

Root Cause Explanation of Outliers under Noisy Mechanisms

Phuoc Nguyen, Truyen Tran, Sunil Gupta, Thin Nguyen, Svetha Venkatesh

A2I2, Deakin University

{phuoc.nguyen,truyen.tran,sunil.gupta,thin.nguyen,svetha.venkatesh}@deakin.edu.au

Abstract

Identifying root causes of anomalies in causal processes is vital across disciplines. Once identified, one can isolate the root causes and implement necessary measures to restore the normal operation. Causal processes are often modelled as graphs with entities being nodes and their paths/interconnections as edge. Existing work only consider the contribution of nodes in the generative process, thus can not attribute the outlier score to the edges of the mechanism if the anomaly occurs in the connections. In this paper, we consider both individual edge and node of each mechanism when identifying the root causes. We introduce a noisy functional causal model to account for this purpose. Then, we employ Bayesian learning and inference methods to infer the noises of the nodes and edges. We then represent the functional form of a target outlier leaf as a function of the node and edge noises. Finally, we propose an efficient gradient-based attribution method to compute the anomaly attribution scores which scales linearly with the number of nodes and edges. Experiments on simulated datasets and two real-world scenario datasets show better anomaly attribution performance of the proposed method compared to the baselines. Our method scales to larger graphs with more nodes and edges.

Introduction

Understanding the root causes behind anomalies in complex network systems holds significant importance across various disciplines, ranging from science to industry (Dhaou et al. 2021; Pool et al. 2020; Yi and Park 2021). Once the causes are identified, one can isolate and implement necessary measures to restore the normal operation of the process. Early fault recognition would prevent costly damage to operations, services, and products. For example, in complex network systems within modern manufacturing industries and information services, the cost of system failure is notably high (Ni et al. 2017), reaching as much as \$20,000 per minute of downtime in an automotive manufacturing plant (Djurdjanovic, Lee, and Ni 2003). The monitoring data and logs obtained from the processing nodes of these systems often contain noise attributed to random fluctuations in nodes and the links between them. In communication systems, for instance, random delays between nodes can arise

from bandwidth limitations and network overhead (Zhang, Branicky, and Phillips 2001). Given the intricate dependencies between monitoring nodes and the substantial volume of data involved, manual analysis of root causes becomes impractical.

Recent methods for root cause analysis (RCA) (Budhathoki et al. 2021, 2022) are based on a given causal structure of the system and a learned functional causal model (Peters, Janzing, and Schölkopf 2017) to explain the anomalous observations at a leaf node. These methods work by first detecting the anomalous leaf, then utilising the causal structure and counterfactual reasoning attribute the outlier scores to ancestor nodes (Budhathoki et al. 2022). The causal structure of a system is a powerful tool enabling the formal analysis of the root cause of unexpected event (Budhathoki et al. 2022). It is a directed acyclic graph (DAG) composed of nodes representing system components and directed edges representing causal links or dependency connections from parent nodes to child nodes. The observational data of nodes are assumed to have additive noise, the edge connections are assumed to be noise-free (Budhathoki et al. 2022). However, in computer networking systems, e.g., the connection between components can be noisy or faulty due to varying workloads, faulty hardware, wear and tear, or signal interference. Therefore, the root cause of an outlier can also arise from an edge in addition to a node. This raises the question of whether an anomalous edge could be detected by existing algorithms. In a recent work, Ni et al. (2017) attempted to detect faulty edges in those networks. Nevertheless, to the best of our knowledge, the study of noisy causal links as root causes has not been considered.

In this paper, we aim to fill this gap by generalising existing frameworks to allow the detection of anomalous edges as well as anomalous nodes. We consider changes at both the individual edges and model them via a Bayesian linear regression, where the noise in the causal edges is represented as the distribution of the regression weights.

In addition, we propose a causal contribution score called Bayesian Integrated Gradient of Edge and Node noise (BIGEN) to attribute a leaf anomaly score to the ancestor nodes and edges. First, given observations of outliers, we infer the noise values of the causal edges represented by Bayesian linear regression coefficients using the mode of the posterior distribution (MAP estimate) (Bishop and Nasrabadi 2006).

Second, we infer the noise values of the nodes and edges given the outlier observations. Third, we use an attribution method to compute the contribution of each noise term to the outlier score. Existing works for explaining outliers use Shapley-based attribution methods that require summing over all possible subsets of the ancestor nodes and edges (Sundararajan and Najmi 2020; Covert and Lee 2021; Budhathoki et al. 2022). This is computationally expensive, even for approximate methods, when applied to large graphs. In this work, we apply the integrated gradient (IG) method (Sundararajan and Najmi 2020) along the path from some references to the target noise for calculating the attribution scores. We show that this attribution method is efficient for scaling to graphs with thousands of nodes in attributing the root causes, compared to the baseline methods based on Shapley values (Sundararajan and Najmi 2020).

Our contributions are:

1. A framework for identifying the sources of changes in both nodes and edges to reduce operational costs.
2. Modelling the causal connections using Bayesian linear regression to enable the inference of edge noises and applying the Shapley-based attribution framework.
3. Introducing a new causal contribution score called B-IGEN to efficiently attribute a leaf anomaly to the ancestor nodes and edges.
4. Demonstrating the effectiveness of the proposed methods on random graph datasets and two real-world scenario datasets.

Preliminaries

Outlier Scores

Based on information theory, Budhathoki et al. (2022) introduced an outlier score that calibrates all probabilistic outlier scores. The authors defined the outlier score by characterising the tail probability of an event $X = x$ based on the distribution of some score space, such as negative log likelihood or z -score, as follow:

$$S_X(x) = -\log P \{-\log p(X) \geq -\log p(x)\} \quad (1)$$

$$\text{or, } S_X(x) = -\log P \{|X - \mu_x| \geq |x - \mu_X|\} \quad (2)$$

Functional Causal Mechanisms (FCMs)

Given a causal graph represented as a DAG, its functional causal mechanisms (Pearl 2009; Peters, Janzing, and Schölkopf 2017) can be described by the following set of equations. For each node j :

$$X_j = f_j(\text{Pa}_j, \epsilon_j)$$

$$\text{or, } X_j = \sum_{i \in \text{Pa}_j} W_{ij} X_i + \epsilon_j, \quad \text{when } f_j \text{ is linear}$$

where Pa_j is the node indices of the parents of node j . These equations represent the dependence of a node X_j on its parent nodes $X_{\text{Pa}_j} = \{X_i : i \in \text{Pa}_j\}$ and ϵ_j which is an additive noise random variable independent of X_i . If we assume f_j is linear, given an observation data matrix X , we can fit a linear regression model for this causal model to learn the weight parameters W_{ij} . A leaf outlier

$X_n = f_n(\text{Pa}_n, \epsilon_n)$ can be recursively expressed as a function depending on all the noise variables as follows:

$$X_n = f(\epsilon_1, \dots, \epsilon_n) = f(\epsilon)$$

where $P_\epsilon = P_{\epsilon_1} \times \dots \times P_{\epsilon_n}$.

Shapley-based RCA of Outliers

Budhathoki et al. (2022) employed the concept of Shapley values (Shapley et al. 1953) from cooperative game theory and randomised experiment to measure the contribution of a noise term to the target outlier score for an observed outlier at the leaf. The Shapley value, in the context of a coalition game, is an axiom-based method which uniquely divides (attributes) a pay-off among the players (the noise variables $\{\epsilon_j\}$ in our study) (Sundararajan and Najmi 2020).

Shapley Classic The originally proposed Shapley value (Shapley et al. 1953) for each player j is computed as the weighted average of its marginal contributions for all possible coalitions (subsets of players), both with and without the player j . Let $P = 1, \dots, d$ denote the set of players, $Q \subseteq P$ be a subset, $|Q|$ denote its cardinality, and v be the value function. The Shapley value for the player j is defined as:

$$\phi_j(v) = \sum_{Q \subseteq P \setminus j} \frac{|Q|!(|P| - |Q| - 1)!}{P!} (v(Q \cup i) - v(Q)) \quad (3)$$

Sampling This method approximates the Shapley values as a solution to the weighted least squares problem (Janzing, Minorics, and Blöbaum 2020; Covert and Lee 2021):

$$\min_{\phi_1, \dots, \phi_n} \sum_{Q \subseteq P} w(Q) (u(Q) - v(Q)) \quad (4)$$

$$w(Q) = \frac{d - 1}{\binom{d}{|Q|} |Q|(d - |Q|)}$$

where the weighting function $w(Q)$ depends only on the cardinality of the subset Q , and $u(S) = \sum_{i \in S} \phi_i$ represents the approximation function.

Permutation This method approximates the Shapley values using Monte Carlo samples of data and random permutations (Štrumbelj and Kononenko 2014). It first draws a random instance ϵ^m from the data, then chooses a random permutation of the independent variables, and finally computes the marginal contribution as:

$$\phi_j(v) = \frac{1}{M} \sum_{m=1}^M (v(\epsilon_{+j}^m) - v(\epsilon_{-j}^m)) \quad (5)$$

where $v(\epsilon_{+j}^m)$ is the prediction for ϵ , but with a random number of variables replaced by the values from ϵ^m , except for ϵ_j , and $v(\epsilon_{-j}^m)$ similarly, but ϵ_j is also replaced by ϵ_j^m .

Value Function The value function v in this study is the outlier score $S(x_n)$ at a leaf node (Budhathoki et al. 2022).

Methods

We first present our framework for RCA of an observed anomaly at a leaf node. We formulate a noisy mechanism to model the noise in the edge connection between component nodes. This allows the leaf node to be expressed as a function of both node and edge noise dependent model. Finally, we present the integrated gradient (IG) based contribution score of ancestor edge and node noises.

Framework

For concreteness, let us consider a computer network system¹ as a running example. We have continuous observations from the network components represented as the variables X_1, \dots, X_n . The causal relationships between these component variables are known and given in the form of a directed acyclic graph (DAG). There is edge noise in each connection from X_j to parent to describe potential noise between each parent-child pair. As example, this can be a random delay in the connection between two servers due to a router malfunction in between. Based on the causal mechanism (Peters, Janzing, and Schölkopf 2017), we introduce noise to each causal link from each parent node $X_i \in \text{Pa}_j$ to X_j in the causal graph to create a noisy mechanism, where Pa_j is the set of parent nodes of node j .

Definition 1. Noisy Causal Mechanisms.

For node j , a noisy mechanism is a generative process,

$$\xi_{.j} \sim \mathcal{N}(0, \alpha^{-1}I) \quad (6)$$

$$W_{.j} = \mu_{.j} + \xi_{.j} \quad (7)$$

$$\epsilon_j \sim \mathcal{N}(0, \beta^{-1}) \quad (8)$$

$$X_j = f_j(W_{.j}^T X_{\text{Pa}_j}) + \epsilon_j \quad (9)$$

where f_j can be a nonlinear function, $\xi_{.j}$ and ϵ_j are Gaussian edge and node noises with fixed variances α^{-1} and $\beta^{-1}I$ (α and β are the precision parameters), $\mu_{.j}$ is the mean of $p(W_{.j})$ with W_{ij} being the causal link from node i to node j for $i \in \text{Pa}_j$, and we use \cdot (dot) to denote the index varying in the set of parent node Pa_j .

We use this noisy mechanism to model the generative process of the system during normal operations. During testing, suppose some node or edge of the system behave strangely, either the node noise ϵ_j or some edge noise ξ_{ij} is interfered with, causing anomalous observations at downstream nodes and leaf nodes.

Definition 2. Abnormal Causal Mechanism

An abnormal causal mechanism of a node j is a similar generative process as the noisy causal mechanism of the node j in Definition 1 but with the noise distribution $p(\epsilon_j)$ or $p(\xi_{ij})$ interfered.

During abnormal operation, as defined in Definition 2, the underlying anomalous noises ξ_{ij} and ϵ_j render the generative processes defined in Eq. 9 to generate anomalous observations of the system. The RCA module then collects these

¹Our method is applicable to all structural causal processes, e.g., in a manufacturing process with sensor readings at each nodes as the variables.

anomalies along with their ancestors' observations, denoted as X' , for analysing and identifying the root causes using an attribution algorithm. Note that this generalises the approach presented in (Budhathoki et al. 2022) which only considers node anomalies.

Fitting the Noisy Mechanism In this paper, we model a noisy mechanism using Bayesian linear regression, allowing us to compute its posterior analytically. We use training data collected during the normal operation of the system to fit the posterior of the noisy mechanism, as described in the following theorem.

Theorem 3. Posterior distribution of noisy mechanism

For each node j , let (X_{Pa_j}, X_j) be the set of input-output observations of the noisy mechanism as defined in Definition 1. Assuming fixed α , and β precision parameters. If the prior for the weights of the causal links and the regression output are respectively:

$$\begin{aligned} W_{.j} &\sim \mathcal{N}(\mu_{.j}^0, \alpha^{-1}I) \\ X_j &\sim \mathcal{N}(W_{.j}^T X_{\text{Pa}_j}, \beta^{-1}I), \end{aligned}$$

then the posterior distribution of the weight vector $W_{.j}$ of the mechanism j is given by:

$$\begin{aligned} W_{.j}|X_j &\sim \mathcal{N}(\mu_{.j}, H) \text{ for} \\ \mu_{.j} &= H^{-1}(\alpha\mu_{.j}^0 + \beta X_{\text{Pa}_j}^T X_j) \\ H &= \alpha I + \beta X_{\text{Pa}_j}^T X_{\text{Pa}_j}. \end{aligned}$$

Proof. Refer to Section 3.3 in (Bishop and Nasrabadi 2006). \square

Proposition 4. MAP estimate of the noisy mechanism

Given the data and the prior in Theorem 3, the maximum a posteriori (MAP) estimate of the weights $W_{.j}$ is

$$W_{.j}^{MAP} = \mu_{.j} = H^{-1}(\alpha\mu_{.j}^0 + \beta X_{\text{Pa}_j}^T X_j)$$

Proof. By definition, the MAP estimate is the mode of the posterior distribution (Bassett and Deride 2019). Since the posterior is a Gaussian, its mode coincides with its mean. \square

Noise Dependent Reparametrisation

When considering RCA of a leaf outlier $X_n = f_n(\text{Pa}_n, \epsilon_n, \xi_n)$, it is convenient to re-parameterise X_n as a function of the node and edge noises (ϵ, ξ) to facilitate the attribution algorithms later on. We can recursively rewrite each ancestor of X_n in Eq. 9 up to the root node and arrive at a function of these noises

$$X_n = g(\epsilon, \xi) \quad (10)$$

where $P_\epsilon = \prod_i P_{\epsilon_i}$, and $P_\xi = \prod P_{\xi_{ij}}$ are independent noises. Using this re-parameterisation, we can attribute the root cause of these anomalies directly to these noise variables. Given a batch of abnormal values X' with x'_n being the observed anomaly at the leaf node X_n , we first infer the noise ξ' , and ϵ' as follows.

Edge noise estimation We employ the MAP estimate of the posterior from Proposition 4 to compute the new edge weights $W'_{.j}$ for each mechanism j . We then use the posterior weights estimated from the training data as the new prior. Finally, the estimated edge noise is defined as:

$$\xi'_{.j} = W'_{.j} - W_{.j} \quad (11)$$

Note that if X' comes from the same distribution as X , i.e., the generative process is normal, the new edge weights $W'_{.j}$ will coincide with the mode $W_{.j}$ of the fitted edge weight posterior, rendering the edge noise being close to a zero vector.

Node noise estimation We use Eq. 9 to estimate the node noise as

$$\epsilon_j = X_j - f_j(W_{.j}^T X_{Pa_j}). \quad (12)$$

Integrated Gradient of Edge and Node Noises

We introduce a Bayesian Integrated Gradient of Edge and Node noise (BIGEN in short) to attribute a leaf anomaly score to ancestor nodes and edges. Instead of using subsets as in Shapley values methods, we use noise reference to explain the root cause. This approach offers a performance gain compared to previous methods. Specifically, we select a reference node noise ϵ' from the normal dataset and the mean edge noise $\xi' = W$. Additional reference points can also be chosen to calculate multiple contribution scores, followed by averaging, to enhance the accuracy of the score.

Let f be the score function in Eq. 1 which is a continuous and differentiable function. Then $\phi_i(x, x', f)$ is defined as the integral of the gradient of f along the straight-line path between x and x' . Formally, we define the path between x and x' as $\gamma(t) = tx + (1-t)x'$ for $t \in [0, 1]$. Then, the Integrated Gradient (IG) for the i th feature x_i is defined as:

$$\begin{aligned} \text{IG}_i(x, x', f) &= (x - x') \int_{t=0}^1 \frac{\partial f(\gamma(t))}{\partial \gamma_i(t)} \frac{\partial \gamma_i(t)}{\partial t} dt \\ &= (x_i - x'_i) \int_{t=0}^1 \frac{\partial f(x' + t(x - x'))}{\partial x_i} dt \end{aligned} \quad (13)$$

This score integrates the gradient along the path from the reference to the inferred noise to calculate the attribution of each node and edge towards the observed anomaly score. In our case, with two noise variables ϵ and ξ , computing the IG for one requires marginalising over the other noise variable. Therefore, we adjust this IG to fit into our attribution to node and edge noises as follows:

$$\text{IG}_i(\epsilon, \epsilon', f) = \mathbb{E}_{\xi} [\epsilon_i - \epsilon'_i] \int_{t=0}^1 \frac{\partial f_t}{\partial \epsilon_i} dt \quad (14)$$

$$\text{IG}_{ij}(\xi, \xi', f) = \mathbb{E}_{\epsilon} [\xi_{ij} - \xi'_{ij}] \int_{t=0}^1 \frac{\partial f_t}{\partial \xi_{ij}} dt \quad (15)$$

This method is more advantageous than subset sampling, as it is linear and only dependent on the number of discretised steps in the path. In contrast, Shapley-based method (Sundararajan and Najmi 2020) requires summing over all possible subsets of the ancestor nodes and edges, which grows exponentially with the number of nodes and edges.

Score function Since the joint distribution of nodes and edges is Gaussian, the conditional distribution of the normal observations at the leaf node is also Gaussian. In this case, the outlier score using the negative log-likelihood feature in Eq. 1 is equivalent to the outlier score with the z -score feature in Eq. 2. An efficient evaluation function can be derived using the error function as follows:

$$\begin{aligned} S_X(x) &= -\log P \{-\log p(X) \geq -\log p(x)\} \\ &= -\log P \left\{ \frac{|X - \mu_X|^2}{2\sigma_X^2} \geq \frac{|x - \mu_X|^2}{2\sigma_X^2} \right\} \\ &= -\log P \left\{ \left| \frac{X - \mu_X}{\sigma_X} \right| \geq \left| \frac{x - \mu_X}{\sigma_X} \right| \right\} \\ &= -\log \{1 - \Phi(z)\} = -\log \Phi(-z) \end{aligned} \quad (16)$$

where $z = \left| \frac{x - \mu_X}{\sigma_X} \right|$, $\Phi(z) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^z e^{-\frac{t^2}{2}} dt$ is the standard normal cumulative distribution function, and (μ_X, σ_X) represents the maximum likelihood estimate of the marginal mean and standard deviation of the marginal distribution of the target node.

Causal graph and noisy FCMs While our focus is not on solving causal discovery, we assume a causal graph is given. For each fixed weight noise vector, we have a set of FCMs with additive node noise and their FCM parameters can be learned from normal observation data (Peters, Janzing, and Schölkopf 2017). Budhathoki et al. (2022) observed that when representing the FCM of a target leaf node as a functional of all node noise, each noise vector takes on the role of selecting a deterministic mechanism. We generalise this idea and represent the FCMs as functionals of the edge and node noises, which similarly play the role of choosing the deterministic mechanisms. Inferred noises of outlier edges or nodes, therefore, will select outlier mechanisms in either cases. We build upon the success of Budhathoki et al. (2022)'s approach, to infer functions and noise from data, and show that counterfactual contribution score by changing the noise term w.r.t. a reference is effective for causal attribution at both the node and edge levels. This approach utilises Pearl's third ladder of causation (Pearl 2009) to enhance our understanding of the system's causes and effects.

Shapley values and IGs Budhathoki et al. (2022) computed Shapley value contributions numerically by averaging over all orderings sets. However for larger number of variables, approximation is needed to be practical. In the experiments, we sample orderings instead of using all orderings and compare Shapley with early stopping (Shapley), subset sampling (Sampling), and methods based on a fixed number of randomly generated permutations (Permutation). For the IGs, however, no subsets of intervention are needed but rather a gradient path is taken along the path from some reference noises to the target noises. For each noise vector, BIGEN requires one forward pass and one backward pass. Since we use a small fixed number of references, this computation scales linearly with the number of nodes and edges in the subgraph. One advantage of BIGEN over Shapley methods is that it can be applied to nonlinear (noisy) FCMs

since the contribution score is still linear w.r.t. the gradients, thus satisfying the efficiency axiom (Shapley et al. 1953).

Related Work

Causal structure discovery-based techniques have been recently used to find the root cause(s) of faults in cloud applications (Arnold, Liu, and Abe 2007; Wang et al. 2018). More recently, Budhathoki et al. (2022) introduced a causal structure based root cause explanation for outlier using Shapley values. The authors represented the dependence of an outlier leaf directly in terms of the noise of ancestor nodes. This approach facilitates intervention on the noise distributions. Their work, however, assumes that the outlier does not originate from the connections in the functional causal models. In real computer network systems, e.g., the connections may be broken or malfunction, causing anomalies in downstream nodes. In our present work, we relax this constraint and consider also changes in the weights as one possible source of the root cause, in addition to the node anomalies.

Budhathoki et al. (2021) model a distribution change of a mechanism for certain nodes and use a new dataset under this new mechanism to fit a new model. In contrast, our work models the uncertainty in the edges of the mechanisms and allows for temporary anomaly edge noise causing a faulty batch of observations. Under Bayesian view, we employ the MAP estimate of the faulty edge noise deviation from the normal edge prior. By estimating noise in both edges and nodes, we can quantify the contribution of each noise term to the outlier score observed at the leaf. This can be viewed as intervention and counterfactual analysis on the nodes and edges. To the best of our knowledge, we are the first to explain anomalies in both nodes and edges based on the given causal structure. A different approach from our work, which combines RCA and causal discovery in one framework, is proposed by Ikram et al. (2022).

In the recent domain of explainable AI, methods based on Shapley value (Shapley et al. 1953; Sundararajan and Najmi 2020) have gained increasing popularity. These methods use an axiomatic approach to design attribute functions with desirable properties, e.g., being fair, unique, and efficient. These methods explain prediction outcomes by attributing the prediction score back through the deep neural networks to the input features (Erion et al. 2021; Yang et al. 2022). Among them, integral-based attribution methods are the most efficient, as they use only a reference to represent the absence of the input signal, rather than a random feature from the training data (Lundberg and Lee 2017; Sundararajan and Najmi 2020; Samek et al. 2021). However, these methods are designed for explaining neural network predictions, which differs from our goal of RCA.

Experiments

We run experiments on random graph datasets and two real-world settings, namely a micro cloud service and a supply chain scenarios. We compare our methods against three baselines described in the Preliminaries section, as well as a naive approach:

1. **Shapley** (classic): This method employs Shapley values as defined in Eq. 3 for the contribution of each node.
2. **Sampling**. This method calculates the Shapley values by selecting random subsets and weighted least squares regression (Janzing, Minorics, and Blöbaum 2020; Covert and Lee 2021), as indicated in Eq. 4.
3. **Permutation**: This method computes the Shapley values through permutation sampling (Štrumbelj and Kononenko 2014), as shown in Eq. 5.
4. **Naive**: This method uses the marginal distribution, i.e., the observational distribution, of each node X_j to compute the Shapley values.
5. **BIGEN** (ours): We use the contribution scores outlined in Eq. 14, 15 to assess the node and edge contributions.

For all methods, we use the outlier score in Eq. 16 and assess the contribution to this score by each ancestor node and edge. For the edge score of the baselines, we use the outer product of the node scores to estimate the edge score $s_{edge} = s_{node} s_{node}^T$, where $s_{node} = (s_1, \dots, s_n)$ represents the contributions of all the ancestor nodes (including the target node). This score quantifies the level of anomaly associated with an edge e_{ij} by combining the scores observed at nodes i and j .

We adhere to the approach of Budhathoki et al. (2022) and employ NDCG@ k (Järvelin and Kekäläinen 2017) to gauge rankings based on graded relevance of outcomes. NDCG@ k yields values within the $[0, 1]$ range, with higher scores indicating that highly “relevant” root causes are assigned higher ranks. To establish the ground truth relevance of all nodes, we assign zero relevance scores to non-root causes, and invert the ranking of injected root causes, akin to the methodology in (Budhathoki et al. 2022).

Random Graph Datasets

We randomly generate 1,000 causal graphs with varying number of nodes in the range from 10 to 10000 nodes. The noisy causal mechanisms follow Definition 1 with $\alpha_j^{-1} = 1$, $\beta_{ij}^{-1} = 0.01$, and $\mu_{ij} \sim |\mathcal{N}(0, 1)|$. We draw normal data from these noisy FCMs, following their generative process. For the abnormal data, we randomly select a target node X_n from this causal graph. We then choose among its ancestor node either $k \in [1, \dots, m]$ root-cause nodes, or $l \in [1, \dots, m]$ root-cause edges, or $k + l$ root-cause nodes and edges. Here, m is chosen to be 10% of the number of nodes in the subgraph. We inject outlier noises into the nodes and edges to create the ground truths as follows: The outlier node noise ϵ_j of node j is randomly drawn from $\mathcal{N}(a, b)$, where a is drawn from $\pm\text{Uniform}(3, 5)$ and b is drawn from $\text{Uniform}(3, 5)$. The outlier edge noise ξ_{ij} of each node j is randomly drawn from $\mathcal{N}(am_j, bs_j)$, where a is drawn from $\pm\text{Uniform}(3, 5)$, b is drawn from $\text{Uniform}(3, 5)$, and m_j represents the maximum magnitude of the current weights w_{ij} (i.e., $m_j = \max_i(|w_{ij}|)$), and s_j represents the maximum standard deviation σ_{ij} (i.e., $s_j = \max_i(\sigma_{ij})$).

Fig. 1 presents a comparison of the results obtained from all methods. On average, BIGEN outperforms all the baselines in detecting the actual root causes of outliers in the top




Figure 1: NDCG@ k ranking of root cause detection on random graphs.

	Nodes	Edges	Nodes + Edges
Shapley*	81.5 \pm 4.5	82 \pm 11	81.8 \pm 7.7
Sampling	84.7 \pm 9.5	77.2 \pm 13.7	80.9 \pm 11.6
Permut.	84.5 \pm 11.5	82 \pm 11.8	83.3 \pm 11.6
Naive	69.1 \pm 13.4	53.6 \pm 30.3	63.3 \pm 22.2
BIGEN	89.8 \pm 8.5	93.8 \pm 10	91.8 \pm 9.5

Table 1: NDCG@ k in percent for root cause detection in random graphs. (*) Due to exponential runtime, for >20 nodes, we compute Shapley values using early stopping when the contribution score does not change much (Blöbaum et al. 2022).

k candidates, at over NDCG 0.9. The Shapley, Sampling, and Permutation methods can overall detect root causes of outliers, at around NDCG 0.83, with similar performance across them. The Shapley method seems to perform better than Sampling and Permutation in node and edge anomaly attributions on average. The Naive method, on the other hand, could not attribute the root causes correctly. Table 1 shows the details of node and edge attribution results for the case when both type of anomaly are present. Overall, the edge score rankings are on average lower than that of the node scores. This suggests that it is greater difficulty in attributing root causes when edge anomalies are present.

Runtime Next, we compare the runtime of BIGEN to Shapley-based methods. Fig. 2 shows the runtime of all methods on a Ubuntu 20.04 workstation with an Intel Xeon E5-1650 v4 CPU and 46Gb RAM. Notably, both the BIGEN and Naive RCA methods exhibit linear time complexities relative to the number of upstream nodes, completing computations in less than a minute for target nodes with roughly 200 ancestors. In contrast, the Shapley method demonstrates exponential complexity, requiring over 2 hours for cases involving more than 20 nodes. This disparity arises due to the Shapley method’s necessity to iterate through all conceivable subsets for the computation of the weighted average of each subset’s contribution score. The Sampling and Permutation methods tend to exhibit polynomial time complexities, taking just slightly over 1 hour to compute for a target outlier node featuring around 200 ancestor nodes. The Naive method independently computes the contribution of each node, while BIGEN capitalises on gradient information and baseline noises to compute the contribution of each noise term within the given context.




Figure 2: Comparing wall-clock runtime (in minutes) between BIGEN and Causal-RCA methods with increasing number of nodes. The attribution complexity of BIGEN and Naive is $O(d + e)$, while Shapley is $O(2^{(d+e)})$. Sampling and Permutation methods tend to have polynomial time.

The key takeaway from this experiment is that BIGEN provides relevant top- k ranking across values of k while reducing the computation time significantly. This is because BIGEN can model the uncertainty in the causal edges to handle situations with noisy mechanisms.

Root Causes of Observed Latencies in Cloud Services


In this experiment, we study the root causes of unexpected observed latencies in a microservice of an online shop. Microservices are building blocks for complex mobile and Internet-of-Things (IoT) applications. Therefore, it is essential to keep service access delay minimal (Guo, Tang, and Tang 2022). Any unusual delay should be quickly studied to identify the root causes and resolve the problem (Ikram et al. 2022). Often, the delay at a node is composed of the communication delay between itself and each parent node (due to network bandwidth, load, and other turbulences) and execution delay of each node (due to the request complexity, node computational power and load).

We use the microservice architecture described in (Blöbaum et al. 2022). We assume the delay noise in each node X_j is a Gaussian noise with unknown variance α_j^{-1} and in each edge W_{ij} is a Gaussian noise with unknown variance β_{ij}^{-1} . We thus make a realistic assumption and allow for noisy connections to account for the above scenarios. The task is to explain the root cause of an unwanted observed latency at the customer end in processing an online order. This service involves multiple other web-services described by a causal dependency graph involving ten other services (Blöbaum et al. 2022), part of the causal graph is shown in Fig. 4 in the Appendix. Assuming that we observe latencies in the order confirmation of the Website leaf node, we also assume all services are synchronised.

For the abnormal data of the target node Website we select among its ancestor node either $k \in [1, \dots, 3]$ root-cause nodes or $l \in [1, \dots, 3]$ root-cause edges or $k + l$ root-cause nodes and edges, then we inject anomalous noises by random sampling the noise outside 3σ ranges of their normal operation distributions to create the ground truths.

We use similar models as in the previous section for this

	Nodes	Edges	Nodes + Edges
Shapley	90.4 ± 4.7	77.8 ± 9.1	84.1 ± 6.9
Sampling	91.8 ± 5	71.5 ± 14	81.6 ± 9.5
Permut.	91.4 ± 4.8	72.6 ± 11.2	82 ± 8
Naive	79.7 ± 10	46.5 ± 19.9	63.1 ± 15
BIGEN	91 ± 6	92.4 ± 7.6	91.7 ± 6.8

Table 2: NDCG@ k in percent for RCA in a micro cloud service.Figure 3: NDCG@ k when varying the number of k , error bars not included to keep the figure uncluttered.

experiment, i.e., assuming the causal graph is given then fit the noisy FCMs models to the training data collected during normal operations.

Results Table 2 displays the outcomes achieved by all methods, revealing the NDCG@ k values for root cause detection within a micro cloud service environment. It shows that the Shapley-based methods can detect the root cause in both the nodes and edges. However, the edge score ranking is lower than that of the node score. This suggests that the deterministic weights methods fall short in high weight noise applications. However, the Shapley method shows better attribution results than Sampling and Permutation. The Naive method can detect node root causes but not edges root causes. The BIGEN, in contrast, can model the uncertainty in the causal links therefore can account better the contribution of the root causes to the observed outliers at leaves. Overall, BIGEN shows better root cause attributions across nodes, edges, and both compared to the baselines.

Fig. 3 shows the NDCG raking at different k between the methods. It shows that BIGEN can rank the root cause better on average than all methods for across top- $k \geq 2$ values. The increasing NDCG@ k scores of all methods with larger k values show that the relevant root causes have more chance to appear in the top- k results.

Root Causes of Outliers in a Supply Chain

In this experiment, we apply our methods to model the noisy interactions between businesses. The behaviour in supply chains usually includes complex interaction of organisation structure, and time delays between decision and implementation (Power 2005). To gain competitive advantage, it is critical that businesses need to optimise their supply chain operations to ensure the flow of physical goods between trading partners. Thus, it is important that delays to this flow

	Nodes	Edges	Nodes + Edges
Shapley	89.1 ± 4.9	92.7 ± 7.2	90.9 ± 6.1
Sampling	89 ± 5	92.5 ± 7.4	90.8 ± 6.2
Permut.	89.2 ± 4.9	92.6 ± 7.2	90.9 ± 6.1
Naive	85.6 ± 5.5	91.5 ± 9.1	88.6 ± 7.3
BIGEN	89 ± 5.2	98 ± 4.5	93.5 ± 4.8

Table 3: NDCG@ k in percent for RCA in a supply chain.

can be identified and rectified in a reliable and timely manner. Often, business decisions are made based on predictions of real demands and constraints and are not a deterministic process (Power 2005).

We study RCA of outliers in a supply chain process shown in Fig. 5 in the Appendix. In this process, a retailer need to submit orders to a vendor based on its inventory constraint of current stock and the demand forecasting of near future sale to fulfil future customers' purchases faster. The vendors will then confirm the retailer's purchase orders with a random delay. When these orders are confirmed, the goods will be sent to the retailer again with some random delay and may arrive at different times. The random delays can be due to various overhead costs such as decisions of the vendors' manager, variable packaging and shipping operations (Croson et al. 2014). We assume noisy linear FCMs for the supply chain process where the noise term of each node is a Gamma distribution to mimic real-world settings, i.e., heavily-tailed behaviour. We assume Gaussian noises for the edges to account for the fluctuations in the causal links between nodes.

For the abnormal data of the target node *Received* we randomly select among its ancestor nodes either $k \in \{1, 2\}$ nodes or $l \in \{1, 2\}$ edges or $k = 1, l = 1$ node and edge, then inject anomalous noises by random sampling the node noise $\epsilon \sim \text{Uniform}(3, 5)$ and edge connection noise $\xi \sim \text{Uniform}(3, 5)$ then collect the outliers.

Results Table 3 presents the outcomes of all methods employed for root cause detection within a supply chain context. The result shows good performance among all methods with our proposed method BIGEN has highest detection ranking followed by Shapley methods, then Naive-RCA. Notably, our method can accurately identify the anomalous link in the process, at NDCG@ k of 0.98 on average. This confirms the effectiveness of the noisy FCM models and the BIGEN attribution method.

Conclusions

We have introduced a framework for identifying the root causes of unexpected events observed at leaf nodes in causal generative processes. We proposed noisy functional causal models that enable the inference of node and edge noises suitable for attribution algorithms. We further developed an efficient attribution score for large graphs based on integrated gradients. Our experimental results demonstrated the effectiveness of our proposed methods.

References

- Arnold, A.; Liu, Y.; and Abe, N. 2007. Temporal causal modeling with graphical granger methods. In *13th ACM SIGKDD*, 66–75.
- Bassett, R.; and Deride, J. 2019. Maximum a posteriori estimators as a limit of Bayes estimators. *Mathematical Programming*, 174: 129–144.
- Bishop, C. M.; and Nasrabadi, N. M. 2006. *Pattern recognition and machine learning*, volume 4. Springer.
- Blöbaum, P.; Götz, P.; Budhathoki, K.; Mastakouri, A. A.; and Janzing, D. 2022. DoWhy-GCM: An extension of DoWhy for causal inference in graphical causal models. *arXiv preprint arXiv:2206.06821*.
- Budhathoki, K.; Janzing, D.; Bloebaum, P.; and Ng, H. 2021. Why did the distribution change? In *AISTATS*, 1666–1674.
- Budhathoki, K.; Minorics, L.; Blöbaum, P.; and Janzing, D. 2022. Causal structure-based root cause analysis of outliers. In *ICML*, 2357–2369.
- Covert, I.; and Lee, S.-I. 2021. Improving kernelshap: Practical shapley value estimation using linear regression. In *International Conference on Artificial Intelligence and Statistics*, 3457–3465. PMLR.
- Crosan, R.; Donohue, K.; Katok, E.; and Serman, J. 2014. Order stability in supply chains: Coordination risk and the role of coordination stock. *Production and Operations Management*, 23(2): 176–196.
- Dhaou, A.; Bertoncello, A.; Gourvéne, S.; Garnier, J.; and Le Penec, E. 2021. Causal and interpretable rules for time series analysis. In *27th ACM SIGKDD*, 2764–2772.
- Djordjanovic, D.; Lee, J.; and Ni, J. 2003. Watchdog Agent an infotronics based prognostics approach for product performance degradation assessment and prediction. *Advanced Engineering Informatics*, 17(3-4): 109–125.
- Erion, G.; Janizek, J. D.; Sturmfels, P.; Lundberg, S. M.; and Lee, S.-I. 2021. Improving performance of deep learning models with axiomatic attribution priors and expected gradients. *Nature Machine Intelligence*, 3(7): 620–631.
- Guo, F.; Tang, B.; and Tang, M. 2022. Joint optimization of delay and cost for microservice composition in mobile edge computing. *World Wide Web*, 25(5): 2019–2047.
- Ikram, A.; Chakraborty, S.; Mitra, S.; Saini, S.; Bagchi, S.; and Kocaoglu, M. 2022. Root Cause Analysis of Failures in Microservices through Causal Discovery. *NeurIPS*, 35: 31158–31170.
- Janzing, D.; Minorics, L.; and Blöbaum, P. 2020. Feature relevance quantification in explainable AI: A causal problem. In *International Conference on artificial intelligence and statistics*, 2907–2916. PMLR.
- Järvelin, K.; and Kekäläinen, J. 2017. IR evaluation methods for retrieving highly relevant documents. In *ACM SIGIR Forum*, volume 51, 243–250. ACM New York, NY, USA.
- Lundberg, S. M.; and Lee, S.-I. 2017. A unified approach to interpreting model predictions. *NeurIPS*, 30.
- Ni, J.; Cheng, W.; Zhang, K.; Song, D.; Yan, T.; Chen, H.; and Zhang, X. 2017. Ranking causal anomalies by modeling local propagations on networked systems. In *2017 IEEE ICDM*, 1003–1008. IEEE.
- Pearl, J. 2009. *Causality: Models, Reasoning, and Inference (2nd ed.)*. Cambridge university press.
- Peters, J.; Janzing, D.; and Schölkopf, B. 2017. *Elements of causal inference: foundations and learning algorithms*. The MIT Press.
- Pool, J.; Beyrami, E.; Gopal, V.; Aazami, A.; Gupchup, J.; Rowland, J.; Li, B.; Kanani, P.; Cutler, R.; and Gehrke, J. 2020. Lumos: A library for diagnosing metric regressions in web-scale applications. In *26th ACM SIGKDD*, 2562–2570.
- Power, D. 2005. Supply chain management integration and implementation: a literature review. *Supply chain management: an International journal*, 10(4): 252–263.
- Samek, W.; Montavon, G.; Lapuschkin, S.; Anders, C. J.; and Müller, K.-R. 2021. Explaining deep neural networks and beyond: A review of methods and applications. *Proceedings of the IEEE*, 109(3): 247–278.
- Shapley, L. S.; et al. 1953. A value for n-person games.
- Štrumbelj, E.; and Kononenko, I. 2014. Explaining prediction models and individual predictions with feature contributions. *Knowledge and information systems*, 41: 647–665.
- Sundararajan, M.; and Najmi, A. 2020. The many Shapley values for model explanation. In *ICML*, 9269–9278.
- Wang, P.; Xu, J.; Ma, M.; Lin, W.; Pan, D.; Wang, Y.; and Chen, P. 2018. Cloudranger: Root cause identification for cloud native systems. In *2018 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID)*, 492–502. IEEE.
- Yang, P.; Akhtar, N.; Wen, Z.; Shah, M.; and Mian, A. S. 2022. Re-calibrating Feature Attributions for Model Interpretation. In *11th ICLR*.
- Yi, J.; and Park, J. 2021. Semi-supervised bearing fault diagnosis with adversarially-trained phase-consistent network. In *27th ACM SIGKDD*, 3875–3885.
- Zhang, W.; Branicky, M. S.; and Phillips, S. M. 2001. Stability of networked control systems. *IEEE control systems magazine*, 21(1): 84–99.