



Counterfactual-Based Root Cause Analysis for Dynamical Systems

Juliane Weibach^{1,2}(✉), Sebastian Gerwinn¹, Karim Barsim¹,
and Martin Fränzle²

¹ Bosch Center for Artificial Intelligence, Renningen, Germany
juliane.weibach@de.bosch.com

² Carl von Ossietzky University of Oldenburg, Oldenburg, Germany

Abstract. Identifying the underlying reason for a failing dynamic process or otherwise anomalous observation is a fundamental challenge, yet has numerous industrial applications. Identifying the failure-causing subsystem using causal inference, one can ask the question: “Would the observed failure also occur, if we had replaced the behaviour of a subsystem at a certain point in time with its *normal* behaviour?” To this end, a formal description of behaviour of the full system is needed in which such counterfactual questions can be answered. However, existing causal methods for root cause identification are typically limited to static settings and focusing on additive external influences causing failures rather than structural influences. In this paper, we address these problems by modelling the dynamic causal system using a Residual Neural Network and deriving corresponding counterfactual distributions over trajectories. We show quantitatively that more root causes are identified when an intervention is performed on the structural equation and the external influence, compared to an intervention on the external influence only. By employing an efficient approximation to a corresponding Shapley value, we also obtain a ranking between the different subsystems at different points in time being responsible for an observed failure, which is applicable in settings with large number of variables. We illustrate the effectiveness of the proposed method on a benchmark dynamic system as well as on a real world river dataset.

Keywords: Dynamic Root Cause Analysis · Counterfactual Inference · Dynamic Systems

1 Introduction

Explaining unexpected behaviour in terms of underlying causes is a difficult challenge with a broad range of applications. Such applications range from identifying

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/978-3-031-70365-2_18.

potential problems in industrial processes to understanding influencing factors in anomalous weather phenomena. For example, within an assembly line of an industrial manufacturing plant, faster identification of root causes of increased scrap rate (the rate at which assembled products fail quality assessment audits) can minimize cost, increase production yield, and increase overall efficiency. If one can observe sufficiently many instances of anomalous behaviour or of faulty traces of a process, one option would be to perform correlation based analysis or causal discovery [14], thereby estimating the influencing factors to the variable “fault” [2,9,17]. Alternatively, causal inference can be used even if only a single anomalous observation is available [5,15]. Here, the identification of root causes is formulated in terms of a counterfactual query: “Would the observed failure also occur, if we had replaced the behaviour of a sub-system at a certain point in time with its *normal* behaviour?”. Although such a causal inference approach can estimate a ranked score of each variable involved of being the underlying root cause, we address three main shortcomings of this approach in this paper:

Static Systems: Root cause analysis based on causal inference has been considered only in static environments [3,5,15]. To address this limitation, we fit a time-discretized version of an Ordinary Differential Equation (ODE) system, thereby obtaining a dynamic model. By deriving counterfactual distributions over trajectories we then employ similar strategies as in the static case.

Structural Influences: Existing causal inference methods using counterfactuals [5,15] focus on additive external influences causing failures rather than structural influences. While [2] also considers structural influences, the method is limited to linear models and does not include single time external influences. In this paper, we address this problem by allowing for interventions on the structural equation and the external influence.

Non-linear Systems: Existing methods for root cause analysis are typically limited to linear dynamic models. Here, we address this problem by allowing transition functions to be non-linear using a simple neural network architecture. Additionally, existing methods are limited to small systems as they rely on the computation of Shapley values, which scales exponentially with the number of variables. This becomes infeasible in a dynamic setting, since the corresponding causal graph – unrolled over time – would have an increasingly large number of nodes. While approximate methods for the computation of Shapley values have been proposed [10], we suggest a simple approximation to the Shapley value, which is applicable in settings with large number of variables.

The remainder of this paper is organized as follows: in Sect. 2, we review the related work mentioned above in more detail, and provide necessary background and notation in Sect. 3. In Sect. 4, we describe our method for identifying root causes. In Sect. 5, we first illustrate the mechanisms of the proposed method in a synthetic linear and non-linear setting before evaluating it on a benchmark

dataset of [2] as well as on real data describing river levels as in [5]. In Sect. 6, we conclude the paper.

2 Related Work

The problem of identifying the root cause of a system failure or anomaly has been addressed in various domains, including healthcare [15], financial income distributions [4], reliability engineering [9], to name a few. In the context of time-series data, causal inference techniques have been used to qualitatively explain outliers using counterfactual trajectories [16]. To detect root causes affecting graphical structure or transition function Assad et al. [2] propose a method based on assessing the direct causal effect. Modelling such causal effect with linear models, Assad et al. [2] show that the total effects change if the underlying causal model changes. In turn, they can use this fact to identify structural changes in the causal model. However, the method is limited to linear models and does not include single time external influences. If more observations of anomalous data are available, the problem of identifying the root causes is also amenable to statistically estimate the correlation or causation of the different variables and time points onto the variable associated with the label “anomalous”. To this end, Tonekaboni et al. [17] introduce feature importance in time (FIT), a scoring mechanism to quantify importance of features in a multi-variate time-series. The authors propose to assess feature importance based on their predictive power w.r.t. the outcome distribution, while accounting for temporal distributional shifts. The approach localizes important features over time and can thus be used to gain useful insights into the behaviour of dynamic systems. However, FIT does not leverage the causal structure of the underlying system and rather provides correlative explanations for the observed outcomes.

Best aligned with our approach, though, is the work by [5] which defines the problem of identifying root causes of a system failure as a counterfactual query. With this reformulation, the authors claim to be the first to propose actionable explanations to anomalous behaviour of underlying systems. In principle, counterfactual reasoning assumes, and leverages, complete causal knowledge of the underlying system in the form of a structural causal model (SCM). More precisely, the work in [5] assumes *invertible* functional causal models: models in which exogenous variables are computable from endogenous system observations. In fact, the authors leverage the default split between endogenous and exogenous variables in a graphical causal model to disentangle a node’s inherited impact from its own contribution. They account for the notion of graded causation [7] and provide order-independent feature scoring using a game-theoretic concept commonly adopted in explainable machine learning [10], namely Shapley values [13]. With its computation complexity, their approach lacks direct applicability to dynamical systems. In our experiments, we compare against the linear model performing interventions on the exogenous variable analogously to [5].

3 Background and Notation

As mentioned in the introduction, we are interested in a counterfactual approach to identify the root cause of a system failure. In this section, we introduce the necessary concept from the literature and also introduce the notation we use throughout the paper. Following the notation from Peters et al. [12], we denote the sequence of observations of the system of interest by d -variate time series $(\mathbf{Y}_t)_{t \in \mathbb{Z}}$ where each \mathbf{Y}_t for fixed t is the vector (Y_t^1, \dots, Y_t^d) . Each Y_t^j represents the j th observable of a system at time t . By some abuse of notation, if we omit super- or subscripts, we refer to the full time series. That is, $\mathbf{Y} = (\mathbf{Y}_t)_{t \in \mathbb{Z}}$, $\mathbf{Y}^j = (Y_t^j)_{t \in \mathbb{Z}}$ and $\mathbf{Y}_t = (Y_t^j)_{j \in \{1, \dots, d\}}$. The full time causal graph \mathcal{G}_t with a node for each time point and signal Y_t^j for $(j, t) \in 1, \dots, d \times \mathbb{Z}$ has theoretically infinitely many nodes and is assumed to be acyclic, while the summary graph \mathcal{G} with nodes Y^1, \dots, Y^d may be cyclic.

Definition 1. (*Structural causal model (SCM)*) [12]

An SCM $\mathcal{M}(\mathcal{S}, P_N, \mathcal{G})$ is defined by a set of structural equations \mathcal{S} , an acyclic graph $\mathcal{G} = (\mathcal{Y}, \mathcal{E})$, and a set of independent noise variables $N^j \sim P_{N^j}, j \in \mathcal{Y}$. The structural equations for each node j are given by:

$$S^j := Y^j = f^j(Y^{PA(j)^{\mathcal{G}}}, N^j)$$

where $\mathcal{S} = \cup_{j \in \mathcal{E}} \{S^j\}$ is a set of structural collections, and $PA(j)^{\mathcal{G}} \subseteq \mathcal{E}$ denotes the parents of the node j according to the graph \mathcal{G} .

To describe dynamic processes, again following [12], we extend the above definition to the dynamic case by unrolling a causal graph over time as follows:

Definition 2. (*Dynamic SCM*)

In analogy to a static SCM, a dynamic SCM $\mathcal{M}(\mathcal{S}_t, P_{N_t}, \mathcal{G}_t)$ is given by an acyclic graph \mathcal{G}_t and exogenous noise influences $N_t^j \sim P_{N_t^j}$ independent over each point in time t and variable j . \mathcal{G}_t is referring to a graph consisting of an unrolled version of a summary graph \mathcal{G} . Following the notation of [12], the structural equations for node Y_t^j are given by:

$$S_t^j := Y_t^j = Y_{t-1}^j + f^j(Y_{t-1}^{PA(j)}, Y_{t-1}^j) + N_t^j$$

with $PA(j)$ being the parents of node j according to the summary graph \mathcal{G} excluding the node itself. A notable difference from static SCMs is that the functional coupling f is constant over time¹.

Definition 3. (*Interventional Dynamic SCM*) Let \mathcal{J} be a set of interventions in which each element ξ can be of the following form:

¹ Note that we restrict ourselves to additive noise in order to realize an *invertible* SCM, see [5].

$$\xi := do(P_{N_t^j}) = \tilde{P}_{N_t^j}, \quad (1) \quad \text{or} \quad \xi := do(S_t^j) = \tilde{S}_t^j \quad (2)$$

where $\tilde{P}_{N_t^j}$ is a new noise distribution and \tilde{S}_t^j is a new structural equation for the node j at time t . The interventional dynamic SCM is then defined by replacing either the noise distribution or structural equation within a given dynamic SCM $\mathcal{M}(S_t, P_{N_t}, \mathcal{G}_t)$. Here Eq. 1 denotes a soft intervention on the noise distribution whereas Eq. 2 denotes an intervention on the structural intervention. We denote the resulting intervened dynamic SCM then by $M_{\mathcal{J}}(S_t, P_{N_t}, G_t)$.

As each SCM (interventional or not) defines structural equations and noise distributions, it can generate a trajectory of observations. We denote the distribution of the observations generated by the SCM as $P_{\mathcal{M}}$ and the distribution of the observations generated by the intervened SCM as $P_{\mathcal{M}_{\mathcal{J}}}$. Given an observed trajectory, we can now also define the counterfactual distribution describing hypothetical trajectories which would have been observed if an (alternative) intervention had been performed.

Abducted and Counterfactual SCMs. Let \mathbf{Y}^F be an observed trajectory and \mathcal{M} a given dynamic SCM. In order to construct a counterfactual dynamic SCM, we define the noise posterior distribution $P_{N_t^j}(N_t^j | \mathbf{Y}^F) = \delta(N_t^j - N_t^{F,j})$ by:

$$N_t^{F,j} = -Y_{t-1}^{F,j} - f^j(Y_{t-1}^{F,PA(j)}, Y_{t-1}^{F,j}) + Y_t^{F,j} \quad (3)$$

where f^j is the structural equation of the node j and $PA(j)$ are the parents of the node j according to the summary graph \mathcal{G} . The resulting dynamic SCM, in which the noise distributions $P_{N_t^j}$ are replaced with the above defined noise posterior distributions, is then denoted as \mathcal{M}^F indicating that the noise distributions are abducted from the observed trajectory \mathbf{Y}^F . In fact, when generating trajectories from this abducted SCM, it only generates the observed trajectory \mathbf{Y}^F due to the above setting of the noise variables. In order to generate new counterfactual trajectories reflecting alternative outcomes, we need to perform an intervention on this abducted SCM, leading to the counterfactual SCM. That is, given an abducted SCM \mathcal{M}^F and a set of interventions \mathcal{J} , we refer to the resulting interventional SCM $\mathcal{M}_{\mathcal{J}}^F$ as the counterfactual SCM. For example, when performing an intervention $do(P_{N_t^j}) = \tilde{P}_{N_t^j}$ on the noise distribution at a specific point in time t and a node j , the counterfactual SCM is defined by the following structural equations:

$$Y_s^e = Y_{s-1}^e + f^e(Y_{s-1}^{PA(e)}, Y_{s-1}^e) + N_s^e, \quad \text{where} \quad (4)$$

$$N_s^e \sim \begin{cases} \tilde{P}_{N_t^j} & \text{if } s = t \text{ and } e = j \\ \delta(N_t^j - N_t^{F,j}) & \text{otherwise} \end{cases} \quad (5)$$

3.1 Root Cause

As we are interested in identifying a root cause, we state here more precisely what we mean by this term. We define a root cause as an intervention according to $M_{\mathcal{J}}(S_t, P_{N_t}, G_t)$ leading to a faulty behaviour. Here, we assume that a faulty behaviour can be detected or defined using a known classifier ϕ . This classifier maps a time series to a binary value, indicating whether the time series is faulty. Such classifier can either be given as a known test function (e.g. corresponding to an end-of-line test in an assembly line, an assertion in a software system, or a medical diagnosis) or can be learned from data (e.g. an outlier-score function learned on normal data).

Definition 4. (*Root cause*) Given a classifier ϕ that determines whether an observed trajectory is faulty, we refer to a (set of) intervention(s) Ξ to be the root cause of a failure associated with the classifier ϕ , if observations $(\mathbf{Y}_{t,t=1,\dots,T}^F)_j$ from the interventional SCM $\mathcal{M}_{\{\Xi\}}$ are leading to an increased failure rate:

$$\mathbb{E}_{\mathbf{Y}_{t,t=1,\dots,T}^F \sim \mathcal{M}_{\Xi}}[\phi(\mathbf{Y}_{t,t=1,\dots,T}^F)] - \mathbb{E}_{\mathbf{Y}_{t,t=1,\dots,T} \sim \mathcal{M}}[\phi(\mathbf{Y}_{t,t=1,\dots,T})] > 0$$

Note that this corresponds to the average treatment effect of an intervention on the external influence or structural intervention. If the probability of a failure for an external intervention on the noise or structure is larger than without any intervention, we assume that the failure has an underlying root cause.

3.2 Shapley Value

Shapley values, originally defined to quantify the contribution of individual players to the outcome of a game, have been used by Budhathoki et al. [5] in a static setup to define a score for nodes being potential root causes of an observed fault. To this end, interventions (or possible root causes) are identified with players in a game whose outcome is determined by a value function that quantifies the degree to which a set of interventions can increase the likelihood of correcting a failure (to be defined below).

Definition 5. (*Shapley value*) The Shapley value [13] of a player i out of a set N of possible players to the outcome of a game characterized by the outcome function v is defined by:

$$Sh(i) := \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(n - |S| - 1)!}{n!} (v(S \cup \{i\}) - v(S))$$

Note that in order to calculate the Shapley value, one has to sum over exponentially many subsets of the set of possible players. This is feasible only for small sets of players. As in the context of root cause analysis in a dynamic setting, the set of players corresponds to the set of possible interventions ranging over all possible times and nodes within the unrolled graph of a dynamic SCM. Due to the exponential growth of the number of possible interventions, exact Shapley value estimation is computationally infeasible for dynamic SCMs, and we have to resort to an approximate version.

4 Method for Identifying Root Causes

Now that we have the necessary background, we can describe our method for identifying root causes in dynamic SCMs. The method is based on the following steps and is illustrated in Fig. 1. We want to identify the root cause that caused an observed failure in a system. To this end, we cast this problem in a counterfactual query: “Would the observed failure also occur if we had replaced the *faulty* behaviour of a sub-system at a certain point in time with its *normal* behaviour?”. To answer this question after we observed a faulty observation $(\mathbf{Y}_{t,t=1,\dots,T}^F)_j$, as illustrated in *Inputs* in Fig. 1, we follow the steps of counterfactual distribution calculation: abduction, action and prediction [11]. However, in order to apply those steps, we need an SCM characterizing the normal and potentially the abnormal system. To characterize the normal system, we assume to have access to data representing the normal behaviour of the system, as shown in *Inputs* in Fig. 1. Additionally, we assume to have at least a summary graph \mathcal{G} of the system. This summary graph can be obtained from expert knowledge

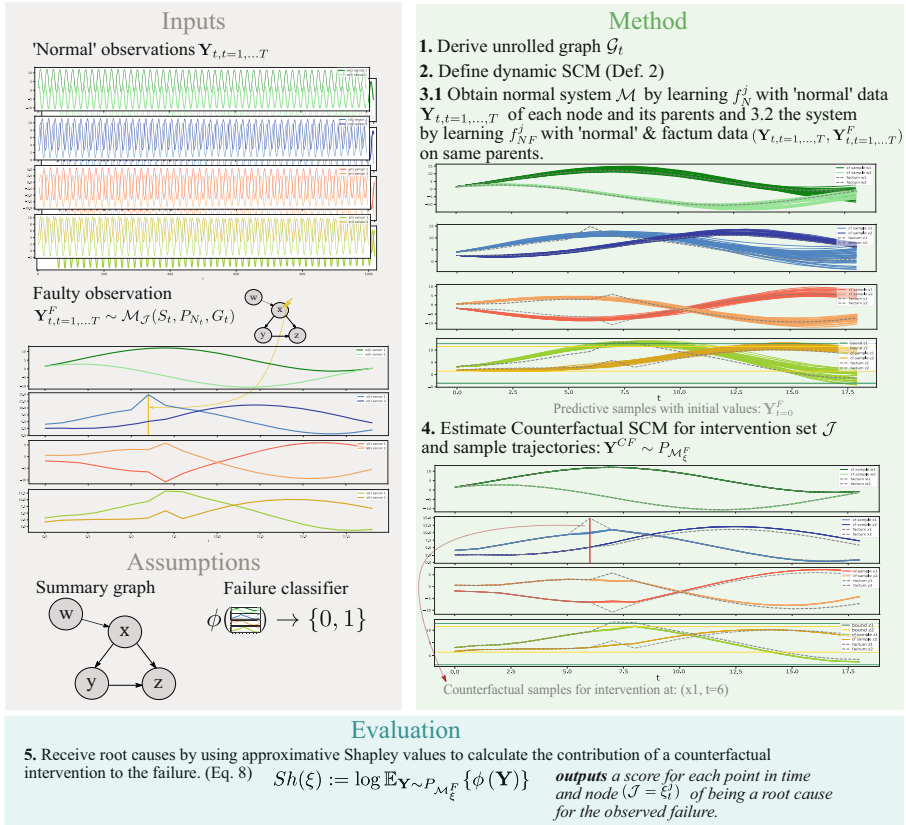


Fig. 1. This figure shows an overview of the individual steps of our method.

or from data. Furthermore, as shown in *Assumptions* in Fig. 1, we assume that we know a function ϕ that classifies an observation into faulty or normal.

Fitting the model in step 3.1 of the *Method* part in Fig. 1 we obtain the normal behaviour system \mathcal{M} by learning the functions f_N^j with the inputs being normal observations Y_t of each node and its parents of the summary graph \mathcal{G} . If for both, normal as well as abnormal data, a node and hence its transition function is not anomalous, the transition function would be identical for both settings. Therefore, in 3.2 we additionally fit a transition function f_{NF}^j with normal and factum data as input on the same parents and children as in 3.1 of the known graph \mathcal{G} and with that we define the SCM \mathcal{FM} . We show predictive samples of \mathcal{M} in the graph under 3.2.

Estimating the Counterfactual: In the abduction step, we first infer the noise distribution corresponding to the observed factum. We refer to the abducted SCMs \mathcal{M}^F and \mathcal{FM}^F by applying the factum as function input to f_N^j and f_{NF}^j and constructing the resulting noise posterior distributions as described in Eq. 3. We need to calculate the noise variables for both SCMs separately, because function couplings and noise variables are coupled. In the action step, we perform an intervention in \mathcal{M} by $\xi_{\mathcal{M}} := \{do(P_{N_t^j}) = \tilde{P}_{N_t^j}\}$ (see Eq. 1), where we use the prediction error of our model to estimate the Gaussian noise variance:

$$\tilde{P}_{N_t^j} = \mathcal{N}(0, \sigma_{val}^2), \quad \sigma_{val}^2 = \frac{1}{V} \frac{1}{T} \sum_v \sum_t (Y_{t+1}^{j,v} - f_j(Y_t^{Pa(j),v}))^2 \quad (6)$$

with $Y^{j,v}$ being a validation trajectory of the normal data, and V the number of validation trajectories. For an intervention in \mathcal{FM} we intervene on the noise as before and we additionally intervene on the structure by $do(S_t^j) = \tilde{S}_t^j$ (see Eq. 2), which replaces the previous transition function f_{FN}^j with a new structural equation \tilde{S}_t^j consisting of the transition function f_N^j originating from the "normal" SCM, obtained purely from training data $\xi_{\mathcal{FM}} := \{do(P_{N_t^j}) = \tilde{P}_{N_t^j}, do(S_t^j) = \tilde{S}_t^j\}$. After the construction of the corresponding counterfactual SCM we can then generate counterfactual trajectories under the different interventions $\mathbf{Y}^{CF} \sim P_{\mathcal{M}_{\xi_{\mathcal{M}}^F}}$, as illustrated in 4. in Fig. 1. If an external influence on node j at time t leads to an abnormal factum, an intervention of the above type should remove the abnormal behaviour and therefore lead to a normal trajectory.

Evaluation: To quantify how close these counterfactual samples are to normal trajectories, the trajectories are processed via a classifier function ϕ (Eq. 4). In turn, we receive a score for each counterfactual sample indicating whether the failure was removed by the counterfactual intervention ξ . We then average over multiple counterfactual samples. To rank interventions at different times and nodes, we can use Shapley values by identifying players with interventions and match-outcomes by the average normality of the counterfactual sample. Shapley values, however, scale exponentially and therefore we use the following simple approximation, which we obtain by ignoring interactions between different interventions, thereby only considering singleton intervention sets. Although mainly

motivated by pure computational tractability, we can alternatively assume that perfectly synchronous occurrence of multiple root causes is very unlikely, thereby justifying the restriction to singleton intervention sets. Consequently we arrive the following simple expression of contribution score of individual interventions ξ for each point in time and node:

$$Sh(\xi) := \log \mathbb{E}_{\mathbf{Y} \sim P_{\mathcal{M}_\xi^F}} \{\phi(\mathbf{Y})\} \quad (7)$$

5 Experiments

In the following experiments, we evaluate the effectiveness of the proposed method for different synthetic and real world data-sets. As for synthetic data-sets, we consider both linear and non-linear dynamic systems with single point external failure-causing disturbances as well as a benchmark data-set for identifying structural causes for anomalies [2]. As for the real-world data-set, we are investigating dynamic water flow rate in rivers [1]. For our synthetic experiments, we perform two meta-experiments which analyze the influence on the model performance of varying root cause injections and how robust the model is against violating the assumption that the causal graph is known. We denote our models, a linear and a non-linear model both performing a counterfactual intervention on the external noise influence and on the structural equation with $Lin(S_t^j, N_t^j)$ and $NLin(S_t^j, N_t^j)$. We compare against a linear layer model with counterfactual noise-influence intervention $Lin(N_t^j)$, similar to [5] as well as against EasyRCA [2] in the benchmarking experiment. For completeness, we additionally provide a nonlinear model $NLin(N_t^j)$ with counterfactual noise-influence intervention. In order to model the non-linear dynamic SCM, for $NLin$ we use a simple three-layer residual neural network (ResNet) with hyperbolic tangent activation functions and 128 neurons as latent layer.

5.1 Experimental Datasets

Linear Synthetic System: In our first data-set, we consider a linear multivariate system with additive Gaussian noise consisting of four nodes (w, x, y, z) , each having two dimensions. The summary graph of the system is shown in *Assumptions* in Fig. 1. The structural equations of the system are of the form:

$$Y_t^j := A^i \mathbf{Y}_{t-1}^j + \sum_{k \in PA(j)} B^k \mathbf{Y}_{t-1}^k + C^l N_t^j, \quad (N_t^j)_d \sim \mathcal{N}(0, 1) \quad \forall d$$

with N_t^j being zero mean standard Gaussian noise. For this system we chose the transition matrices such that they generate a stable system by using eigenvalues smaller than 1 (see Appendix). To simulate a root cause, we inject an additive constant term at a single dimension of a node j at time t to the equation above. Instead of a learned anomaly scoring function, in this experiment, we assume

to have access to a function that checks the validity of a given observation, similarly as it would be in a manufacturing scenario, in which an end of line test is performed [6]. Therefore, we examine if a failure on the “last” node in a manufacturing line (here “last” node in the summary graph is z) has occurred. To this end, we use a threshold function, fixed over time for each dimension of node z . More precisely, this classifier can be applied to any time-series observation $(\mathbf{Y}_t^j)_{t \in \{1, \dots, T\}, j \in \{w, x, y, z\}}$:

$$\phi(\mathbf{Y}) = 1 - \frac{1}{D_z} \sum_{k=1}^{D_z} \mathbb{1}_{[(\mu_z)_k - (\sigma_z)_k, (\mu_z)_k + (\sigma_z)_k]}(\mathbf{Y}_k^z)$$

Here, the dimension of node z is denoted with D_z . Note that this function provides a gradual feedback of how many of the dimensions in node z are outside of the pre-specified corridor given by the threshold function.

FitzHugh-Nagumo System: Next, to allow for non-linear dynamic behaviour, we are generating data of the FitzHugh-Nagumo system (FHN), which is cyclic with regard to its summary graph, but acyclic in the unrolled graph \mathcal{G}_t . Although being a multivariate system, as the two dimensions interact, the corresponding dynamic SCM consists of one node x with two dimensions:

$$\dot{x}_1 = 3(x_1 - x_1^3/3 + x_2), \quad \dot{x}_2 = (0.2 - 3x_1 - 0.2x_2)/3$$

We chose the initial values as in [8] but with slightly reduced additive Gaussian noise variance $\sigma^2 = 0.0025$. The root cause is simulated similarly to the linear system by adding a constant to the difference equation at one dimension and time point. We classify an observation as faulty, if it deviates too much from a normal observation. As we have, in this setting, access to the ground truth, a normal observation is represented by a trajectory generated from the ground truth system. Consequently, the classifier consists of a time-varying threshold bound around each dimension of the normal observation without the injected root cause of node x . Denoting the expected trajectory from the system by \mathbf{E} a given observation \mathbf{Y} is then classified to be faulty if it does not deviate more than 10 standard deviations at any point in time from the expected trajectory: $\phi(\mathbf{Y}) = 1 - \prod_t \mathbb{1}_{[\mathbf{E}_t^x - 10\sigma_x, \mathbf{E}_t^x + 10\sigma_x]}(\mathbf{Y}_t^x)$.

5.2 Evaluation

When we have drawn counterfactual samples from our model, we calculate the approximate Shapley values (see Eq. 4) and use the ϕ function to evaluate each performed intervention based on whether it corrected the failure. The root cause is the intervention of the node j at time t that has the highest influence on the failure. If all counterfactual samples lead to the same ϕ evaluation for all interventions, then no unique root cause could be identified. However, due to random sampling of the counterfactual, this is an unlikely scenario (see for example Fig. 5.) Nevertheless, for the evaluation, we only require that the ground truth

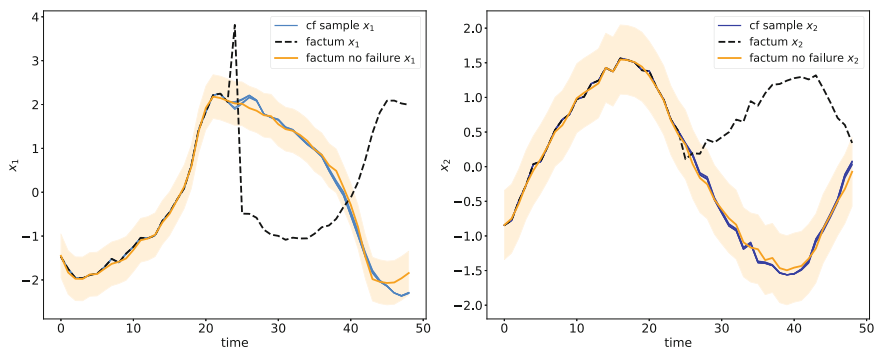


Fig. 2. The figure shows the counterfactual samples for the FHN system with injected root cause at $(j = x_1, t = 24)$. The injected root cause disrupts the system observation heavily (black dashed line). However, the counterfactual intervention performed by our model $NLin(S_t^j, N_t^j)$ corrects the failure in both dimensions, such that it lies inside the threshold region (orange area). (Color figure online)

root cause is within the set of identified root causes. In Fig. 2 we show five counterfactual samples for the nonlinear FHN system at the actual root cause injection point. Although the injected root cause is fairly large with regard to the interval of the normal observation without failure (drawn as orange line), the counterfactual intervention performed by our model $NLin(S_t^j, N_t^j)$ corrects the failure for both dimensions of x . In order to analyze root cause injections and how the identification capabilities of our model behave under varying injections, we performed an *Injection experiment* for the synthetic linear and nonlinear FHN system. External disturbances in dynamic systems may be propagated and thereby increase their impact. Alternatively, if the system is robust against incremental noise (as it is the case in the defined systems above due to the external noise influence even under the 'normal' conditions), it is not obvious how large an external influence at which point in time is noticeable. In Fig. 3, we show varying root cause injections for the linear synthetic system (varying constant added to the structural equation) over 20 randomly sampled facta with $T = 20$. It can be seen that the models intervening on the structure and the noise achieve a significantly higher identification score for large added constants. This could be due to a large root cause, in this setting leading to a factum with high distance to the normal data, which may lead to a divergence over time of the normal behaviour system \mathcal{M} . Note that we did a similar injection experiment for the FHN system, which can be found in the Appendix.

Assumption Violation. We probe our models on violation of the causal graph assumption for the linear synthetic system. For this, we modify the causal graph used by the underlying model through adding or removing random edges, while keeping the original summary graph for data generation. We use the same facta generated as $\sigma = 500$ in Fig. 3. In Table 1 it can be seen that removing edges for all models has a stronger impact on predictive performance than adding. As

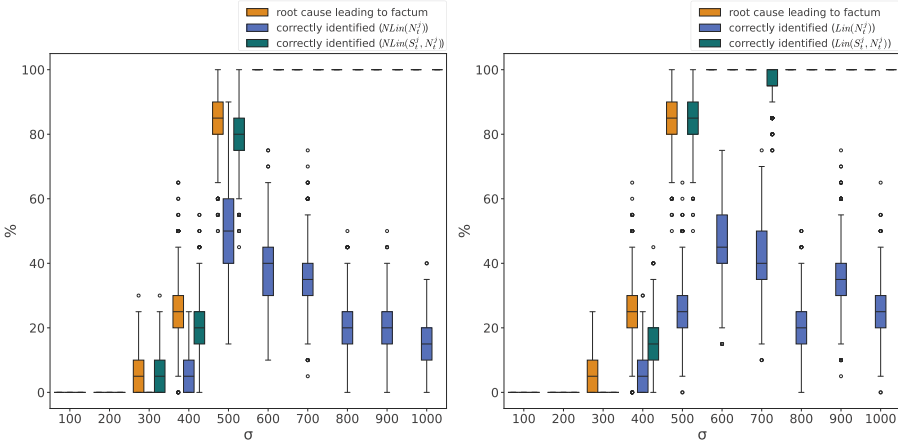


Fig. 3. The root cause was injected at a random node $j = x_1$ at $t = 6$ with varying constants in $[1, 10]$. The horizontal axis shows the injected constant in relation to the noise standard deviation denoted by σ . We report how many root causes could be identified in %.

Table 1. We show the **Accuracy** for the setup $\sigma = 500$ of the linear synthetic system (see Fig. 3) with varying number of *removed or added edges* of the summary graph \mathcal{G} used by the models.

	$NLin(S_t^j, N_t^j)$	$NLin(N_t^j)$	$Lin(S_t^j, N_t^j)$	$Lin(N_t^j)$
nr. of removed edges				
1	0.47 ± 0.25	0.23 ± 0.18	0.47 ± 0.24	0.18 ± 0.15
2	0.29 ± 0.20	0.06 ± 0.06	0.47 ± 0.24	0.12 ± 0.10
nr. of added edges				
1	0.82 ± 0.15	0.12 ± 0.10	1.0 ± 0.0	0.18 ± 0.15
2	0.88 ± 0.10	0.0 ± 0.0	0.88 ± 0.10	0.06 ± 0.06

expected, $Lin((S_t^j, N_t^j))$ performs best on this linear system (closely followed by $NLin((S_t^j, N_t^j))$). It must be mentioned that in a graph with only four edges, removing an edge is a major incision in the model assumption.

Linear EasyRCA Benchmark. We compare against the linear univariate benchmark of [2] consisting of six nodes and two types of root causes. The parametric root cause meaning they change the coefficient of the parent nodes to a random uniform sampled value. As a special case of the parametric setting, they inject structural root causes, which set the coefficient of the parent nodes to zero. Since EasyRCA excludes single time point root causes, in order to do a fair comparison, we only rank sets of interventions, where we intervene on all times for a given node and evaluating it accordingly by $Sh(\xi_0^j, \dots, \xi_T^j)$. In their work they inject on two nodes, where one is always the root node of the system and the other one a

randomly chosen node. As in their benchmark comparison the root node root cause is excluded, we exclude it from the evaluation as well. In the evaluation they distinguish for parametric and structural root causes, but because our model makes no prediction about the type of root cause, it is sufficient if EasyRCA predicted root causes contain the true root cause, regardless of the type. To rate the normality of a given trajectory \mathbf{Y} , we make use of the learned dynamical SCM \mathcal{M} which was fitted on normal observations of the system. More precisely, for the EasyRCA benchmark as well as the following River experiment, we used an outlier score similarly to [3], based on the learned dynamic SCM. That is, given a dynamic SCM \mathcal{M} consisting of N nodes and providing the conditional distribution $p(\mathbf{Y}_t^j | \mathbf{Y}_t^{PA(j,t)})$ via the dynamics equation learned from normal observational data $(\mathbf{Y})_k$, we can define the following outlier score:

$$\phi(\mathbf{Y}) = \frac{1}{NT} \sum_{j,t} \log p(\mathbf{Y}_t^j | \mathbf{Y}_t^{PA(j,t)}) \quad (8)$$

In Table 2, it can be seen that in general the intervention (S_t^j, N_t^j) is preferable to an intervention only on (N_t^j) . For the linear systems, the accuracy of $NLin(S_t^j, N_t^j)$ and $Lin(S_t^j, N_t^j)$ are similarly good, while *EasyRCA* shows lower performance in the factum $T = 100$ experiments. However, $Lin(S_t^j, N_t^j)$ is inadequate for addressing the complexities of the nonlinear problem (Fig. 3).

Table 2. We report the **Accuracy** over 20 facta of the summary graph on a linear system and the FHN oscillator. In the lower part of the table we present the experimental results of the EasyRCA benchmark [2] comparing the accuracy for one factum over 30 graphs for different factum lengths T (here, normal data has the same size T . We excluded the 2000 factum length experiment of the EasyRCA benchmark for computational reasons. Additionally, note that since EasyRCA is univariate, it can not be applied to our synthetic systems.)

	NLin		Lin		EasyRCA
	(S_t^j, N_t^j)	(N_t^j)	(S_t^j, N_t^j)	(N_t^j)	
Lin. system	0.94 ± 0.05	0.59 ± 0.24	1.0 ± 0.0	0.29 ± 0.20	–
FHN oscillator	0.90 ± 0.09	0.65 ± 0.23	0.20 ± 0.16	0.15 ± 0.13	–
Lin. Parametric					
Factum-100	1.0 ± 0.0	0.97 ± 0.03	1.0 ± 0.0	0.97 ± 0.03	0.87 ± 0.12
Factum-200	0.97 ± 0.03	0.93 ± 0.06	1.0 ± 0.0	0.93 ± 0.06	0.93 ± 0.06
Factum-500	1.0 ± 0.0	0.97 ± 0.03	1.0 ± 0.0	0.97 ± 0.03	1.0 ± 0.0
Factum-1000	1.0 ± 0.0	1.0 ± 0.0	0.97 ± 0.03	0.83 ± 0.14	0.97 ± 0.03
Lin. Structural					
Factum-100	1.0 ± 0.0	0.87 ± 0.12	1.0 ± 0.0	0.83 ± 0.14	0.8 ± 0.16
Factum-200	0.90 ± 0.09	0.27 ± 0.20	0.70 ± 0.21	0.53 ± 0.25	0.90 ± 0.09
Factum-500	1.0 ± 0.0	1.0 ± 0.0	0.87 ± 0.12	1.0 ± 0.0	1.0 ± 0.0
Factum-1000	1.0 ± 0.0	0.8 ± 0.16	1.0 ± 0.0	1.0 ± 0.0	0.97 ± 0.03

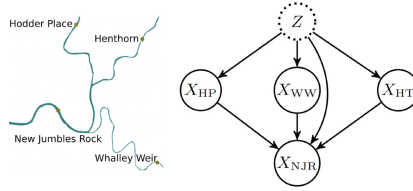


Fig. 4. With the geographical knowledge of the river flow, a summary graph can be inferred (Figure taken from [3]).

Real World River Experiment. We analyse our method on real-world data considering a univariate river experiment consisting of four nodes. The nodes represent measuring stations of the Ribble River in England (data is from [1]). These measuring stations are influenced by unknown external influences such as for example rain. For this reason, the summary graph includes an unobserved confounder Z that influences all nodes. This unobserved confounder affects the accuracy of our model when learning the normal system \mathcal{M} from observational data. The nodes represent stations of the Ribble River that measure the flow rate. Although this data-set has been investigated in [3], as a result of our dynamic viewpoint, we consider a slightly different factum. They consider four time points as static facta and infer the root causes for these. In contrast, we consider an entire time series as factum and infer the root cause. In addition, we use a finer time resolution of 15-minute intervals instead of averaged daily values, which has the advantage that the resulting SCM is less prone to instantaneous effects due to aggregation within a time-window. The finer resolution means that we consider a shorter period of time, namely the three days from 16.03.2019 to 19.03.2019 in which the flow rate is particularly high. As training data, we use the same time span as [3] from 01.01.2010 to 31.12.2018. They provide a z-score threshold for the New jumbles rock station, which we use as ϕ in 8, see also Fig. 5. We find the Shapley values with the highest scores at station Henthorn, which is an upstream station of the New jumbles rock station. Although no ground truth root cause exist for this experiment since it is a real world example, the result is plausible both geographically and with regard to the time point. Nevertheless, the counterfactual intervention cannot correct the failure, as the counterfactual sample is not below the z-score threshold. This could be due to the fact that the influence of the unobserved confounders is particularly high.

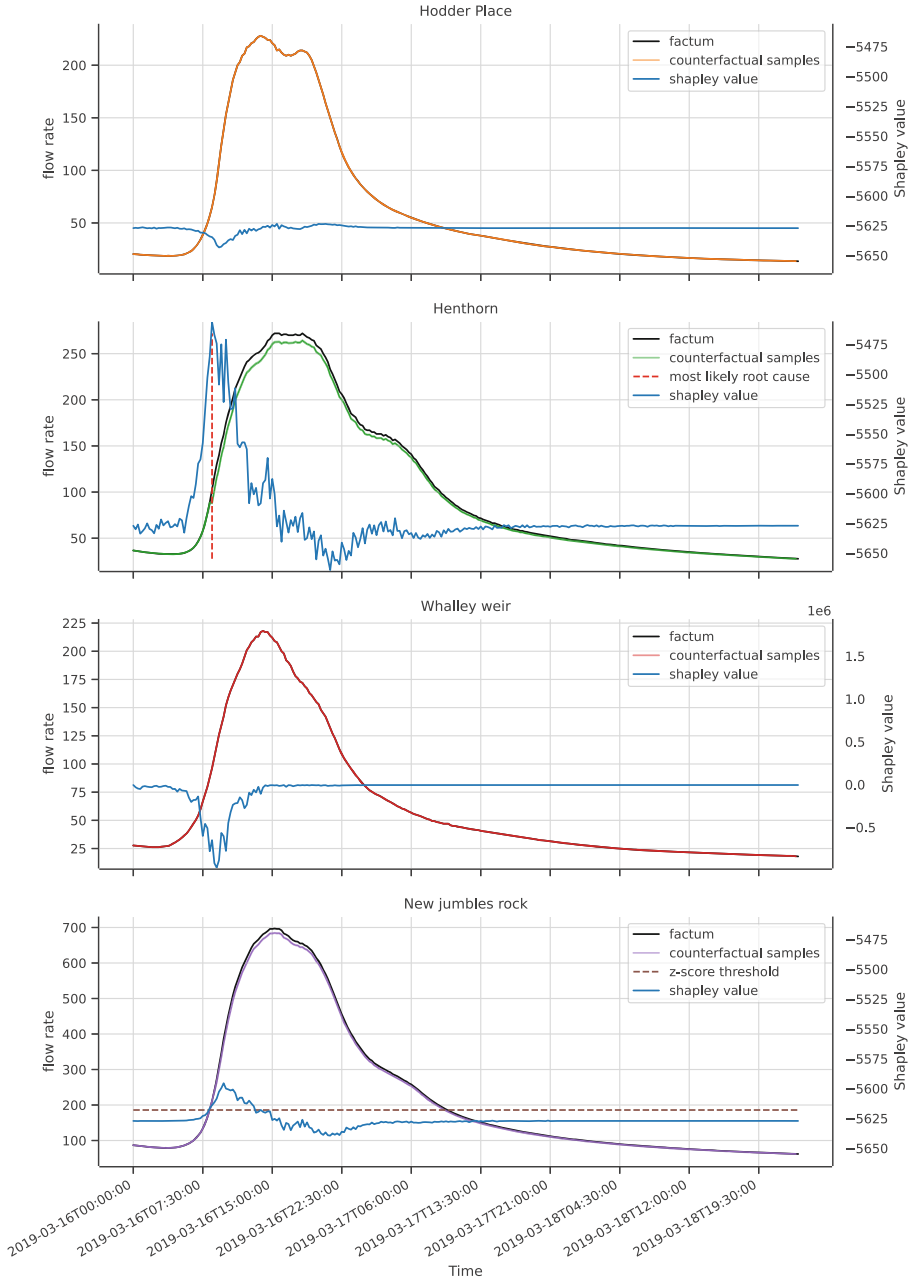


Fig. 5. We show five counterfactual samples (for each station) of our model $NLin(S_t^j, N_t^j)$ with the intervention at the predicted root cause at 08:30 on 16.03.2019. Additionally, we illustrate the resulting Shapley values for each time point, showing that right before the failure occurs the Shapley values increase.

6 Conclusion

In this paper, we have presented a method for identifying root causes in dynamic systems based on counterfactual reasoning. As the proposed method ranks individual interventions corresponding to individual nodes or sensors at particular times within a trajectory, our method is capable of exploiting not only the causal structure but also the natural direction of causality over time. By modelling temporal transitions with a non-linear neural network and a Shapley value approximation, we are able to remove important limitations of current counterfactual root cause analysis methods. While we demonstrated both on synthetic as well as real data the effectiveness of our method in identifying root causes in dynamic systems, there are several directions for further improvement. For example, our method is current limited to the assumption that the root cause consists of a single intervention and that the causal graphical structure is known as well as the absence of latent confounders. In future work, we plan to extend our method to identify multiple root causes and to include uncertainties in the graphical structure as well as potential latent confounders.

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

References

1. Hydrology data explorer. <https://environment.data.gov.uk/hydrology/explore>. Accessed 18 Mar 2024
2. Assaad, C.K., Ez-Zejjari, I., Zan, L.: Root cause identification for collective anomalies in time series given an acyclic summary causal graph with loops. PMLR (2023). <https://proceedings.mlr.press/v206/assaad23a.html>
3. Budhathoki, K., Janzing, D., Bloebaum, P., Ng, H.: Why did the distribution change? PMLR (2021). <https://proceedings.mlr.press/v130/budhathoki21a.html>
4. Budhathoki, K., Michailidis, G., Janzing, D.: Explaining the root causes of unit-level changes (2022). <https://arxiv.org/abs/2206.12986>
5. Budhathoki, K., Minorics, L., Bloebaum, P., Janzing, D.: Causal structure-based root cause analysis of outliers. PMLR (2022). <https://proceedings.mlr.press/v162/budhathoki22a.html>
6. Göbler, K., Windisch, T., Drton, M., Pychynski, T., Sonntag, S., Roth, M.: `causalAssembly`: generating realistic production data for benchmarking causal discovery (2024)
7. Halpern, J.Y., Hitchcock, C.: Graded causation and defaults. *Brit. J. Phil. Sci.* (2015)
8. Hegde, P., Çağatay Yıldız, Lähdesmäki, H., Kaski, S., Heinonen, M.: Variational multiple shooting for bayesian odes with gaussian processes (2022)
9. Ikram, A., Chakraborty, S., Mitra, S., Saini, S., Bagchi, S., Kocaoglu, M.: Root cause analysis of failures in microservices through causal discovery (2022)
10. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions (2017). https://papers.nips.cc/paper_files/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html

11. Pearl, J., Mackenzie, D.: The book of why: the new science of cause and effect (2018). <https://books.google.de/books?id=EmY8DwAAQBAJ>
12. Peters, J., Janzing, D., Schölkopf, B.: Elements of causal inference - foundations and learning algorithms. In: Adaptive Computation and Machine Learning Series (2017)
13. Shapley, L.S.: A Value for N-Person Games, Santa Monica, CA (1952)
14. Spirtes, P., Glymour, C., Scheines, R.: Causation, Prediction, and Search (2000)
15. Strobl, E.V., Lasko, T.A.: Identifying patient-specific root causes with the heteroscedastic noise model (2022)
16. Sulem, D., et al.: Diverse counterfactual explanations for anomaly detection in time series. arXiv preprint [arXiv:2203.11103](https://arxiv.org/abs/2203.11103) (2022)
17. Tonekaboni, S., Joshi, S., Campbell, K., Duvenaud, D.K., Goldenberg, A.: What went wrong and when? instance-wise feature importance for time-series black-box models (2020)