

CAUSAL LINK ESTIMATION UNDER HIDDEN CONFOUNDING IN ECOLOGICAL TIME SERIES

Violeta Teodora Trifunov^{1,2}, Maha Shadaydeh¹, Jakob Runge², Veronika Eyring^{4,5}, Markus Reichstein^{3,6}, Joachim Denzler^{1,3}

Abstract—Understanding the causes of natural phenomena is a subject of continuous interest in many research fields such as climate and environmental science. We address the problem of recovering nonlinear causal relationships between time series of ecological variables in the presence of a hidden confounder. We suggest a deep learning approach with domain knowledge integration based on the Causal Effect Variational Autoencoder (CEVAE) which we extend and apply to ecological time series. We compare our method’s performance to that of vector autoregressive Granger Causality (VAR-GC) to emphasize its benefits.

I. INTRODUCTION

Many research fields such as climate and environmental sciences [1], [2] are continuously striving to understand the causes of natural phenomena. The complex nature and the continuously changing climate system both contribute to the slow advances in this field. This issue was shown to be amenable through the development of data-driven methodologies that are guided by theory to produce more accurate models [3]. We propose to mitigate the problem of recovering nonlinear causal relationships between time series of ecological variables in the presence of a hidden confounder. We suggest a deep learning approach with domain knowledge integration in the form of the ground-truth causal graph, shown in Fig. 1, for mending this issue. Our approach is based on the Causal Effect Variational Autoencoder (CEVAE) [4] which we extend by modelling an intervention for time series of confounded ecological

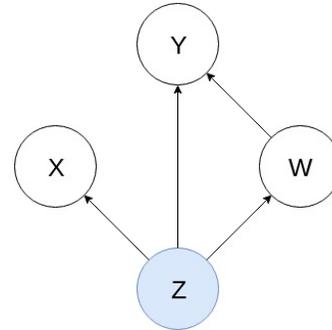


Fig. 1. Graphical model portraying hidden confounding with one proxy. Y denotes an outcome, W an intervention variable, Z an unobserved confounder and X denotes a proxy variable providing noisy views on the hidden confounder Z .

variables. We compare our method’s performance to that of the vector autoregressive Granger Causality (VAR-GC) [5], [6] and find that our approach is indeed capable of recovering nonlinear causal relationships under hidden confounding in contrast to VAR-GC baseline. In the study of ecosystems for example, considering confounders is important when trying to determine a causal link between the air temperature (T_{air}) and the ecosystem respiration (R_{eco}). Since both of these variables are influenced by the global radiation (R_g), one can not know with certainty that the causal link between T_{air} and R_{eco} is not affected by R_g . Therefore, not taking confounding into consideration may lead to erroneous conclusions. Two variables, W and Y , are said to be *confounded* if there exists another variable Z that is a cause for both W and Y . In order to confirm if the confounder is influencing the link between W and Y , one needs to intervene on W in the sense of *do*-Calculus [7] and thereby remove any influence of Z on W . If the intervention on W induces no change in the outcome Y , it is evident that the causal link between W and Y is solely influenced by the hidden confounder Z itself. In the case of an observed confounder, a conventional approach for accounting for its effect is to “control” for it. This is done for instance

Corresponding author: Violeta Teodora Trifunov, violetateodora.trifunov@uni-jena.de ¹Computer Vision Group, Friedrich Schiller University Jena, Jena, Germany ²Climate Informatics Group, Deutsches Zentrum für Luft- und Raumfahrt e.V. (DLR), Institute for Data Science, Jena, Germany ³Michael Stifel Center Jena for Data-Driven and Simulation Science, Jena, Germany ⁴Deutsches Zentrum für Luft- und Raumfahrt e.V. (DLR), Institut für Physik der Atmosphäre, Oberpfaffenhofen, Germany ⁵University of Bremen, Institute of Environmental Physics (IUP), Bremen, Germany ⁶Max Planck Institute for Biogeochemistry, Jena, Germany

by covariate-adjusted regression or propensity score regression [8]. However, if a confounder is hidden, it is impossible to estimate the effect of the intervention on the outcome without making further assumptions [7]. Figure 1 depicts a form of this problem when there is a single proxy variable. A proxy is an observed variable that describes the unobserved confounder and it is then used in causal link estimation between the confounded variables instead of the hidden confounder itself. For the cases of more general proxy models and the conditions under which they can be identified, please refer to [9].

II. RELATED WORK

A standard causality analysis method is Granger Causality (GC) [5] applied in the setting of no hidden confounding [10]. The main assumption of this concept is that causes always come before their effects in time. This means that if one time series causes another series, knowing the former series should be helpful for predicting future values of the latter series after influences of all other variables have been considered. Ecological variables often contain trends or periodic components such as diurnal or seasonal cycles which act as a hidden confounder. In [11], [12] authors have used parametric spectral representation for inferring the cause-effect relationships between ecological variables, assuming no hidden confounder. They have shown that time domain causality analysis of ecosystem variables based on VAR-GC [6], may result in spurious causal links due to the above-mentioned periodic components of ecological variables and thus proposed to use the parametric frequency domain representation of VAR-GC instead. In [13] it was further shown using a deep learning approach for causal inference based on a Causal Effect Variational Autoencoder (CEVAE) [4], that cause-effect analysis can be done in the presence of a periodic component acting as a hidden confounder in the time domain.

In regards to other deep learning methods for causal inference, several approaches have been suggested recently. One such is applied to inference of interactions between variables while learning the dynamics in an unsupervised manner [14]. Moreover, Causal Effect Network (CEN) [15] has been proposed for assessing causal relationships of time series, as well as their time delays between different processes. However, most causal inference methods cannot be applied if hidden confounders are present [2]. Another research branch dealing with the modelling of the latent variable space using deep graphical models was introduced in the recent years. Autoencoders are a class of deep learning

methods that can be combined with directed probabilistic graphical models for efficient inference in the presence of continuous latent variables with intractable posterior distributions, such as the Variational Autoencoder (VAE) [16]. Moreover, VAE represents the fundamental building block of the CEVAE [4], which allows for the estimation of the unknown latent space and inference of causal links between the confounded variables. Our work extends the capabilities of the CEVAE to time series that are based on real observations of global radiation and provides a comparison to VAR-GC.

III. METHODOLOGY

In the first part of this section, we describe the main deep graphical model our method relies on. Then, we explain how our method builds upon that graphical model. Finally, we provide a brief introduction to the VAR-GC method, which we use as a baseline.

A. Causal effect variational autoencoder

Based on the VAE [16] and the TARnet [17] generative model structure, the CEVAE [4] is a deep learning method that addresses hidden confounding by estimating the latent space and summarizing the causal effect of discrete or continuous, non-sequential variables. This is accomplished through the use of a noisy proxy related to the confounder, as shown in Fig. 1. In its original application to medical data, W from Fig. 1 denotes treatment, Y an outcome of the treatment, whereas a hidden confounder Z represents the socio-economic status of each patient. Its proxy X represents patient's income for the previous year and a place of residence. The main objective was, therefore, recovering the Individual Treatment Effect (ITE) and the Average Treatment Effect (ATE) defined in (1) and (2), respectively:

$$ITE(x) := \mathbb{E}(Y|X = x, do(W = w^1)) - \mathbb{E}(Y|X = x, do(W = w^0)) \quad (1)$$

$$ATE := \mathbb{E}(ITE(x)) \quad (2)$$

The metrics from Eq. 1 and 2 are defined for each individual value x of variable X , and by w^1 we denote the provided treatment, while w^0 denotes the values of W when no treatment is provided. To obtain the ITE, we need to recover the joint probability $p(Z, X, W, Y)$, as shown by Theorem 1 in [4]. Obtaining this joint distribution is done through a model network of the CEVAE by estimating the true posterior over Z which depends on X , W and Y , where Z is modelled by the standard normal distribution since it is unobserved and

an assumption about its distribution has to be made. The estimate of the posterior is then inferred via TARnet [17] by calculating it for each intervention group in W . It is then possible to construct a single objective for the inference and model networks, i.e. the *variational lower bound*

$$\mathcal{L} = \sum_{i=1}^N \mathbb{E}_{q(z_i|x_i, w_i, y_i)} \left(\log p(z_i) - \log q(z_i|x_i, w_i, y_i) + \log p(x_i, w_i|z_i) + \log p(y_i|w_i, z_i) \right), \quad (3)$$

of the graphical model from Fig. 1. By x_i we denote an input data point, by w_i each treatment assignment, by y_i the outcome of the specific treatment, by z_i the latent confounder and by q we denote estimates of the true probability distributions p which are computationally intractable. Finally, since it is necessary to know the intervention assignment w along with its outcome y before inferring the posterior distribution over Z , two auxiliary distributions are introduced, helping to predict w_i and y_i for new samples, so the variational lower bound becomes

$$\mathcal{F}_{\text{CEVAE}} = \mathcal{L} + \sum_{i=1}^N \left(\log q(w_i = w_i^* | x_i^*) + \log q(y_i = y_i^* | x_i^*, w_i^*) \right), \quad (4)$$

where x_i^* , w_i^* , y_i^* are, respectively, the observed values for the input, intervention and outcome variables in the training set.

B. CEVAE for ecological time series

Ecological time series often encompass nonlinearly related variables, as well as influence by a hidden confounder. In contrast to the conventional CEVAE setting, our intervention variable W , as well as variables X , Y and Z , are time series. Doing an intervention in real-world climate data is most often not feasible, so an alternative approach has to be followed. One such approach could be separating the values of an intervention variable W during day and night time or during winter and summer, for instance. As intervention we consider those values of W for which the influence of the confounder Z , the global radiation, is the weakest and thus arguably negligible. We provide more detail on modelling the intervention in Section IV. We also adjust several probability distributions to accommodate our problem setting. Namely, we model a conditional distribution of W given Z as follows:

$$p(W|Z) = \mathcal{N}(\mu_w, \sigma_w^2), \quad [\mu_w, \sigma_w] = f_1(Z). \quad (5)$$

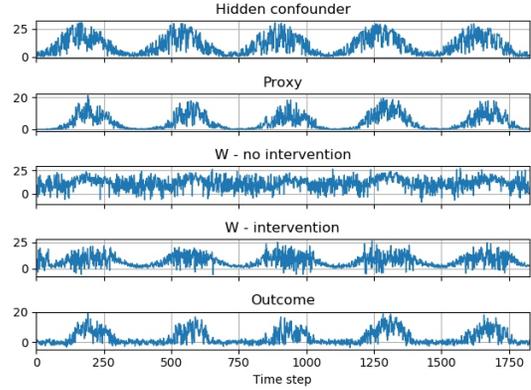


Fig. 2. Synthetic data. The first row shows real observation of global radiation R_{gobs} as the hidden confounder Z ; the second row shows the noisy GPP as the proxy X for $c = 0.03196$, $e = 0.3$, $\tau_3 = 23$ and $\tau_4 = 14$; the third row shows the variable T_{air} as W for $a = 0.4632$, $b = 0.48078$, $\tau_1 = 22$ and $\tau_2 = 17$ without intervention; the fourth row shows variable W with intervention; the fifth row shows R_{eco} as Y for $e = 0.198087$, $\alpha = 0.61078$, $\tau_5 = 21$ and $\tau_6 = 21$.

Estimation of this distribution is obtained through the use of the proxy X :

$$q(W|X) = \mathcal{N}(\hat{\mu}_w, \hat{\sigma}_w^2), \quad [\hat{\mu}_w, \hat{\sigma}_w] = f_2(X). \quad (6)$$

Functions f_1 and f_2 are feedforward neural networks with three layers. To measure the intervention effect of W on Y , we extend ITE (Eq. 1) to the case of a sequential intervention and define the Interval Intervention Effect (IIE) and the Average Intervention Effect (AIE):

$$\text{IIE}(I_i) := \mathbb{E}(Y|X \in I_i, do(W = w^1)) - \mathbb{E}(Y|X \in I_i, do(W = w^0)) \quad (7)$$

$$\text{AIE} := \mathbb{E}(\text{IIE}(I_i)_{i=0, \dots, m-1}) \quad (8)$$

By interval $I_i = [x_i, x_{i+1}]$ we denote the i -th uniform quantization level of X , with x_i and x_{i+1} being its limits, for $i = 0, \dots, m-1$ with $m = 256$. In Eq. (8), by $\text{IIE}(I_i)_{i=0, \dots, m-1}$ we denote an m -dimensional vector whose elements are $\text{IIE}(I_i)$ for each $i = 0, \dots, m-1$. In this manner we have extended the CEVAE to the setting of a continuous intervention variable, since intervals of W and Y are taken into account for calculating the causal effect between them, instead of the individual values, as seen in Eq. (1) and Eq. (2).

C. Vector autoregressive Granger causality

The main assumption of Granger causality (GC) [5] is that causes precede their effects and can be used for their prediction. Let u_i , $i = 1, \dots, N$ be the time series of N ecological variables. Each time series $u_i(t)$, $t =$

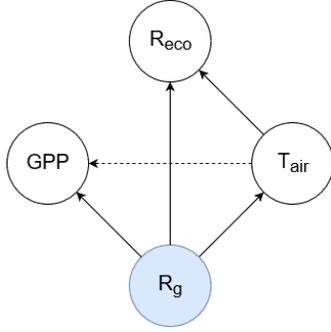


Fig. 3. Causal graphical model portraying causal relationships between ecological variables defined in equations (11)-(14). Respiration of the ecosystem R_{eco} denotes the outcome, air temperature T_{air} denotes an intervention variable, the global radiation R_g is assumed to be the hidden confounder and GPP denotes a proxy variable providing noisy views on the unobserved confounder R_g .

$1, \dots, k$ is a realization of length k real valued discrete stationary stochastic process $U_i, i = 1, \dots, N$. These N time series can be represented by a p th order vector autoregressive model (VAR(p)) of the form

$$\begin{bmatrix} u_1(t) \\ \vdots \\ u_N(t) \end{bmatrix} = \sum_{r=1}^p A_r \begin{bmatrix} u_1(t-r) \\ \vdots \\ u_N(t-r) \end{bmatrix} + \begin{bmatrix} \epsilon_1(t) \\ \vdots \\ \epsilon_N(t) \end{bmatrix}. \quad (9)$$

The residuals $\epsilon_i, i = 1, \dots, N$ form a white noise stationary process with covariance matrix Σ . The model parameters at time lags $r = 1, \dots, p$ comprise the matrix $A_r = [a_{ij}(r)]_{N \times N}$. Let Σ_j be the covariance matrix of the residual ϵ_j associated to u_j using the model in (9), and let Σ_j^{i-} denote the covariance matrix of this residual after excluding the i th row and column in A_r . The time domain VAR-GC of u_i on u_j conditioned on all other variables is defined by [6]

$$\gamma_{i \rightarrow j} = \ln \frac{|\Sigma_j^{i-}|}{|\Sigma_j|}. \quad (10)$$

IV. EXPERIMENTAL RESULTS

We have applied the proposed method to a synthetic data generated from 1825 real observations of the global radiation ($R_{g_{obs}}$) measured at the flux tower in Heinrich National Park - Germany, over the period of five years as suggested by [18]:

$$R_g(t) = R_{g_{obs}}(t) \quad (11)$$

$$T_{air}(t) = a \cdot T_{air}(t - \tau_1) + b \cdot R_g(t - \tau_2) + \eta_1(t) \quad (12)$$

$$GPP(t) = c \cdot R_g(t - \tau_3) \cdot T_{air}(t - \tau_4) + \eta_2(t) \quad (13)$$

$$R_{eco}(t) = e \cdot R_g(t - \tau_5) \cdot \alpha \frac{T_{air}(t - \tau_6)}{20} + \eta_3(t). \quad (14)$$

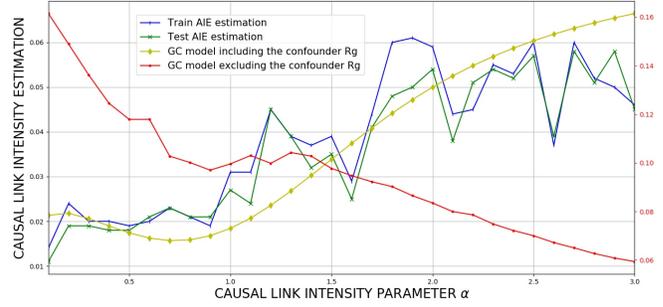


Fig. 4. Causal link estimation results of our method in comparison to the vector autoregressive Granger causality (VAR-GC). Blue and green curves show our methods estimation of the AIE during training and test, respectively. The yellow curve represents the VAR-GC method's estimate $\gamma_{T_{air} \rightarrow R_{eco}}$ when the confounding variable R_g is included along with T_{air} , R_{eco} and GPP. The red curve represents the VAR-GC method's estimate $\gamma_{T_{air} \rightarrow R_{eco}}$ when the proxy variable GPP is used along with T_{air} and R_{eco} , instead of the confounder itself. We note that in real ecological data $\alpha < 1$.

In equations (11)-(14), by η_1, η_2, η_3 we denote mutually uncorrelated Gaussian noise with mean $\mu = 0$ and variance $\sigma = 1$, by $\tau_i \in \mathbb{N}$ for $i = 1, \dots, 6$ we denote time lags, and by $a, b, c, e \in (0, 1)$ we denote constants. The plots of all above-mentioned variables along with the exact parameters used is shown in Fig. 2. The causal link intensity parameter is denoted by $\alpha \in \mathbb{R}$ and it represents the value of our highest interest. Namely, throughout our experiments, we increase the parameter α and observe the AIE estimation results in accordance to this increase. The portion of data we consider spans from April to September of each year, which results in 900 samples of real global radiation observations, corresponding to the hidden confounder Z . Synthetically generated variables T_{air} , GPP and R_{eco} , correspond to W , X and Y , respectively, as shown in Fig. 3. We note that even though there is a causal link from T_{air} to GPP, it does not change any of the original problem settings, nor any of the probability distributions involved. During those months the influence of T_{air} to R_{eco} should be more pronounced as the temperatures are higher. We model the intervention by considering a certain threshold of the proxy mean instead of the hidden confounder. More precisely, we consider values of the intervention variable T_{air} corresponding to the daily values of the proxy GPP that are smaller than the threshold of 0.7 of its mean. To create w^1 from Eq. (7), we use values of W corresponding to the proxy values smaller than the said threshold. This alone would lead to w^1 having missing values for time steps that correspond to the proxy values greater than the threshold. In order to prevent that, we concatenate all

present values and replicate them successively until the sample size is reached. We define w^0 from Eq. (7) in an analogue fashion. This type of intervention was chosen to simulate the properties of *do*-calculus as closely as possible. Moreover, it allows for a more easily adaptable application of our method to real data. Namely, after intervention, causal link between T_{air} and its parent R_g should either be removed or so small, that it can be neglected. This naturally occurs during winter or night time.

As far as the neural network architecture is concerned, we closely followed [4]. We used feedforward neural networks, namely f_1 and f_2 with 3 hidden layers and the ELU [19] nonlinearity. We note, however, that more hidden layers as well as different types of networks can be used. We modelled variable Z as normally distributed with 20 dimensions, due to its latency. We used a small weight decay term for all parameters with $\lambda = 0.0001$. For optimization, Adamax [20] was utilized with a learning rate of 0.01. Furthermore, early stopping according to the lower bound on a validation set was performed. For obtaining the outcomes $p(y|x_i \leq X \leq x_{i+1}, do(W = w^1))$ and $p(y|x_i \leq X \leq x_{i+1}, do(W = w^0))$ we averaged over 100 samples from the approximate posterior $q(Z|X) = \sum_w \int q(Z|w, y, X)q(y|w, X)q(w|X)dy$. As domain knowledge in this work, we consider the ground truth causal relationships between time series of ecological variables which are used to verify our method's results. When the causal graph is not partly or entirely known, methods like [1] should be used for obtaining it to a greater extent. Our goal is to estimate the causal link intensity between T_{air} and R_{eco} in the presence of the unobserved confounder R_g . We do so by running our method for different values of parameter α from Eq. (14), function of which is proportional to the causal link intensity between the variables in question. The results for each value of α from 0.1 to 3 are obtained as the average of the outputs of ten different realizations of the data. In real ecological time series, usually $\alpha < 1$. We note that the increase of α , yields the increase of the estimate of the causal link strength measured by the absolute value of AIE, as shown by the blue and green curves in Fig. 4. We also note that the curves describing the relation between the causal link intensity parameter α and the AIE estimation during training and testing is nonlinearly increasing. This means we are able to estimate the nonlinear causal relationship between T_{air} and R_{eco} from Eq. (14) under hidden confounding. Furthermore, for the purpose of a fair comparison, we applied the VAR-GC method to all

four variables $u_1 = T_{\text{air}}$, $u_2 = R_{\text{eco}}$, $u_3 = \text{GPP}$ and $u_4 = R_g$ from Eq. (9), over the entire 1825 data samples. More specifically, we included the otherwise hidden confounder R_g . This way, we could also reproduce the nonlinear causal link between T_{air} and R_{eco} , as seen in Fig. 4. To test if VAR-GC can detect the increase of the nonlinear causal link's intensity between T_{air} and R_{eco} without using R_g , we omitted it and observed as shown by a decreasing red curve in Fig. 4 that the desired causal link intensity increase could not be detected.

V. CONCLUSION AND FUTURE WORK

The goal of this work was to infer nonlinear causal relationships between time series of ecological variables in the presence of a hidden confounder using an extended version of the CEVAE. We provided a comparison to the baseline VAR-GC method with and without the use of a hidden confounder and concluded that our method is much more suitable for the task when the confounder is unobserved. This is since we use the proxy in estimating the AIE metric instead of the confounder, influence of which is removed after the intervention on T_{air} . We intend to enhance our method by doing a quantitative evaluation of the said comparison and incorporating recurrent neural networks into the current architecture.

REFERENCES

- [1] J. Runge, "Causal network reconstruction from time series: From theoretical assumptions to practical estimation," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 28, no. 7, p. 075310, 2018.
- [2] J. Runge, S. Bathiany, E. Bollt, G. Camps-Valls, D. Coumou, E. Deyle, and C. Glymour, "Inferring causation from time series of earth system sciences," *Nature Communications*, vol. 10, p. 2553, 2019.
- [3] J. H. Fahmous and V. Kumar, "A big data guide to understanding climate change: The case for theory-guided data science," in *Big Data*, vol. 2, pp. 155–163, 2014.
- [4] C. Louizos, U. Shalit, J. Mooij, Z. Sontag, D. R., and M. Welling, "Causal effect inference with deep latent-variable models," in *Advances in Neural Information Processing Systems 30*, pp. 6446–6456, 2017.
- [5] C. W. J. Granger, "Investigating causal relations by econometric models and cross-spectral methods," *Econometrica - Journal of the Econometric Society*, vol. 37, no. 3, pp. 424–438, 1969.
- [6] J. Geweke, "Measurement of linear dependence and feedback between multiple time series," in *Journal of the American statistical association*, vol. 77, pp. 304–313, 1982.
- [7] J. Pearl, *Causality*. Cambridge University Press, 2009.
- [8] L. Li, K. Kleinman, and M. W. Gillman, "A comparison of confounding adjustment methods with an application to early life determinants of childhood obesity," *Journal of developmental origins of health and disease*, vol. 5, no. 6, pp. 435–447, 2014.



- [9] W. Miao, Z. Geng, and E. Tchetgen Tchetgen, “Identifying causal effects with proxy variables of an unmeasured confounder,” in *arXiv preprint arXiv:1609.08816*, 2016.
- [10] M. Eichler, *Causal inference in time series analysis*, pp. 327–354. Wiley Series in Probability and Statistics, United States: John Wiley & Sons, 2012.
- [11] M. Shadaydeh, Y. Guanhe, M. Mahecha, M. Reichstein, and J. Denzler, “Causality analysis of ecological time series: a time-frequency approach,” in *Climate Informatics Workshop 2018*, 2018.
- [12] M. Shadaydeh, J. Denzler, Y. Guanhe, and M. Mahecha, “Time-frequency causal inference uncovers anomalous events in environmental systems,” in *German Conference on Pattern Recognition (GCPR)*, 2019.
- [13] V. T. Trifunov, M. Shadaydeh, J. Runge, V. Eyring, M. Reichstein, and J. Denzler, “Nonlinear causal link estimation under hidden confounding with an application to time series anomaly detection,” in *German Conference on Pattern Recognition (GCPR)*, 2019.
- [14] T. Kipf, E. Fetaya, K.-C. Wang, M. Welling, and R. Zemel, “Neural relational inference for interacting systems,” in *International Conference on Machine Learning 2018 (ICML)*, *arXiv:1802.04687v2 [stat.ML]*, 2018.
- [15] M. Kretschmer, D. Coumou, J. F. Donges, and J. Runge, “Using causal effect networks to analyze different arctic drivers of midlatitude winter circulation,” *Journal of Climate*, vol. 29, pp. 4069–4081, 2016.
- [16] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” in *Proceedings of the 2nd International Conference on Learning Representations (ICLR)*, *arXiv: 1312.6114 [stat.ML]*, 2014.
- [17] U. Shalit, F. Johansson, and D. Sontag, “Estimating individual treatment effect: generalization bounds and algorithms,” in *arXiv:1606.03976v5 [stat.ML]*, 2016.
- [18] C. Krich, M. Mahecha, J. Runge, M. Reichstein, and D. G. Miralles, “Revealing causal dependencies between land-surface fluxes and meteorological variables,” *Geophysical Research Abstracts*, vol. 20, 2018.
- [19] D.-A. Clevert, T. Unterthiner, and S. Hochreiter, “Fast and accurate deep network learning by exponential linear units (elus),” *arXiv:1511.07289 [cs.LG]*, 2016.
- [20] D. P. Kingma and J. B. Adam, “A method for stochastic optimization,” *International Conference on Learning Representations (ICLR)*, 2015.