as follows: In Section 2, we review a neural network (Belghazi et al. 2018) that has been proposed previously to estimate mutual information. We describe AdOR and AdOSE methods in Section 3 and Section 4, respectively. We provide experimental results in Section 5 and conclude the paper in Section 6.

2 Mutual Information Neural Estimation

In this section, we describe the neural network proposed in (Belghazi et al. 2018) for estimating mutual information based on an alternative representations of KL-divergence. This representation will be exerted as the loss network in Section 3 and 4.

Let P and Q be two distributions on some compact domain $\Omega \subset \mathbb{R}^d$. The KL-divergence between them is defined as:

$$D_{KL}(P||Q) := \mathbb{E}_P \left[\log \frac{dP}{dQ} \right]. \quad (1)$$

One of the representation of KL-divergence, which we focused on, is Donsker-Varadhan representation (Donsker and Varadhan 1983):

$$D_{KL}(P||Q) = \sup_{T: \Omega \rightarrow \mathbb{R}} \mathbb{E}_P [T] - \log (\mathbb{E}_Q [e^T]), \quad (2)$$

where the supremum is taken over all functions T such that the two expectations are finite.

Let X and Y denote two continuous random variables with distributions P_X and P_Y , respectively. Mutual information between X and Y is denoted by $I(X; Y)$, which is a measure for the dependence of them. mutual information has some multiple forms, and one form is defined as the KL-divergence between the joint distribution P_{XY} , and the product of marginal distributions $P_X P_Y$:

$$I(X; Y) = D_{KL}(P_{XY}||P_X P_Y). \quad (3)$$

Let $\mathcal{F} = \{T_\theta\}_{\theta \in \Theta}$ be the set of functions parametrized by a neural network (i.e. weights, biases, batch normalization parameters, etc.). Mutual Information Neural Estimator (MINE, see definition 3.1 of (Belghazi et al. 2018)) is defined as:

$$\hat{I}(X; Y) = \sup_{\theta \in \Theta} \mathbb{E}_{P_{XY}} [T_\theta] - \log (\mathbb{E}_{P_X P_Y} [e^{T_\theta}]). \quad (4)$$

As the class of all functions in (2) is restricted to neural network class \mathcal{F} in (4), we have the following lower bound:

$$I(X; Y) \geq \hat{I}(X; Y). \quad (5)$$

Theoretical properties of $\hat{I}(X; Y)$ are provided in (Belghazi et al. 2018). In MINE, samples from joint distribution P_{XY} are fed as the inputs of a neural network and an optimizer like stochastic gradient descent, updates the parameters θ so as to maximize the right hand side of (4). Ultimately, as the parameters converge, the loss value of network is the estimated mutual information. For more details on the implementation of MINE, please refer to Algorithm 1 of (Belghazi et al. 2018).

3 Adversarial Orthogonal Regression

Let Z and U represent the scalar response and regressor vector, respectively. The regression problem is to find \hat{f} :

$$Z = \hat{f}(U) + \varepsilon, \quad (6)$$

such that the residual ε is independent of U .

In AdOR method, the regression network (R) is pitted against the loss network where a mutual information estimator (MI) learns to find any high order dependencies (see the top block diagram of Figure 1). In regression part, $\hat{Z} = \hat{f}(U; \theta_R)$ is a differentiable function represented by a multilayer perceptron, and parametrized with θ_R , in which \hat{Z} is the regression output. The residual $\varepsilon = Z - \hat{Z}$ and the regressor vector U are fed as inputs to MI , and the output $T(\varepsilon, U; \theta_{MI})$ is also a differentiable function represented by a multilayer perceptron with parameters θ_{MI} . $L(R, MI) = \mathbb{E}_{P_{\varepsilon U}} [T] - \log(\mathbb{E}_{P_{\varepsilon} P_U} [e^T])$ denotes the mutual information between U and ε . R is trained to minimize the dependency between residual and regressors. MI is simultaneously trained to tighten the gap between $I(U; \varepsilon)$ and $\hat{I}(U; \varepsilon)$ in order to achieve more accurate estimate of mutual information. In other words, R and MI play the following two-player minimax game:

$$\min_R \max_{MI} L(R, MI) = \mathbb{E}_{P_{\varepsilon U}} [T] - \log(\mathbb{E}_{P_{\varepsilon} P_U} [e^T]) \quad (7)$$

At equilibrium point, the value of loss $L(R, MI)$ is mutual information between U and ε . We provide experimental results in Section 5 that show convergence to the equilibrium point. In practice, the game in (7) is implemented by an iterative approach, in which the gradient of loss ∇L_B for mini-batch B is used via back-propagation procedure. As mentioned in (Belghazi et al. 2018), the second term in the mini-batch's gradient ∇L_B leads to a biased estimate of the full-batch gradient ∇L . To overcome this issue, Adam optimizer (Kingma and Ba 2014) can be utilized where the history of gradients is also considered in the next update.

Algorithm 1 shows AdOR training. In forward path, $2b$ examples are fed to R , and residuals $\varepsilon^{(i)}$ are computed in line 3. The first b pairs $\varepsilon^{(i)}$ and $u^{(i)}$ are jointly sampled; while, the second b pairs $\varepsilon^{(i+b)}$ and $u^{(i)}$ are marginal samples. Output of MI is computed twice: once by joint samples, and once by marginal samples in line 4. Finally, mini-batch loss L_B is computed in line 5 based on mean of samples computed in line 4. In backward path, parameters of each network are updated while the ones of other network is fixed. Note that in each iteration, R and MI are updated k_R and k_{MI} times, respectively.

4 Adversarial Orthogonal Structural Equation Model

In (6), the noise ε is modeled as an additive term. However, in general, the exogenous noise can affect the variable Z in a non-linear form, such as in structural equation models (SEM, see (Peters, Janzing, and Schölkopf 2017)). Thus, we assume here that the true model is: $Z = f(U, \varepsilon)$. In AdOSE, we propose a new method to estimate both the nonlinear

Algorithm 1: AdOR

for number of iterations **do**

Forward path:

1. Draw $2b$ minibatch samples
 $\{(u^{(1)}, z^{(1)}), \dots, (u^{(2b)}, z^{(2b)})\}$

2. Evaluate regression output

$$\hat{z}^{(i)} = \hat{f}(u^{(i)}; \theta_R); i = 1, \dots, 2b$$

3. Compute residual

$$\varepsilon^{(i)} = z^{(i)} - \hat{z}^{(i)}; i = 1, \dots, 2b$$

4. Evaluate output of MI twice

$$T^{(i)} = T(\varepsilon^{(i)}, u^{(i)}; \theta_{MI}); i = 1, \dots, b$$

$$T_{sh}^{(i)} = T(\varepsilon^{(i+b)}, u^{(i)}; \theta_{MI}); i = 1, \dots, b$$

5. Compute loss

$$L_B(\theta_R, \theta_{MI}) = \frac{1}{b} \sum_{i=1}^b T^{(i)} - \log\left(\frac{1}{b} \sum_{i=1}^b e^{T_{sh}^{(i)}}\right)$$

Backward path:

for k_R steps **do**

Update R by descending its stochastic gradient

$$\nabla_{\theta_R} L_B$$

end

for k_{MI} steps **do**

Update MI by ascending its stochastic gradient

$$\nabla_{\theta_{MI}} L_B$$

end

end

function f and also the joint distribution P_{UZ} . Hence, our goal is to obtain a function \hat{f} :

$$\hat{Z} = \hat{f}(U, \varepsilon), \quad (8)$$

such that \hat{Z} is similar as possible as to the response Z , with the same U ; i.e. $D_{KL}(P_{UZ} || P_{U\hat{Z}}) = 0$.

In AdOSE, similar to AdOR, the regression network (R) is pitted against the loss network: a KL-divergence estimator (KL) that learns to match the joint distribution $P_{U\hat{Z}}$ to distribution P_{UZ} (see the bottom diagram of Figure 1). Inspired by GAN (Goodfellow et al.(2014)), in AdOSE, the noise n_G is generated by a random Gaussian generator and transformed to the noise ε through a one-hidden layer perceptron $RanTrans$; i.e. $\varepsilon = RT(n_G)$. Then, regressors U and generated noise ε are passed to the regression network R , similar to AdOR; $\hat{Z} = \hat{f}(U, \varepsilon; \theta_R)$. Afterwards, pairs (U, Z) and (U, \hat{Z}) are passed through KL by a differentiable transformation T , and the outputs are $T(U, Z; \theta_{MI})$ and $T(U, \hat{Z}; \theta_{MI})$, respectively. Based on (2), the KL-distance is estimated by $L(R, KL) = \mathbb{E}_{P_{UZ}} [T] - \log(\mathbb{E}_{P_{U\hat{Z}}} [e^T])$, and two networks play the following minimax game:

$$\min_R \max_{KL} L(R, KL) = \mathbb{E}_{P_{UZ}} [T] - \log(\mathbb{E}_{P_{U\hat{Z}}} [e^T]). \quad (9)$$

At equilibrium, the value of loss $L(R, KL)$ is zero. After training, instead of having a nonlinear mapping between regressors and response, we have a nonlinear transformation for each samples of $U = u$, that assigns a distribution for

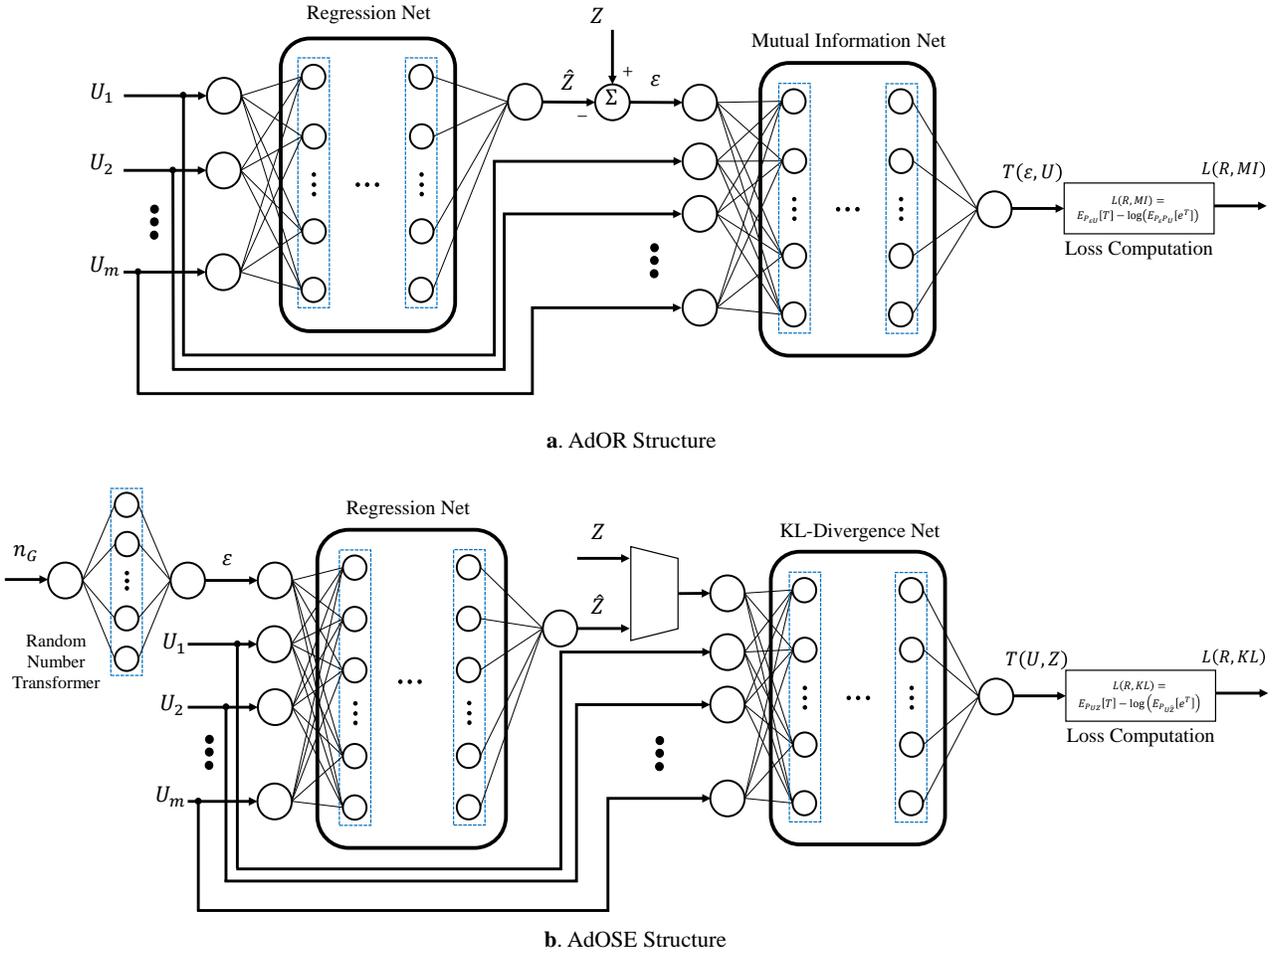


Figure 1: Block diagram of AdOR and AdOSE. **a.** AdOR structure: (U_1, \dots, U_m) are the input regressors, \hat{Z} is the predict, and ε is the residual. The output loss $L(R, MI)$ is the estimated mutual information between them. **b.** AdOSE structure: n_G is generated by Gaussian generator, ε is the exogenous noise, (U_1, \dots, U_m) and ε are fed as inputs to R . KL computes the output twice: once using (U_1, \dots, U_m, Z) , and once by $(U_1, \dots, U_m, \hat{Z})$. The output $L(R, KL)$ is the KL-Divergence.

Z ; i.e. $\hat{Z} \sim P(Z|U = u)$. Indeed, as the true value of n_G is unknown, we can not obtain single predict for each input sample u ; while, we can draw output samples by feeding different values of n_G . Since training AdOSE is more trickier than AdOR, we provide some implementation details in Section 5 to avoid divergence of the algorithm.

Algorithm 2 shows the training procedure of AdOSE. In forward path, b Gaussian samples are drawn and fed to *RanTrans*. The regression output is computed in line 3. As MI in AdOR, KL evaluates T twice: once by using $u^{(i)}$ and true responses $z^{(i)}$, and once by $u^{(i)}$ and predicted responses $\hat{z}^{(i)}$ (line 4). Mini-batch loss L_B is then computed using mean of true and estimated T . Similar to AdOR, in backward path, k_R and k_{KL} control the training of two networks. Furthermore, they play the main rule in convergence of the algorithm; if the loss is large, R has bad predicts and k_R should be increased, and if it is small, KL can not distinguish between true and predicted values and k_{KL} should

be increased.

Applications in Causal Inference

AdOR and AdOSE can be used in causal models that assume there is a structural model between child and parents. For instance, consider the additive noise model (ANM) between the cause variable C and the effect variable E : $E = f(C) + \varepsilon$. In (Hoyer et al. 2009), it has been shown that there exist no function g and noise $\tilde{\varepsilon}$ almost surely such that $C = g(E) + \tilde{\varepsilon}$ and E and $\tilde{\varepsilon}$ are independent. Hence, we can utilize AdOR to infer causal direction between two variables X and Y . To do so, we regress each variable on the other one and pick the direction with minimum loss $L(R, MI)$. Moreover, one can use AdOR as the class of functions for TiMINO (Peters, Janzing, and Schölkopf(2013)) for inferring causal direction in time series. At last, the causally conditioned DI (Quinn, Kiyavash, and Coleman 2015) of each child on each candidate parent can also be estimated by

regress the child twice, one on all variables, and the other on all variables except the candidate parent. The difference of two residuals’ entropy is DI from parent to child.

Algorithm 2: AdOSE

```

for number of iterations do
  Forward path:
  1. Generate  $b$  Gaussian samples  $\{n_G^{(1)}, \dots, n_G^{(b)}\}$ 
     Feed them to  $RanTrans$ :  $\varepsilon^{(i)} = RT(n_G^{(i)})$ 
  2. Draw  $b$  minibatch examples
      $\{(u^{(1)}, z^{(1)}), \dots, (u^{(b)}, z^{(b)})\}$ 
  3. Evaluate regression output
      $\hat{z}^{(i)} = \hat{f}(\varepsilon^{(i)}, u^{(i)}; \theta_R); i = 1, \dots, b$ 
  4. Evaluate output of  $KL$  twice
      $T^{(i)} = T(z^{(i)}, u^{(i)}; \theta_{KL}); i = 1, \dots, b$ 
      $T_{es}^{(i)} = T(\hat{z}^{(i)}, u^{(i)}; \theta_{KL}); i = 1, \dots, b$ 
  5. Compute loss
      $L_B(\theta_R, \theta_{KL}) = \frac{1}{b} \sum_{i=1}^b T^{(i)} - \log\left(\frac{1}{b} \sum_{i=1}^b e^{T_{es}^{(i)}}\right)$ 
  Backward path:
  for  $k_R$  steps do
    Update  $R$  and  $RanTrans$  by descending its
    stochastic gradient  $\nabla_{\theta_R} L_B$ 
  end
  for  $k_{KL}$  steps do
    Update  $KL$  by ascending its stochastic gradient
     $\nabla_{\theta_{KL}} L_B$ 
  end
end

```

5 Experiments

In this section, we first evaluate the performance of proposed regression methods on synthetic data and compare with the method in (Mooij et al. 2009) and some other nonlinear regression methods. Then, we apply the proposed method to find the causal direction in some real-world bilinear data (Mooij et al. 2016).

Implementation Details

The main point in training both AdOR and AdOSE is that the two networks R and MI (KL in AdOSE) should be trained simultaneously. As discussed before, Adam optimizer (Kingma and Ba 2014) is used, and all weights and biases initialized using Xavier initializer (Glorot and Bengio 2010). The number of layers, learning rate, and batch size are chosen similar in both networks.

In AdOR, we use three hidden layers with *tanh*, *sigmoid* and *leaky-ReLU* activation functions for R and three hidden layers with *leaky-ReLU* activation for MI . Note that adding a bias term to $\hat{f}(U)$ in (6) does not change mutual information, so bias term is removed from output layer of R . Similarly, adding a constant term to $T(\varepsilon, U; \theta_{MI})$ does not change the computed loss $L(R, MI)$ in (7), and we omit

the bias term from output layer of MI . Instead, the maximum mini-batch value $\max_{i=1, \dots, b} \{T^{(i)}, T_{sh}^{(i)}\}$ is reduced from whole $T^{(i)}$ and $T_{sh}^{(i)}$ in order to obtain a stable computation of loss.

The structure of AdOSE layers are designed similar to AdOR. The noise n_G is generated by normal Gaussian distribution, and $RanTrans$ has a hidden layer with *leaky-ReLU* activation. The bias term is added to the output layer of R , and biases in KL are similar to MI . Finding the stable solution of AdOSE is more trickier than AdOR. The optimizer might diverge in the first few iterations, because one of networks R or KL outstrips the other. To avoid this, we adjust steps k_R and k_{KL} by looking at the value of loss L_B in each iteration in order to stabilize the training procedure. A simple choice of steps has a linear feedback form $k_R = \lfloor a + bL_B \rfloor$ and $k_{KL} = \lfloor a - bL_B \rfloor$. We used $a = 30$ and $b = 10$ in our simulations.

Toy Examples

In this part, AdOR and AdOSE are compared with four regression methods: Support Vector Regression (Smola and Schölkopf 2004), neural network with same structure as AdOR with MSE loss minimization, HSIC regression proposed by Mooij et al.(2009), and Gaussian Process regression (Williams and Rasmussen 1996) with RBF kernel. The model has a simple form of $Y = f(X) + \varepsilon$. In each test, 300 samples are drawn from uniform distribution $X \sim U(-1, 1)$. The function $f(\cdot)$ is nonlinear and ε is generated from different non-Gaussian distributions. Note that for AdOSE, the averaged $\mathbb{E}_\varepsilon [Y|X = x]$ is plotted by feeding 5000 samples of n_G at each $X = x$. Figure 2 shows the output of different methods for the case of $f(X) = X^2$ and $\varepsilon \sim Exponential(1)$.

Comparison between methods is shown in Table 1 for different performance measures of Mean Squared Error (MSE), Mean Absolute Error (MAE) between predictions and responses, and Integral Squared Error (ISE) between estimated function and $f(x)$. As can be seen, in each case, AdOR has the worst MSE and MAE among the others; in contrast, its performance is much better in terms of ISE measure. In fact, we expect that AdOR/AdOSE do not have better performance in terms of MSE/MAE, compared to regression methods minimizing squared losses (or similar losses) since the goal of such methods is actually to minimize MSE while AdOR tries to minimize the mutual information between the residual and regressors. Moreover, it is not guaranteed that regression methods with square loss error estimate the underlying function statistically efficiently in cases other than Gaussian additive noise. In such cases, although the squared loss is minimized, the result might be dependent on the regressors. For instance, in Figure 2, in which the additive noise has an exponential distribution, it can be seen that AdOR finds the best approximation of the true function while the estimates given by NN-MSE and SVR are not close enough to it.

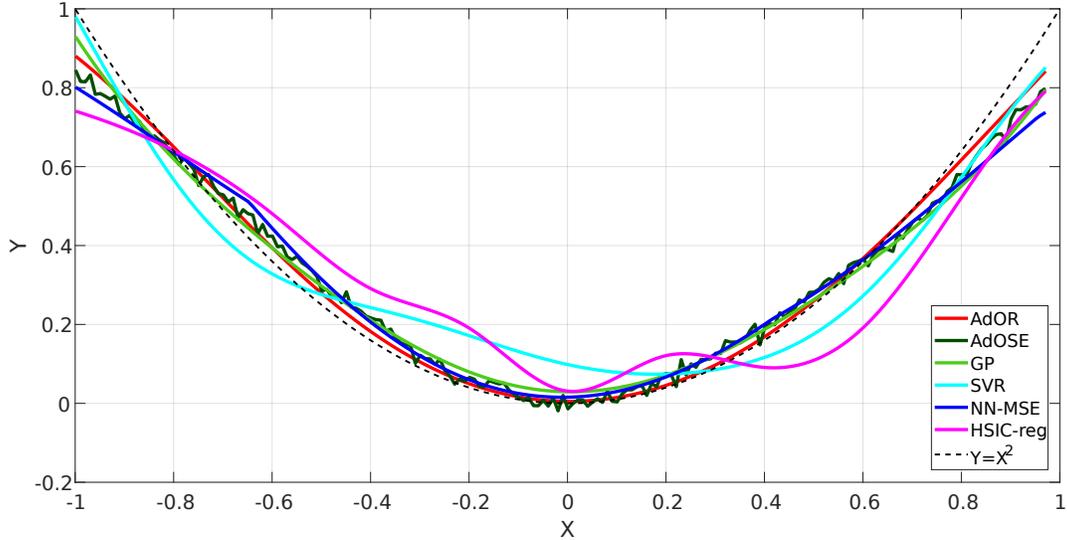


Figure 2: Example of different methods' outputs. $Y = X^2 + \varepsilon$ and $\varepsilon \sim \text{Exp}(1)$.

Table 1: Comparison of Different methods.

Model $f(X)$	X^2	$\sin(\pi X)$	e^{2X}	$\text{sigmoid}(5X)$	
	Noise $\varepsilon \sim \text{Exp}(1)$	$\varepsilon \sim \text{ChiSqr}(3)$	$\varepsilon \sim \text{Rayl}(4)$	$\varepsilon \sim \text{BioNom}(20, 0.3)$	
MSE	SVR	8.320e-01	5.651e+00	6.320e+00	3.849e+00
	HSIC-reg	8.419e-01	5.688e+00	6.386e+00	3.878e+00
	NN-MSE	8.373e-01	5.548e+00	6.226e+00	3.707e+00
	GP	8.262e-01	5.586e+00	6.228e+00	3.846e+00
	AdOSE	8.301e-01	7.658e+00	9.708e+00	5.112e+00
	AdOR	9.299e-01	9.740e+00	1.209e+01	4.073e+00
MAE	SVR	6.945e-01	1.926e+00	2.031e+00	1.555e+00
	HSIC-reg	7.049e-01	1.933e+00	2.051e+00	1.561e+00
	NN-MSE	7.061e-01	1.909e+00	2.037e+00	1.531e+00
	GP	6.975e-01	1.918e+00	2.035e+00	1.559e+00
	AdOSE	7.008e-01	2.071e+00	2.406e+00	1.527e+00
	AdOR	7.426e-01	2.492e+00	2.768e+00	1.626e+00
ISE	SVR	1.160e-02	1.363e-01	2.197e-01	1.669e-02
	HSIC-reg	2.666e-02	1.985e-01	2.551e-01	1.209e-01
	NN-MSE	8.430e-03	2.634e-01	1.078e-01	2.866e-01
	GP	5.247e-03	1.422e-01	3.491e-02	2.059e-02
	AdOSE	5.109e-03	6.633e-02	8.289e-02	1.554e-02
	AdOR	1.734e-03	4.974e-02	1.908e-02	2.232e-03

Distribution Estimation with AdOSE

Now suppose a non-additive model $Z = \varepsilon U$ with $U \sim \log\text{Normal}(1, 0.6) + 1$ and $\varepsilon \sim \text{Uniform}(-1, +1)$. 1000 number of samples are drawn from this model, and AdOSE is trained by these samples. Afterwards, 10^5 samples are drawn from the learned model by feeding different noise n_G for each value of $U = u$. The conditional distribution $P(\hat{Z}|U = u)$ is then estimated by naive binning for each u in the valid range. We also trained the model proposed by Sugiyama et al.(2010). True conditional distribu-

tion $P(Z|U = u)$ is depicted versus two estimated distributions $P(\hat{Z}|U = u)$ in Figure 3. As can be seen, the AdOSE has a great capacity to model distributions even in the regions with few samples.

Causal Direction Discovery in Real-World Datasets

Cause-effect (version 1.0) pairs (Mooij et al. 2016) is a collection of 108 real-world datasets, each with different sample size from 94 to 16382, where we considered 94 number of these datasets. Each dataset consists of samples of two

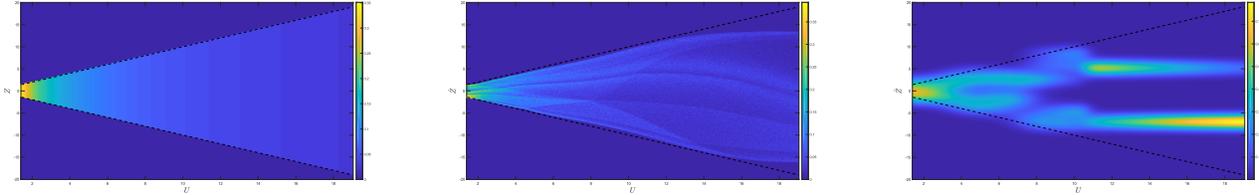


Figure 3: **Left:** True conditional distribution $P(Z|U)$. **Center:** Estimated conditional distribution $P(\hat{Z}|U)$ by AdOSE. **Right:** Estimated conditional distribution $P(\hat{Z}|U)$ by (Sugiyama et al. 2010).

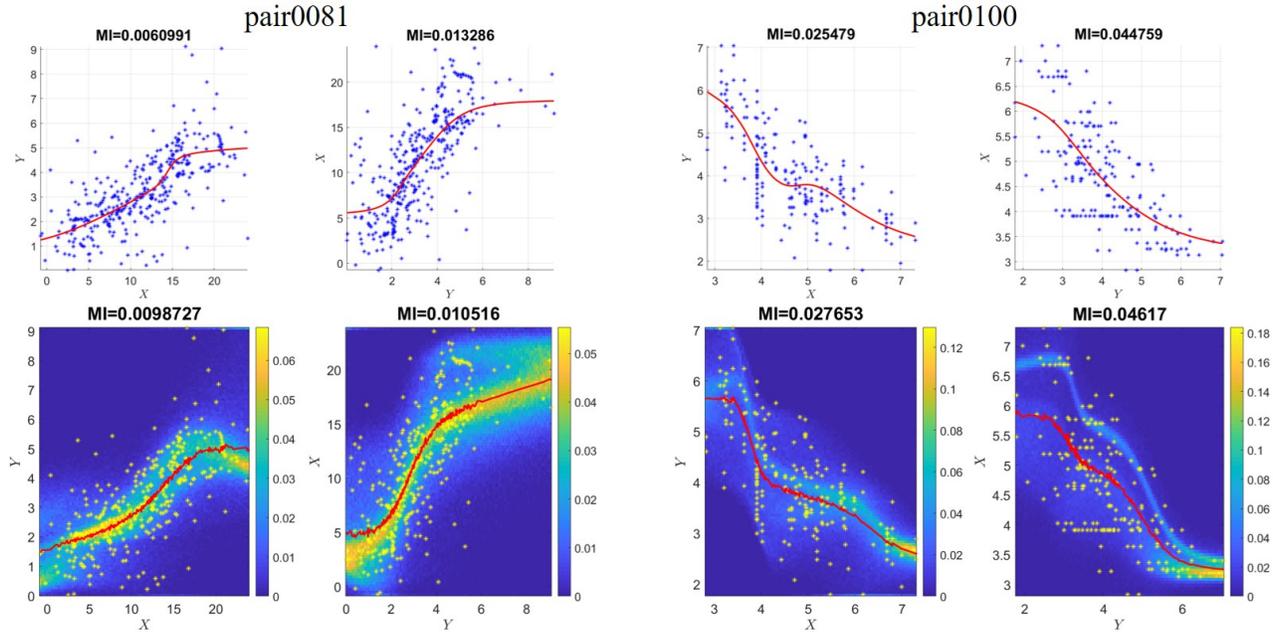


Figure 4: Results of proposed methods on real datasets. **Top row:** AdOR **Bottom row:** AdOSE

statistically dependent random variables X and Y , where one variable is known to causally influence the other. The task is to infer which variable is the cause and which one is the effect.

AdOR is trained with each dataset twice: once when Y is response and X is regressor and once in the reverse direction. The direction with lower mutual information is considered as the true direction. Experimental results show that we can infer the true direction for 66/94 (70.2%) fraction of datasets. We defined the score $S_i = MI(\varepsilon_{Y \rightarrow X}, Y) - MI(\varepsilon_{X \rightarrow Y}, X)$ for each dataset i . AdOR has AUPR (Area Under Precision-Recall Curve) of 78.68 based on the scores S_i and its performance is similar to the best AUPR achieved by the previous methods considered in (Goudet et al. 2017). In the same manner, AdOSE is trained twice in forward and reverse directions. The estimated function $\hat{f}(x) = \mathbb{E}_\varepsilon [Y|X=x]$ is computed by feeding 5000 samples of n_G at each $X=x$. AdOSE can infer the true direction for 63/94 (67.0%) fraction of datasets with AUPR of 74.32 based on the scores S_i . Figure 4 shows the estimated functions of AdOR and AdOSE on two pairs pair0081

and pair0100. The results for other datasets are given in the supplementary material.

6 Conclusions

We introduced two novel regression methods: AdOR which minimizes mutual information between the residual and the regressors, and AdOSE which produce response that mimics the true output by reducing distance between joint distributions. Conducted with details, we implemented our methods through adversarial neural networks and showed their great potential for inferring causal influences in models with unknown noise distributions. As a future work, one can extend these methods to the cases with categorical variables or utilize them in other causal learning problems such as learning causal structures.

References

Belghazi, M. I.; Baratin, A.; Rajeswar, S.; Ozair, S.; Bengio, Y.; Courville, A.; and Hjelm, R. D. 2018. Mine: mutual information neural estimation. *arXiv preprint arXiv:1801.04062*.

- Bonneau, R.; Reiss, D. J.; Shannon, P.; Facciotti, M.; Hood, L.; Baliga, N. S.; and Thorsson, V. 2006. The inferelator: an algorithm for learning parsimonious regulatory networks from systems-biology data sets de novo. *Genome biology* 7(5):R36.
- Darbellay, G. A., and Tichavsky, P. 2000. Independent component analysis through direct estimation of the mutual information. In *ICA*, volume 2000, 69–75.
- Donsker, M. D., and Varadhan, S. S. 1983. Asymptotic evaluation of certain markov process expectations for large time. iv. *Communications on Pure and Applied Mathematics* 36(2):183–212.
- Faith, J. J.; Hayete, B.; Thaden, J. T.; Mogno, I.; Wierzbowski, J.; Cottarel, G.; Kasif, S.; Collins, J. J.; and Gardner, T. S. 2007. Large-scale mapping and validation of escherichia coli transcriptional regulation from a compendium of expression profiles. *PLoS biology* 5(1):e8.
- Glorot, X., and Bengio, Y. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, 249–256.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.
- Goudet, O.; Kalainathan, D.; Caillou, P.; Guyon, I.; Lopez-Paz, D.; and Sebag, M. 2017. Causal generative neural networks. *arXiv preprint arXiv:1711.08936*.
- Granger, C. W. J. 1963. Economic processes involving feedback. *Information and control* 6(1):28–48.
- Granger, C. W. 1969. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society* 424–438.
- Haurly, A.-C.; Mordelet, F.; Vera-Licona, P.; and Vert, J.-P. 2012. Tigress: trustful inference of gene regulation using stability selection. *BMC systems biology* 6(1):145.
- Hausser, J., and Strimmer, K. 2009. Entropy inference and the james-stein estimator, with application to nonlinear gene association networks. *Journal of Machine Learning Research* 10(Jul):1469–1484.
- 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.
- James, W., and Stein, C. 1992. Estimation with quadratic loss. In *Breakthroughs in statistics*. Springer. 443–460.
- Jazayeri, M., and Afraz, A. 2017. Navigating the neural space in search of the neural code. *Neuron* 93(5):1003–1014.
- Kay, S. M. 1993. *Fundamentals of statistical signal processing*. Prentice Hall PTR.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kraskov, A.; Stögbauer, H.; and Grassberger, P. 2004. Estimating mutual information. *Physical review E* 69(6):066138.
- Liu, Y.; Aviyente, S.; and Al-khassaweneh, M. 2009. A high dimensional directed information estimation using data-dependent partitioning. In *2009 IEEE/SP 15th Workshop on Statistical Signal Processing*, 606–609. IEEE.
- Marbach, D.; Costello, J. C.; Küffner, R.; Vega, N. M.; Prill, R. J.; Camacho, D. M.; Allison, K. R.; Aderhold, A.; Bonneau, R.; Chen, Y.; et al. 2012. Wisdom of crowds for robust gene network inference. *Nature methods* 9(8):796.
- Marko, H. 1973. The bidirectional communication theory—a generalization of information theory. *IEEE Transactions on communications* 21(12):1345–1351.
- Miller, E. G. 2003. A new class of entropy estimators for multi-dimensional densities. In *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings (ICASSP'03)*, volume 3, III–297. IEEE.
- Mooij, J.; Janzing, D.; Peters, J.; and Schölkopf, B. 2009. Regression by dependence minimization and its application to causal inference in additive noise models. In *Proceedings of the 26th annual international conference on machine learning*, 745–752. ACM.
- 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.
- Murin, Y. 2017. k -nn estimation of directed information. *arXiv preprint arXiv:1711.08516*.
- Peters, J.; Janzing, D.; and Schölkopf, B. 2013. Causal inference on time series using restricted structural equation models. In *Advances in Neural Information Processing Systems*, 154–162.
- Peters, J.; Janzing, D.; and Schölkopf, B. 2017. *Elements of causal inference: foundations and learning algorithms*. MIT press.
- Quinn, C. J.; Kiyavash, N.; and Coleman, T. P. 2015. Directed information graphs. *IEEE Transactions on information theory* 61(12):6887–6909.
- Shadlen, M. N.; Britten, K. H.; Newsome, W. T.; and Movshon, J. A. 1996. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *Journal of Neuroscience* 16(4):1486–1510.
- 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.
- Smola, A. J., and Schölkopf, B. 2004. A tutorial on support vector regression. *Statistics and computing* 14(3):199–222.
- Sugiyama, M.; Takeuchi, I.; Suzuki, T.; Kanamori, T.; Hachiya, H.; and Okano, D. 2010. Least-squares conditional density estimation. *IEICE Transactions on Information and Systems* 93(3):583–594.
- Ver Steeg, G., and Galstyan, A. 2012. Information transfer in social media. In *Proceedings of the 21st international conference on World Wide Web*, 509–518. ACM.
- Ver Steeg, G., and Galstyan, A. 2013. Information-theoretic measures of influence based on content dynamics. In *Proceedings of the sixth ACM international conference on Web search and data mining*, 3–12. ACM.
- Wang, Z., and Bovik, A. C. 2009. Mean squared error: Love it or leave it? a new look at signal fidelity measures. *IEEE signal processing magazine* 26(1):98–117.
- Williams, C. K., and Rasmussen, C. E. 1996. Gaussian processes for regression. In *Advances in neural information processing systems*, 514–520.
- Zellner, A. 1988. Causality and causal laws in economics. *Journal of econometrics* 39(1-2):7–21.