# Unsupervised Anomaly Detection on Microservice Traces through Graph VAE

**Zhe Xie**[1], Haowen Xu, Wenxiao Chen, Wanxue Li, Huai Jiang, Liangfei Su, Hanzhang Wang, Dan Pei

1. Presenter. Email: xiez22@mails.tsinghua.edu.cn

# What is a Trace

https://example.com/checkout

External Request

# What is a Trace

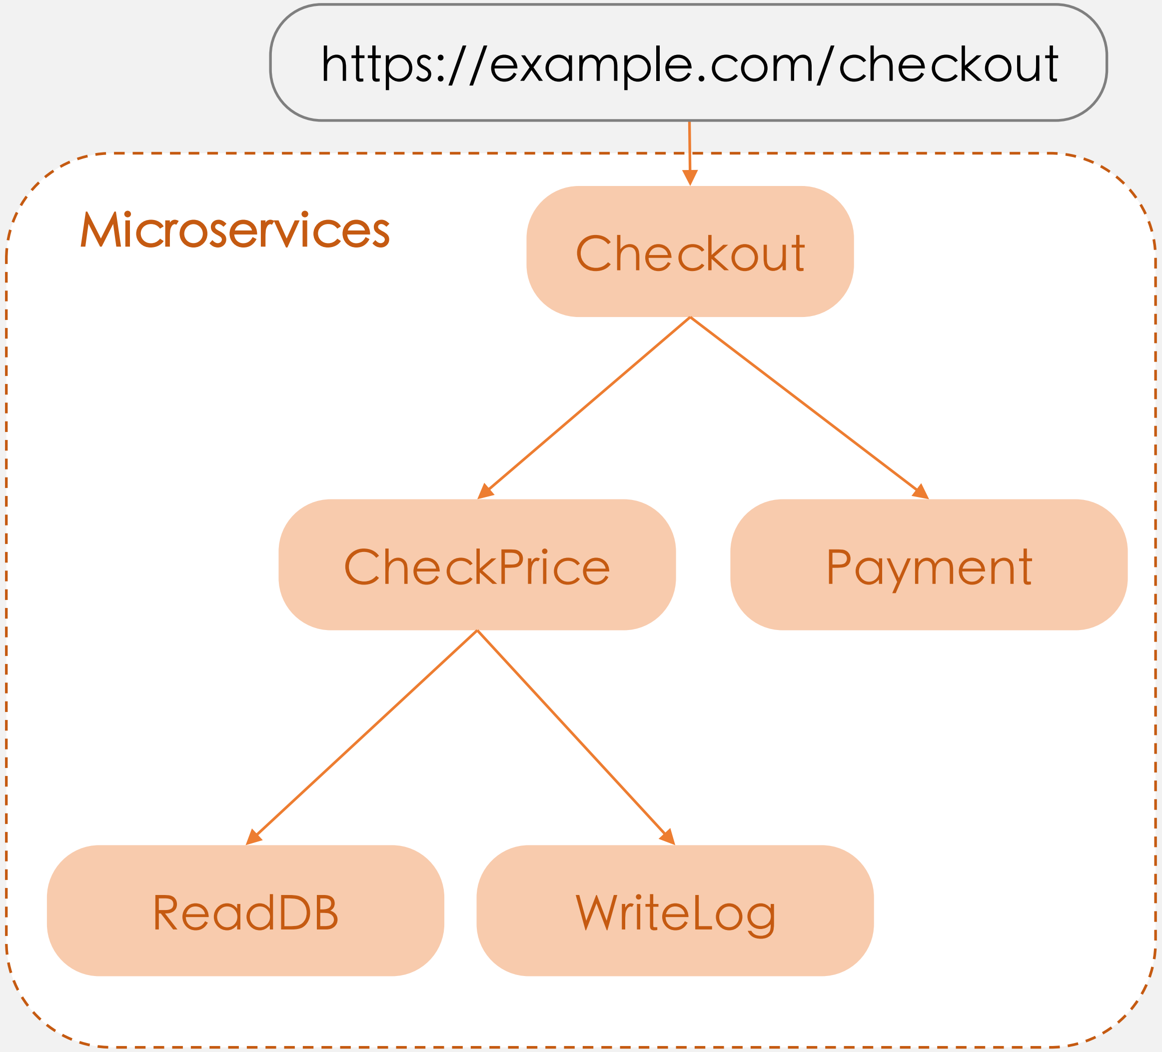


External Request

Inter-service Requests

# What is a Trace

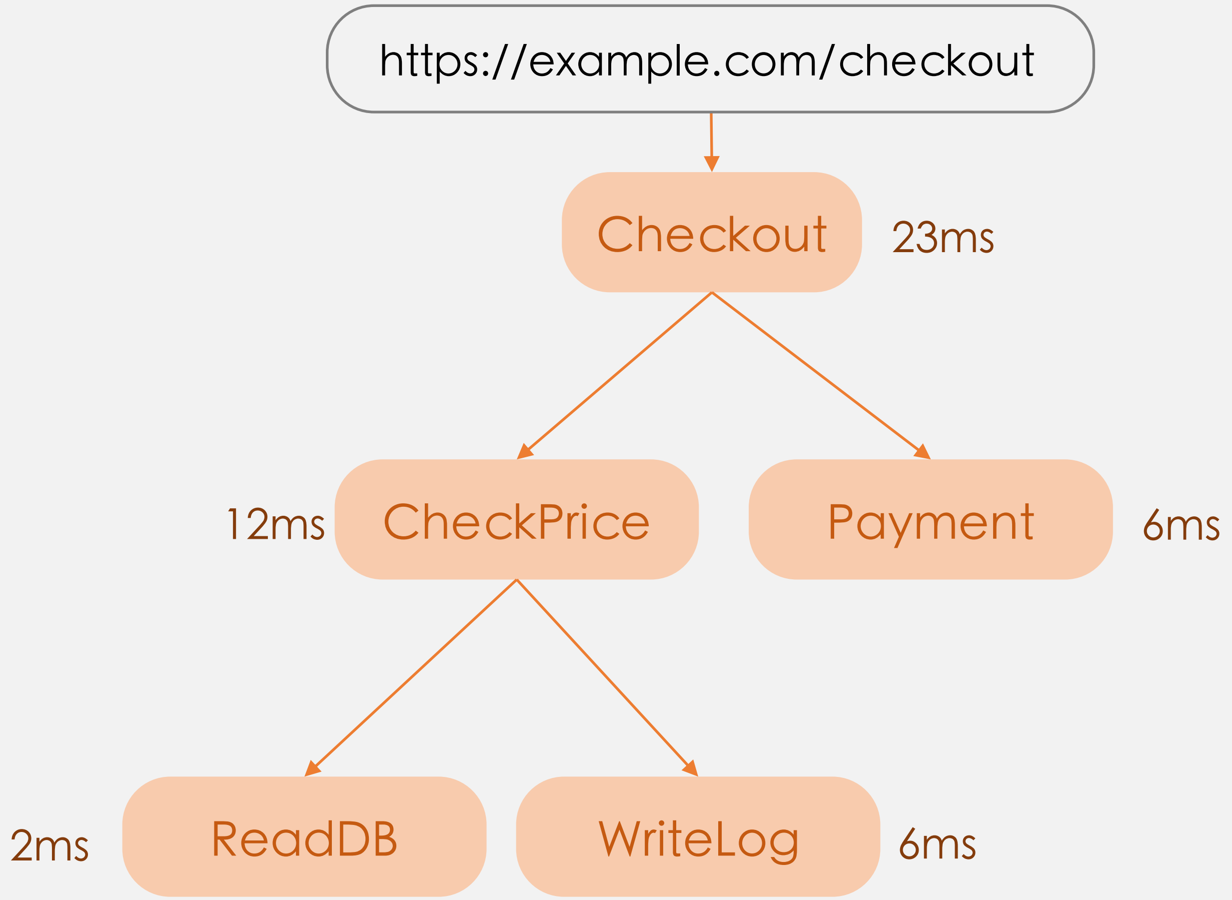


**Example of a Trace Record**

| ID | Parent | API | Time | Code |
|---|---|---|---|---|
| 0 | - | Checkout | 23 | 200 |
| 1 | 0 | CheckPrice | 12 | 200 |
| 2 | 1 | ReadDB | 2 | 200 |
| 3 | 1 | WriteLog | 6 | 200 |
| 4 | 0 | Payment | 6 | 200 |

A **trace** records all these invocations along with some additional information.

# Anomalies in Traces

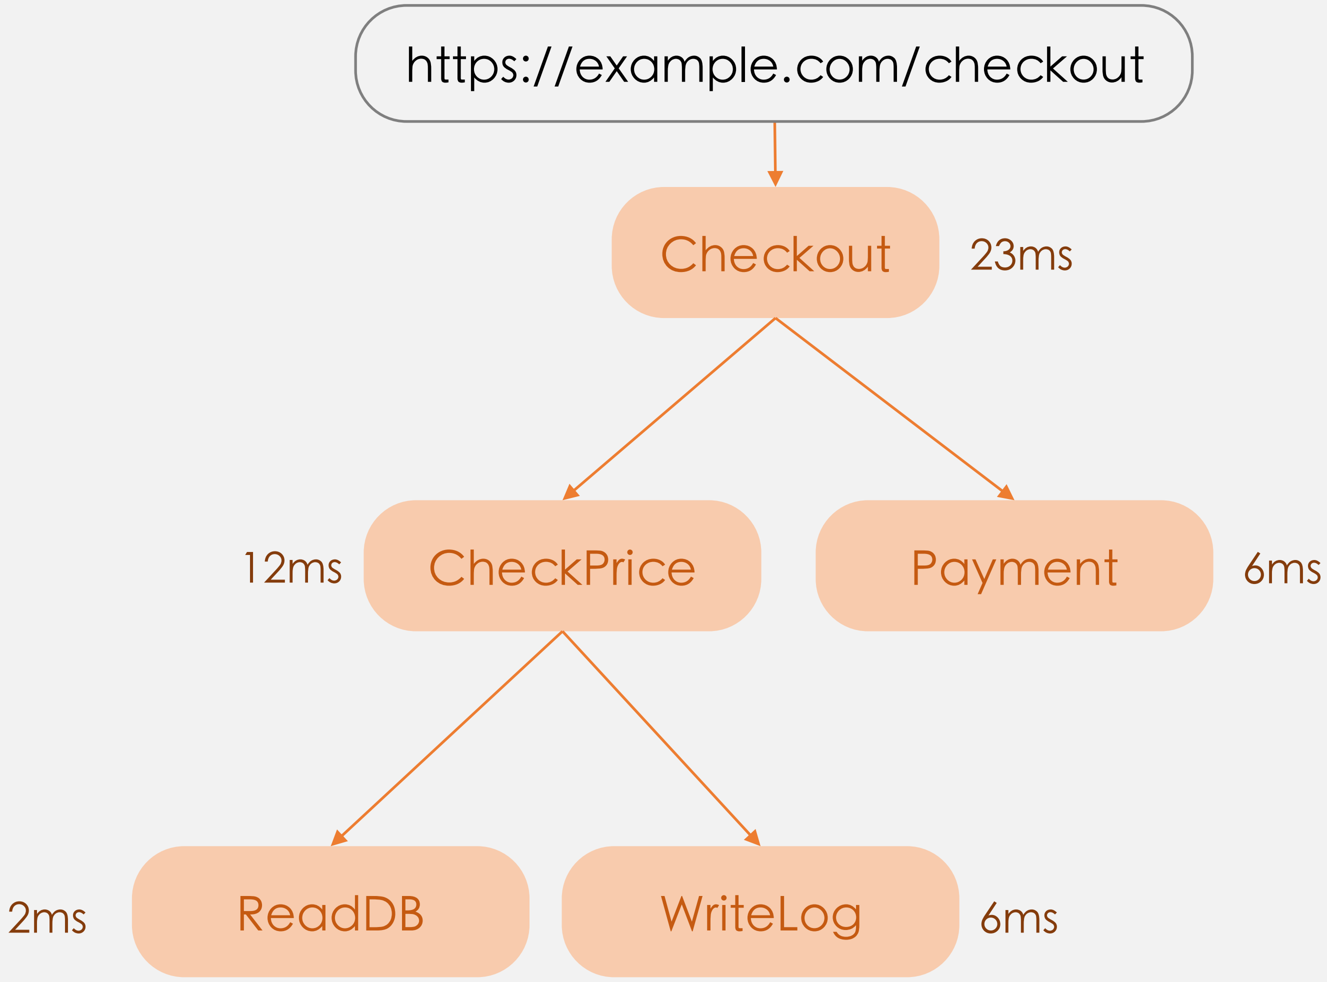

# Anomalies in Traces

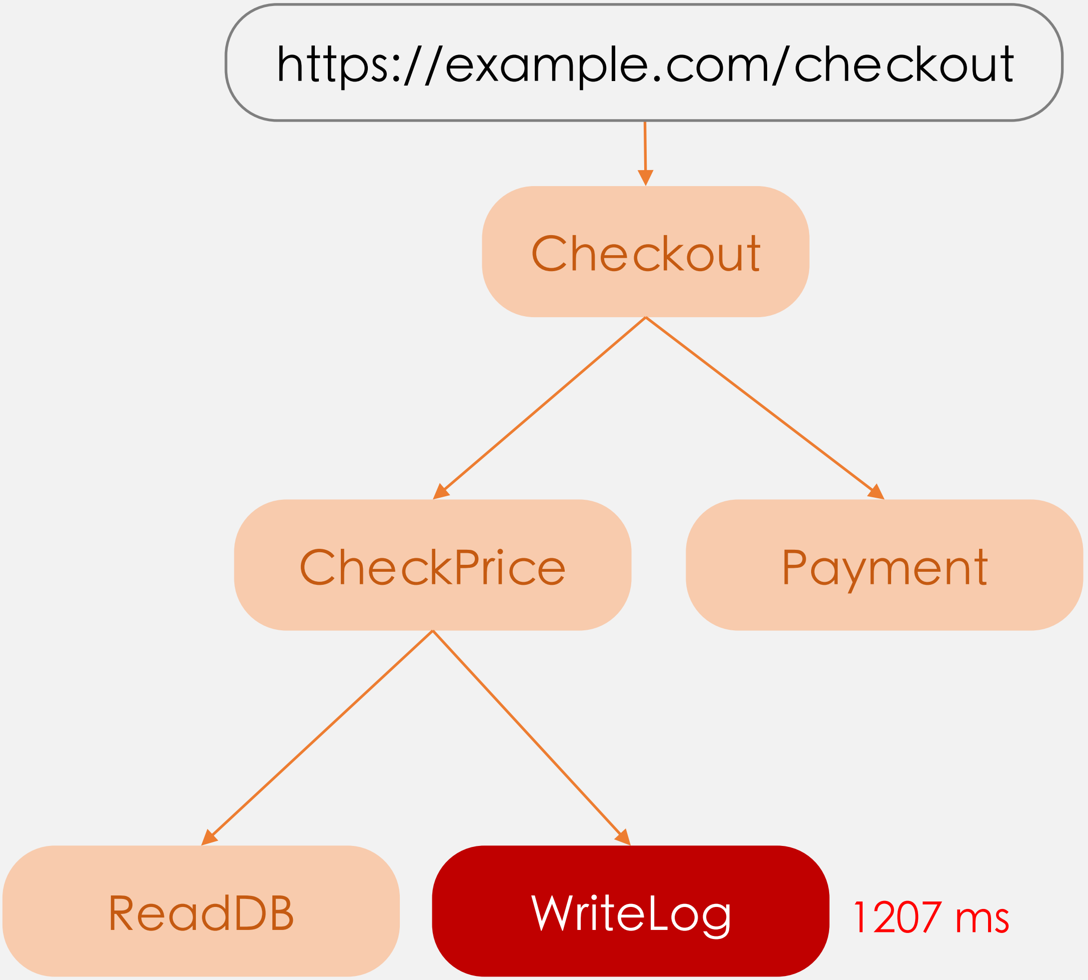

# Anomalies in Traces

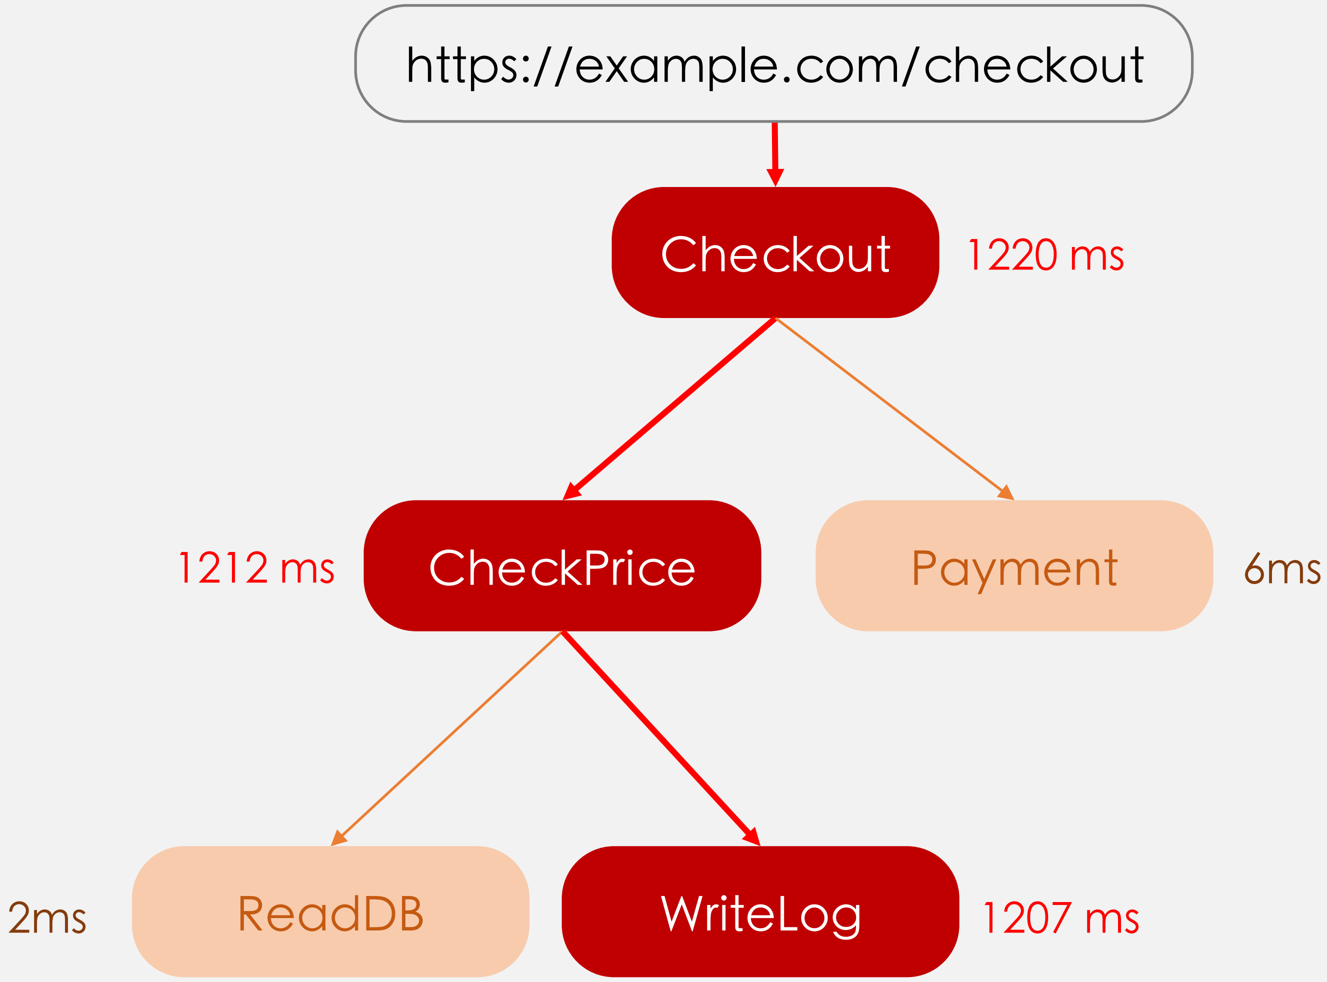


Trace records the path of fault propagation.

# Types of Trace Features

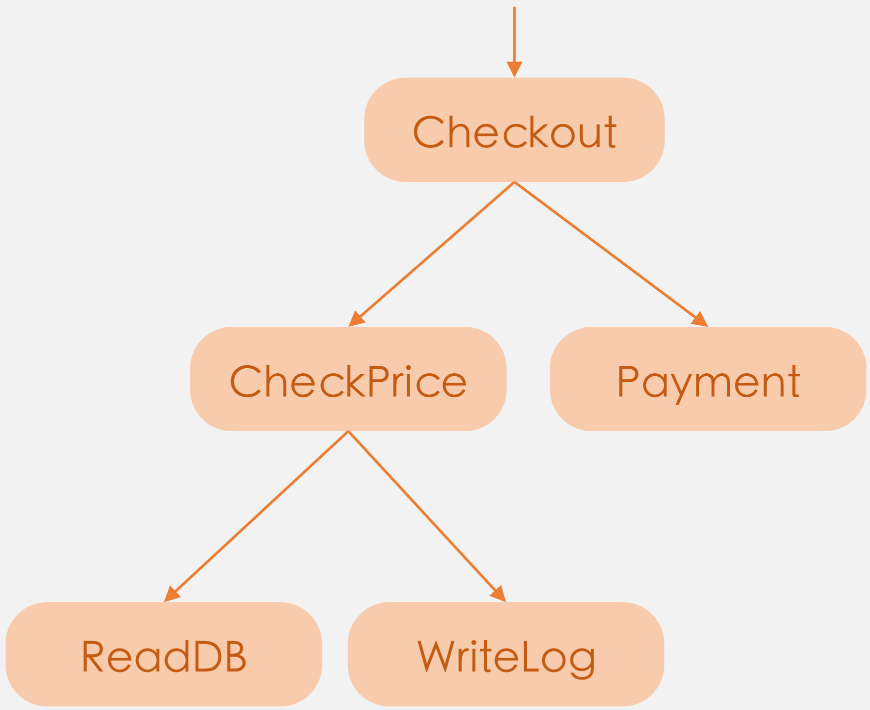


Structure Features

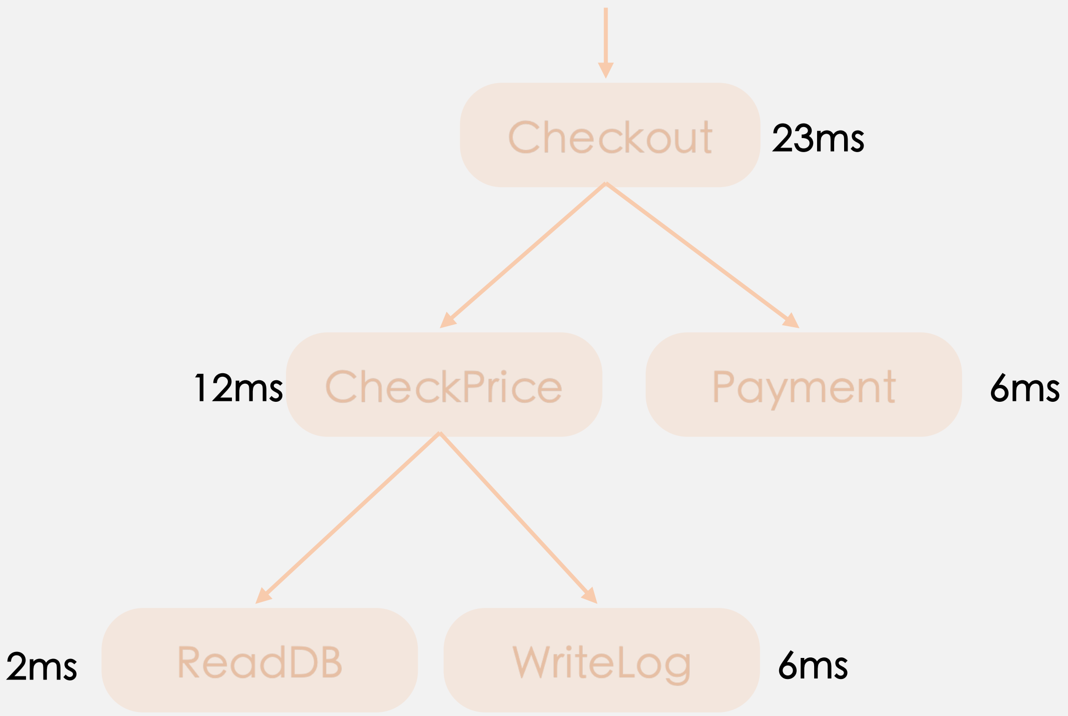


Time Features

# Types of Trace Anomalies

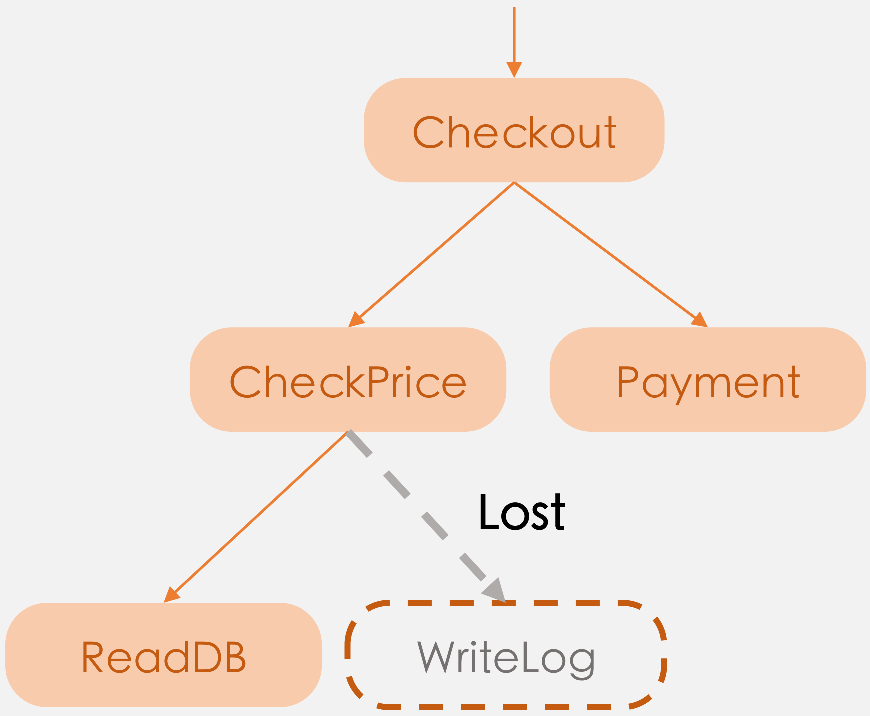


Structure Anomaly

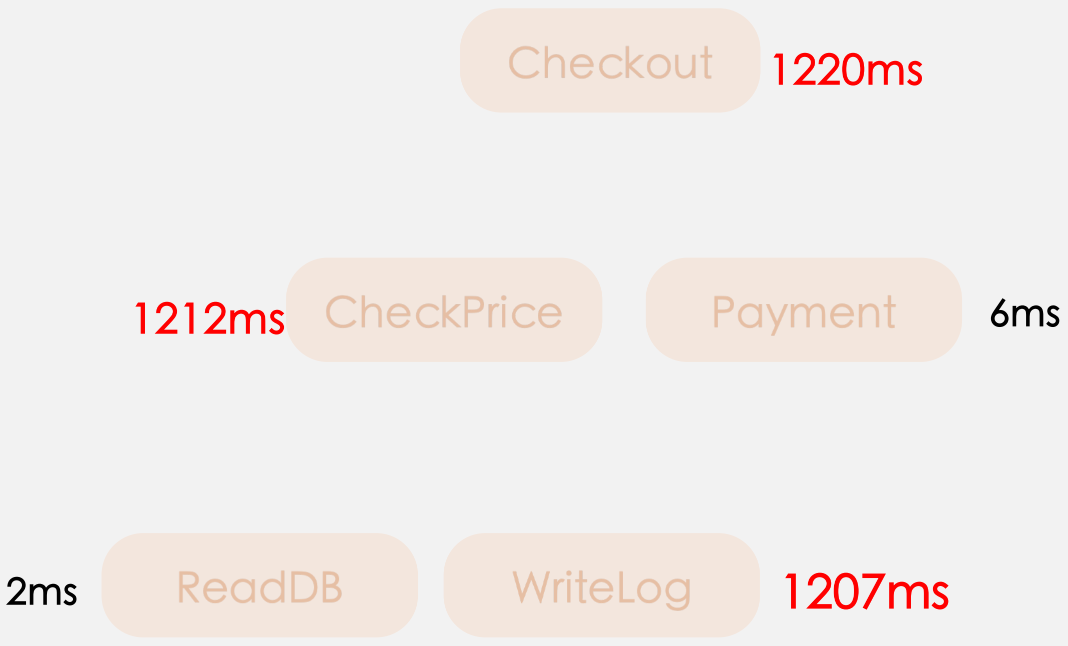


Time Anomaly

# Trace anomaly detection is not that easy!

Microservice invocations can be complex!

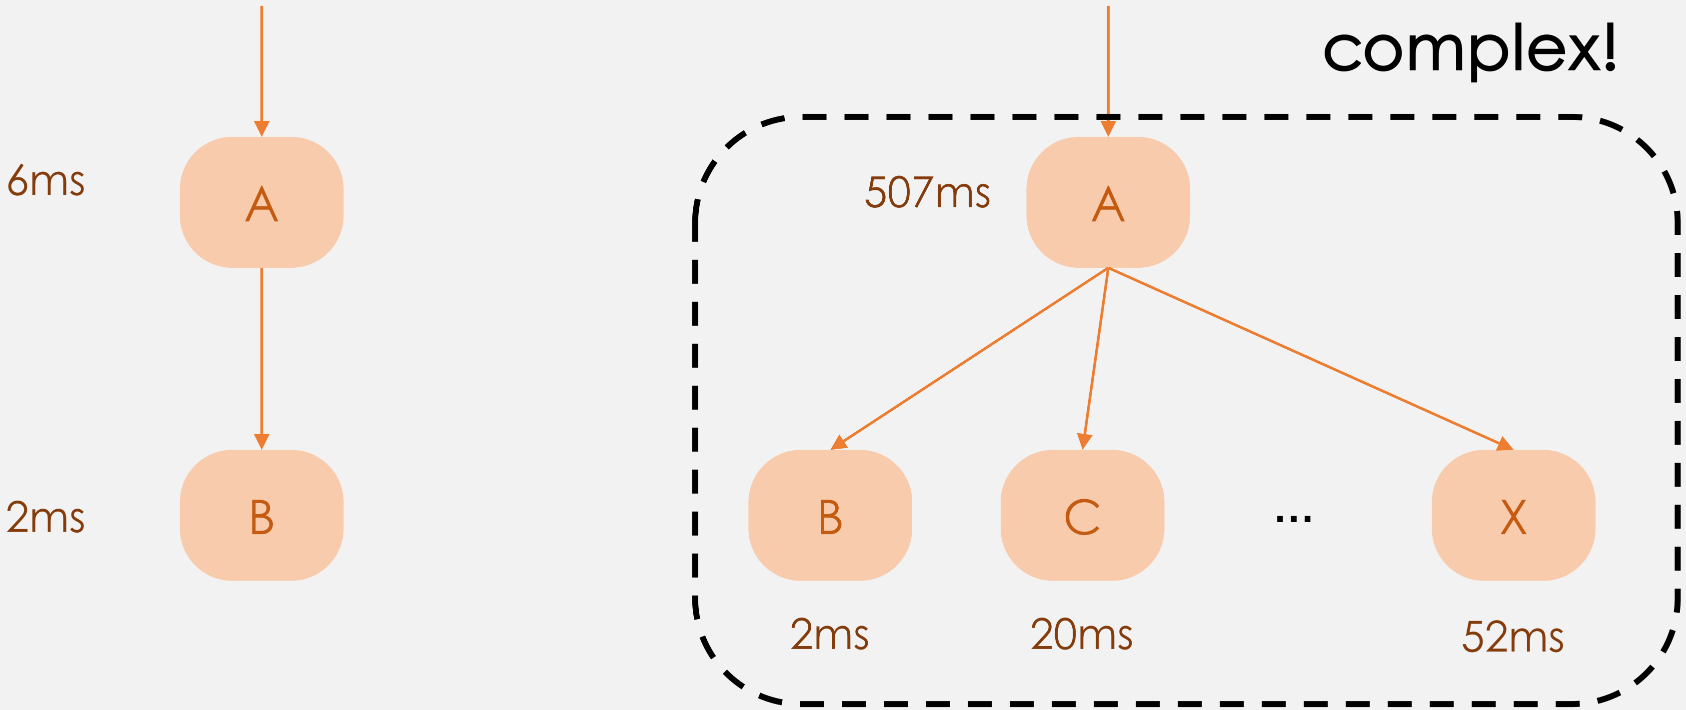


Single Downstream Invocation

Multiple Downstream Invocations

# Trace anomaly detection is not that easy!

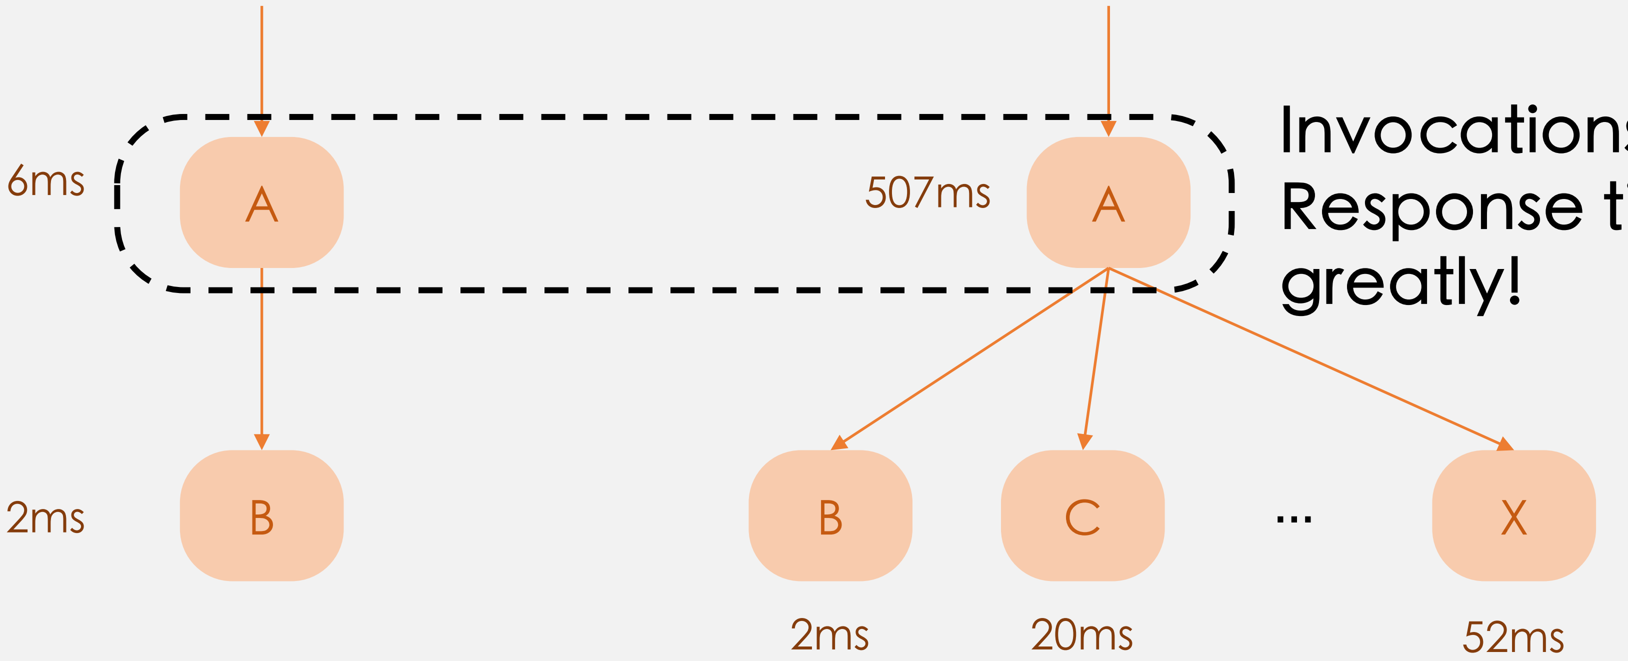


Invocations are correlated. Response time differs greatly!

Single Downstream Invocation

Multiple Downstream Invocations

# Trace anomaly detection is not that easy!


## How to capture the correlations?

Single Downstream Invocation

Multiple Downstream Invocations

# Trace Graph

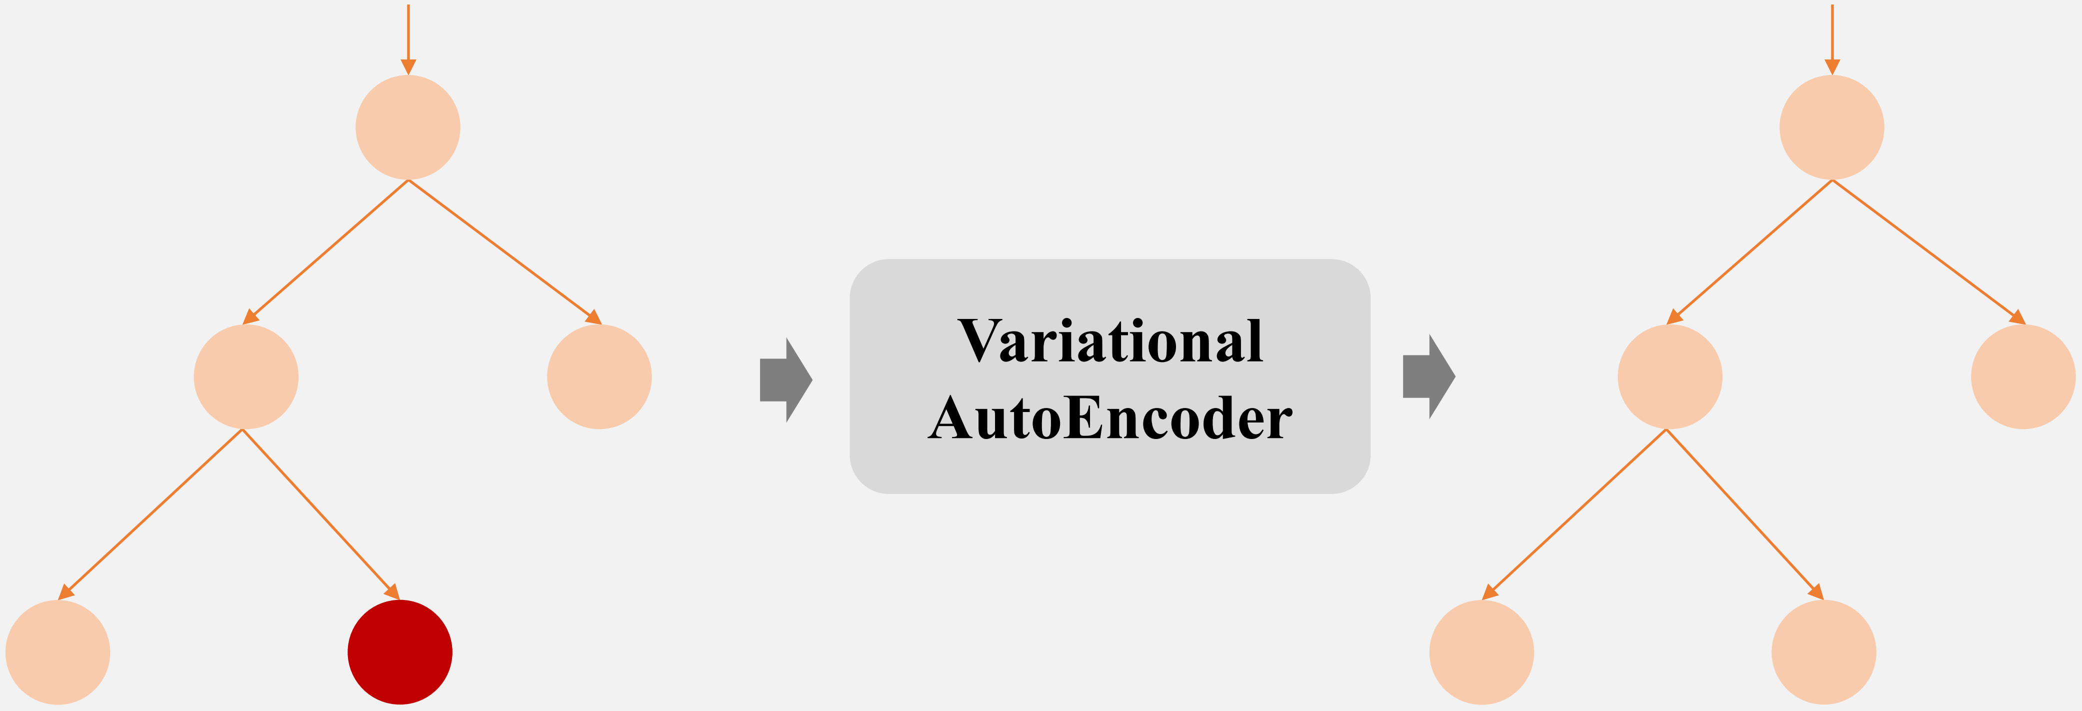


Model a Trace as a Graph

Reconstructed with VAE

# Dual-Variable Graph VAE

TraceVAE

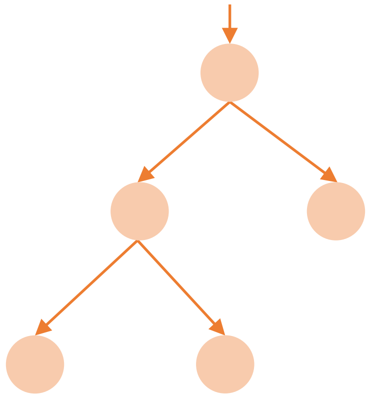

**A** Adjacency Matrix **X** Time Features **Y** Structure Features

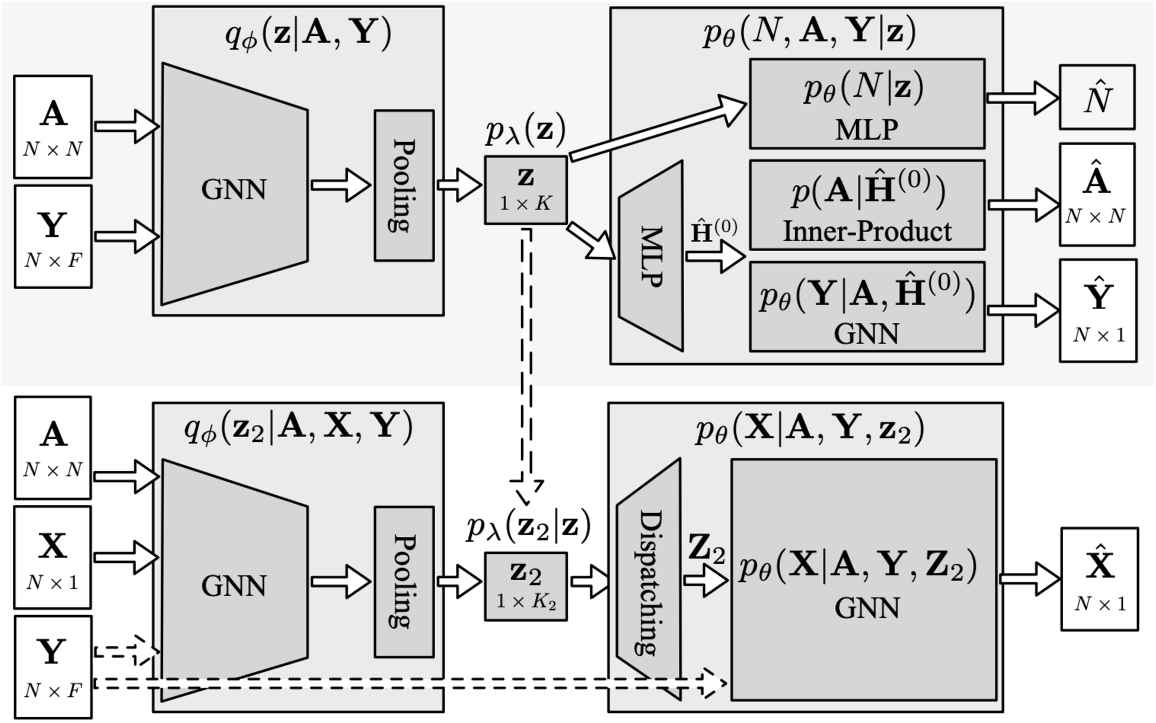


Structure VAE

Time VAE

Trace

Input Features

VAE Model

Reconstructed Features

# TraceVAE - GNN Layers

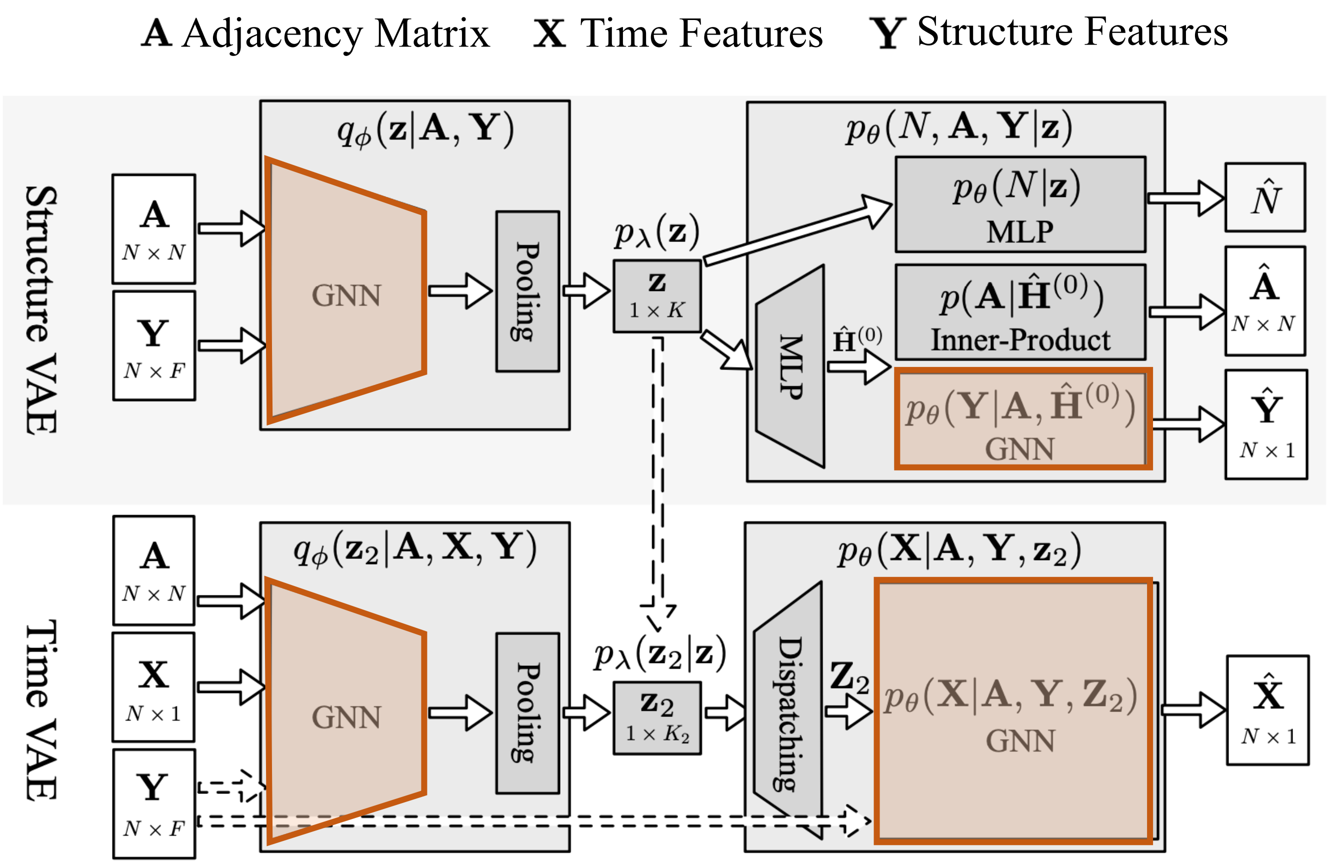


## GNN Layers

Capture the correlations in trace graphs with **Graph Neural Network (GNN)**

# TraceVAE - Dispatching Layer

## Dispatching Layer

Enhance the generalizability of time VAE

$z_2$ as a shared context for all nodes in a trace

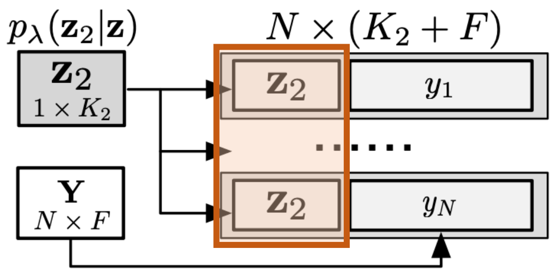


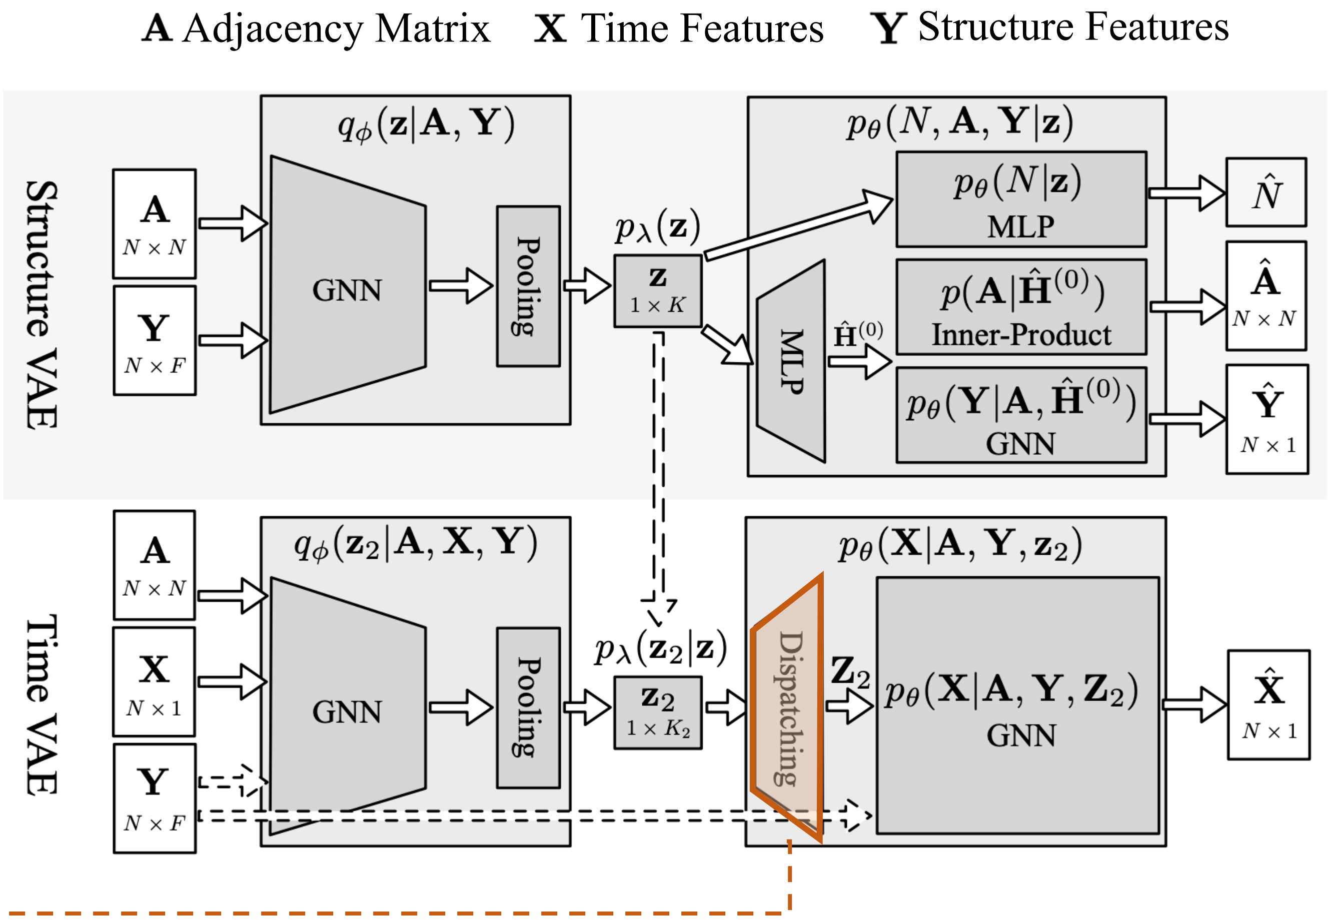

# TraceVAE - Anomaly Score

$$
\begin{aligned}
NLL_G &= -\log p_{\text{model}}(G) \\
&= -\log \mathbb{E}_{q_\phi(\mathbf{z},\mathbf{z}_2|G,N)}\left[\frac{p_{\theta,\lambda}(G,N,\mathbf{z},\mathbf{z}_2)}{q_\phi(\mathbf{z},\mathbf{z}_2|G,N)}\right]
\end{aligned}
$$

Negative Log-Likelihood (NLL)
as Anomaly Score

# TraceVAE - Anomaly Score

$$NLL_G = -\log p_{\text{model}}(G)$$

$$= -\log \mathbb{E}_{q_\phi(\mathbf{z},\mathbf{z}_2|G,N)}\left[\frac{p_{\theta,\lambda}(G,N,\mathbf{z},\mathbf{z}_2)}{q_\phi(\mathbf{z},\mathbf{z}_2|G,N)}\right]$$

Negative Log-Likelihood (NLL)
as Anomaly Score

- Traces have different numbers of nodes
- Is NLL suitable for trace anomalies?

# Inversion of NLL

**Normal:** Normal traces with structure-matched TraceVAE reconstructions

**Mismatched:** TraceVAE reconstructions are structured differently from the inputs

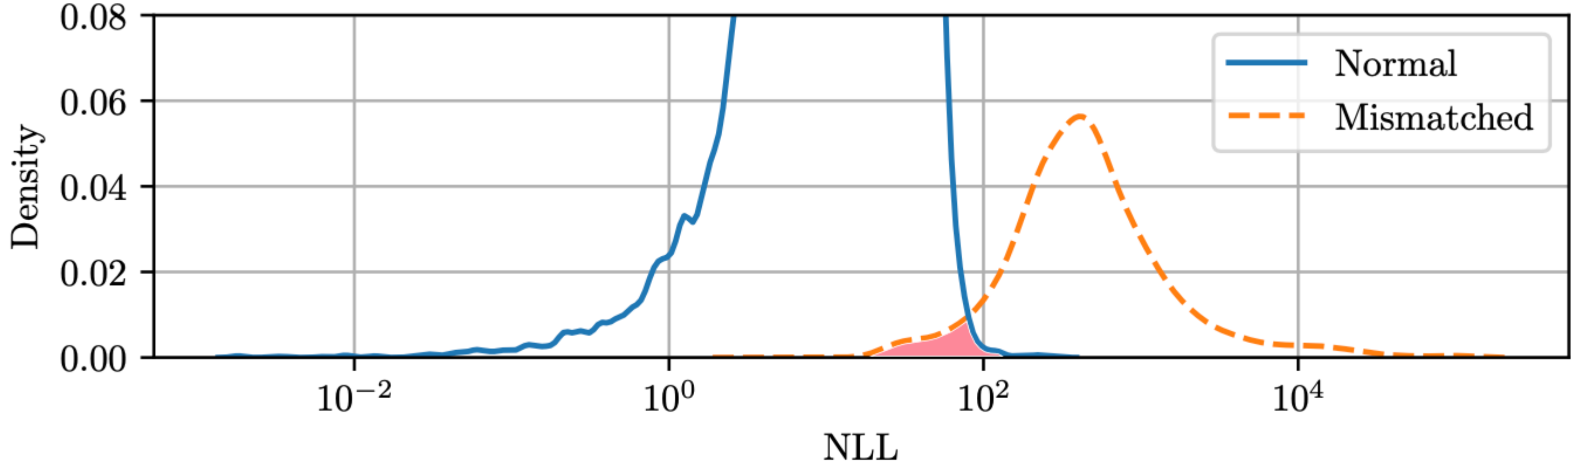


Large Intersection Area!

NLLs of many mismatched traces are even smaller than many normal traces.

# Go further with NLL

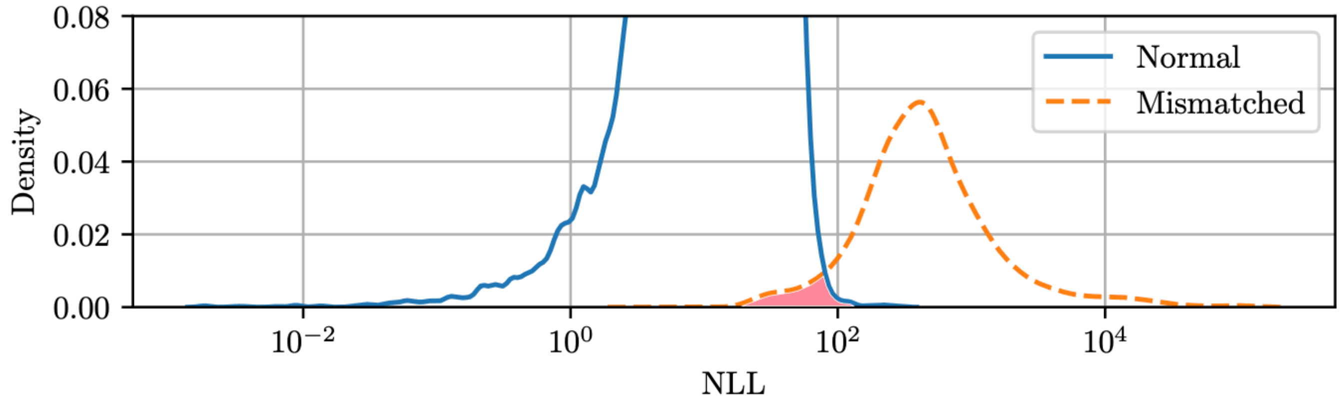


Inversion of NLL

$$\overline{NLL_x} \approx \mathrm{KLD}\left[p^\star \,\|\, p_{\mathrm{model}}\right] + \mathbb{H}\left[p^\star\right]$$

KL Divergence

Difference between real and model distributions

# Go further with NLL

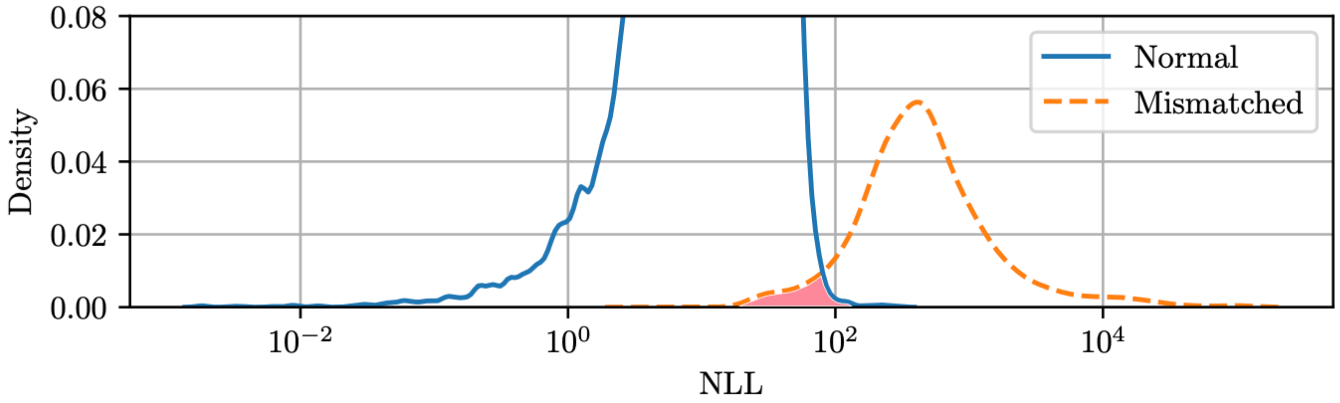


Inversion of NLL

$$\overline{NLL_x} \approx \text{KLD}[p^\star \,\|\, p_{\text{model}}] + \mathbb{H}\left[p^\star\right]$$

KL Divergence

Difference between real and model distributions

Entropy Term

Average Level of "Uncertainty"

# Entropy Gap

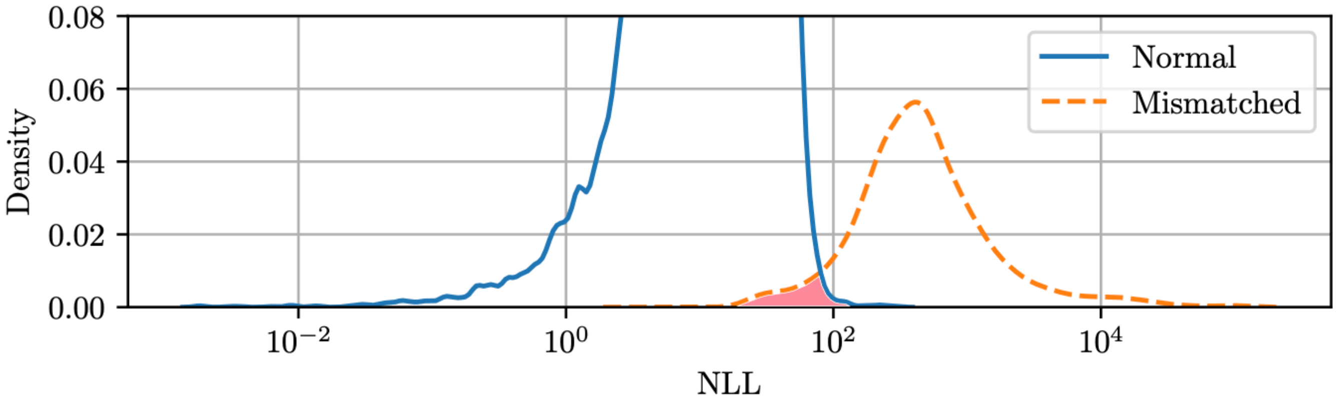


Inversion of NLL

$$\overline{NLL_x} \approx \mathrm{KLD}[p^\star \,\|\, p_{\mathrm{model}}] + \mathbb{H}\left[p^\star\right]$$

Average Level of "Uncertainty"

Traces with different sizes may have a large difference in their entropy.

⬇

Entropy Gap

# Entropy Gap

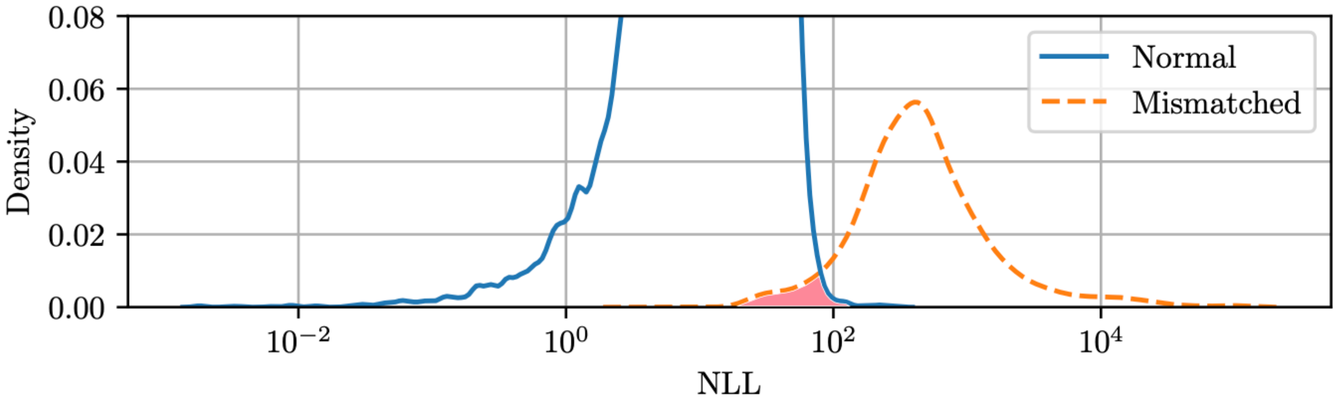


Inversion of NLL

$$\overline{NLL_x} \approx \mathrm{KLD}[p^\star \,\|\, p_{\text{model}}] + \mathbb{H}\left[p^\star\right]$$

Average Level of "Uncertainty"

Traces with different sizes may have a large difference in their entropy.

Entropy Gap (Unable to Estimate)

# Techniques to Reduce the Impact of Entropy Gap

## Bernoulli & Categorical Scaling

- Scaling the NLLs of anomalous nodes

## Node Count Normalization

- Normalize the entropy of traces with different #nodes

## Gaussian Std-Limit

- Further enlarge the NLLs of anomalous nodes

# Datasets

- 5 datasets collected from eBay’ s microservices system
- 1 dataset containing real online anomalies
- 4 datasets containing synthetic anomalies injection
- Collected from different business domains

# Comparison with Baselines

**Table 1: Best F-Scores of TraceVAE and the Baselines**

| | | $\mathcal{A}$ | | | $\mathcal{B}$ | | | $\mathcal{C}$ | | | $\mathcal{D}$ | | | $\mathcal{E}$ | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | Total | Struct | Time | Total | Struct | Time | Total | Struct | Time | Total | Struct | Time | Total | Struct | Time |
| Baselines | FSA | 0.664 | 0.497 | 0.497 | 0.737 | 0.583 | 0.583 | 0.813 | 0.685 | 0.685 | 0.527 | 0.358 | 0.358 | 0.199 | 0.090 | 0.134 |
| | LSTM-AD | 0.745 | 0.470 | 0.872 | 0.710 | 0.420 | 0.820 | 0.565 | 0.184 | 0.881 | 0.758 | 0.558 | 0.927 | 0.442 | 0.213 | 0.513 |
| | TraceAnomaly | 0.560 | 0.091 | 0.832 | 0.570 | 0.105 | 0.812 | 0.528 | 0.182 | 0.717 | 0.530 | 0.090 | 0.775 | 0.410 | 0.048 | 0.565 |
| | CRISP | 0.438 | 0.164 | 0.502 | 0.416 | 0.179 | 0.520 | 0.526 | 0.092 | 0.769 | 0.334 | 0.090 | 0.382 | 0.344 | 0.041 | 0.422 |
| | VGAE | 0.275 | *N/A* | 0.454 | 0.261 | *N/A* | 0.408 | 0.631 | *N/A* | 0.682 | 0.387 | *N/A* | 0.625 | 0.450 | *N/A* | 0.529 |
| TraceVAE Ablation Study | TraceVAE | **0.954** | 0.935 | **0.945** | 0.944 | 0.903 | **0.940** | 0.923 | 0.911 | **0.911** | **0.980** | 0.988 | **0.965** | **0.791** | **0.813** | **0.772** |
| | TraceVAE-FC | 0.936 | 0.889 | 0.938 | 0.925 | 0.877 | 0.936 | 0.915 | 0.903 | 0.907 | 0.975 | 0.983 | 0.959 | 0.729 | 0.742 | 0.677 |
| | TraceVAE-SingleZ | 0.854 | 0.849 | 0.829 | 0.888 | 0.921 | 0.816 | 0.919 | 0.881 | 0.894 | 0.946 | 0.943 | 0.931 | 0.632 | 0.702 | 0.507 |
| | TraceVAE-DimEx | 0.789 | 0.768 | 0.777 | 0.818 | 0.705 | 0.863 | 0.841 | 0.892 | 0.762 | 0.897 | 0.901 | 0.882 | 0.579 | 0.268 | 0.692 |
| Techniques to Reduce Entropy Gap | TraceVAE-NLL | 0.918 | 0.930 | 0.867 | 0.928 | 0.954 | 0.879 | 0.927 | 0.902 | 0.885 | 0.957 | 0.969 | 0.937 | 0.645 | 0.769 | 0.561 |
| | TraceVAE-BCScale | 0.925 | **0.967** | 0.868 | 0.931 | **0.971** | 0.878 | 0.925 | **0.925** | 0.883 | 0.965 | **0.990** | 0.935 | 0.662 | 0.813 | 0.558 |
| | TraceVAE-NCNorm | 0.918 | 0.877 | 0.916 | 0.904 | 0.891 | 0.882 | 0.873 | 0.798 | 0.880 | 0.965 | 0.964 | 0.955 | 0.687 | 0.731 | 0.627 |
| | TraceVAE-StdLimit | 0.940 | 0.910 | 0.928 | **0.947** | 0.924 | 0.934 | **0.930** | 0.892 | 0.904 | 0.963 | 0.957 | 0.957 | 0.732 | 0.769 | 0.680 |

- TraceVAE achieves the best results on all these datasets.

# Ablation Study

**Table 1: Best F-Scores of TraceVAE and the Baselines**

| | | $\mathcal{A}$ | | | $\mathcal{B}$ | | | $\mathcal{C}$ | | | $\mathcal{D}$ | | | $\mathcal{E}$ | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | Total | Struct | Time | Total | Struct | Time | Total | Struct | Time | Total | Struct | Time | Total | Struct | Time |
| Baselines | FSA | 0.664 | 0.497 | 0.497 | 0.737 | 0.583 | 0.583 | 0.813 | 0.685 | 0.685 | 0.527 | 0.358 | 0.358 | 0.199 | 0.090 | 0.134 |
| | LSTM-AD | 0.745 | 0.470 | 0.872 | 0.710 | 0.420 | 0.820 | 0.565 | 0.184 | 0.881 | 0.758 | 0.558 | 0.927 | 0.442 | 0.213 | 0.513 |
| | TraceAnomaly | 0.560 | 0.091 | 0.832 | 0.570 | 0.105 | 0.812 | 0.528 | 0.182 | 0.717 | 0.530 | 0.090 | 0.775 | 0.410 | 0.048 | 0.565 |
| | CRISP | 0.438 | 0.164 | 0.502 | 0.416 | 0.179 | 0.520 | 0.526 | 0.092 | 0.769 | 0.334 | 0.090 | 0.382 | 0.344 | 0.041 | 0.422 |
| | VGAE | 0.275 | *N/A* | 0.454 | 0.261 | *N/A* | 0.408 | 0.631 | *N/A* | 0.682 | 0.387 | *N/A* | 0.625 | 0.450 | *N/A* | 0.529 |
| TraceVAE Ablation Study | TraceVAE | **0.954** | 0.935 | **0.945** | 0.944 | 0.903 | **0.940** | 0.923 | 0.911 | **0.911** | **0.980** | 0.988 | **0.965** | **0.791** | **0.813** | **0.772** |
| | TraceVAE-FC | 0.936 | 0.889 | 0.938 | 0.925 | 0.877 | 0.936 | 0.915 | 0.903 | 0.907 | 0.975 | 0.983 | 0.959 | 0.729 | 0.742 | 0.677 |
| | TraceVAE-SingleZ | 0.854 | 0.849 | 0.829 | 0.888 | 0.921 | 0.816 | 0.919 | 0.881 | 0.894 | 0.946 | 0.943 | 0.931 | 0.632 | 0.702 | 0.507 |
| | TraceVAE-DimEx | 0.789 | 0.768 | 0.777 | 0.818 | 0.705 | 0.863 | 0.841 | 0.892 | 0.762 | 0.897 | 0.901 | 0.882 | 0.579 | 0.268 | 0.692 |
| Techniques to Reduce Entropy Gap | TraceVAE-NLL | 0.918 | 0.930 | 0.867 | 0.928 | 0.954 | 0.879 | 0.927 | 0.902 | 0.885 | 0.957 | 0.969 | 0.937 | 0.645 | 0.769 | 0.561 |
| | TraceVAE-BCScale | 0.925 | **0.967** | 0.868 | 0.931 | **0.971** | 0.878 | 0.925 | **0.925** | 0.883 | 0.965 | **0.990** | 0.935 | 0.662 | 0.813 | 0.558 |
| | TraceVAE-NCNorm | 0.918 | 0.877 | 0.916 | 0.904 | 0.891 | 0.882 | 0.873 | 0.798 | 0.880 | 0.965 | 0.964 | 0.955 | 0.687 | 0.731 | 0.627 |
| | TraceVAE-StdLimit | 0.940 | 0.910 | 0.928 | **0.947** | 0.924 | 0.934 | **0.930** | 0.892 | 0.904 | 0.963 | 0.957 | 0.957 | 0.732 | 0.769 | 0.680 |

- 2.6% - 23.6% improvement on time anomalies
- 0.5% - 5.7% improvement on structure anomalies

# NLL Distribution with Different Techniques

Normal

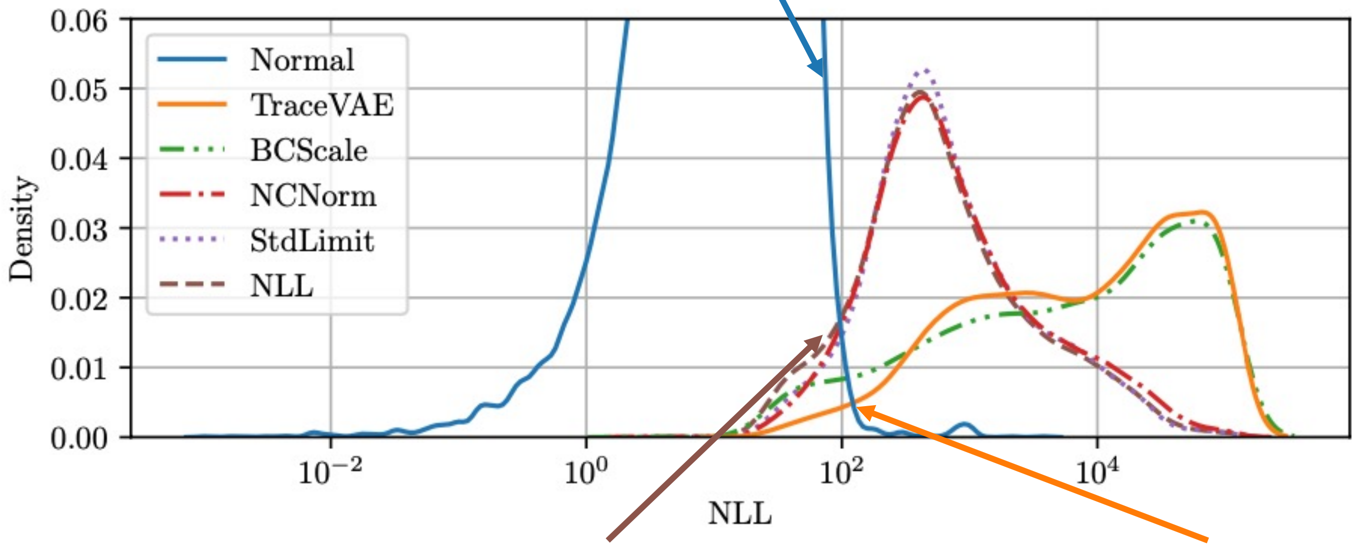


w/o Entropy Techniques (Mismatched)

TraceVAE (Mismatched)

## Intersection area becomes smaller

The proposed techniques effectively reduce the impact of the entropy gap.

# Contributions

- TraceVAE: Dual-Variable Graph VAE for Trace Modeling
- NLL Inversion and Entropy Gap in Trace Anomaly Detection
- Techniques to Reduce the Impact of Entropy Gap

# Thank you!

Code (https://github.com/NetManAIOps/TraceVAE.git)