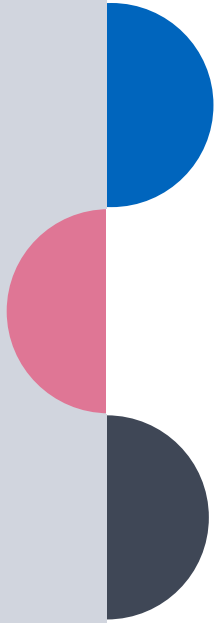

# Detecting Anomalous Event Sequences with Temporal Point Processes

**Oleksandr Shchur**[1], Ali Caner Türkmen[2], Tim Januschowski[2],

Jan Gasthaus[2], Stephan Günnemann[1]

[1]Technical University of Munich, [2]Amazon Research

Neural Information Processing Systems 2021

# Temporal point process (TPP)

Probabilistic model for continuous-time event data

- Server logs
- User activity traces
- Financial transactions

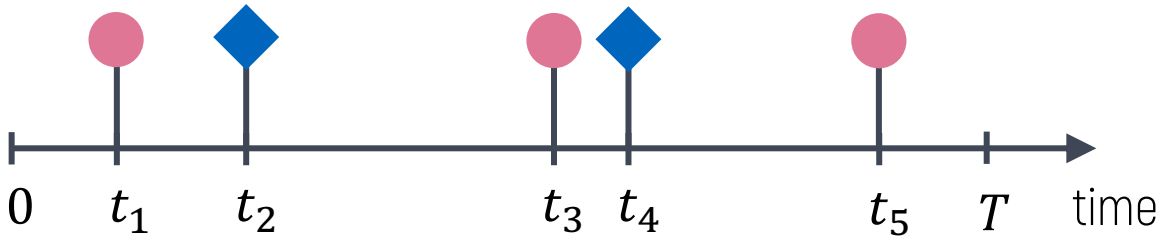

# Out-of-distribution (OoD) detection

Given: Normal sequences $\mathcal{D}_{\text{train}} = \{X_1, \ldots, X_M\}$ that were generated by some unknown TPP $\mathbb{P}_{\text{data}}$

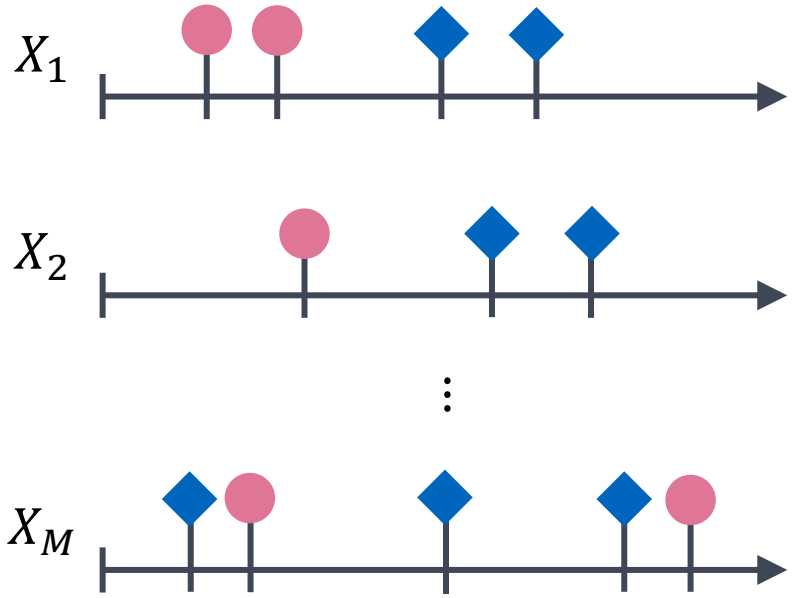


Question: Is the new sequence $X$ normal or anomalous?

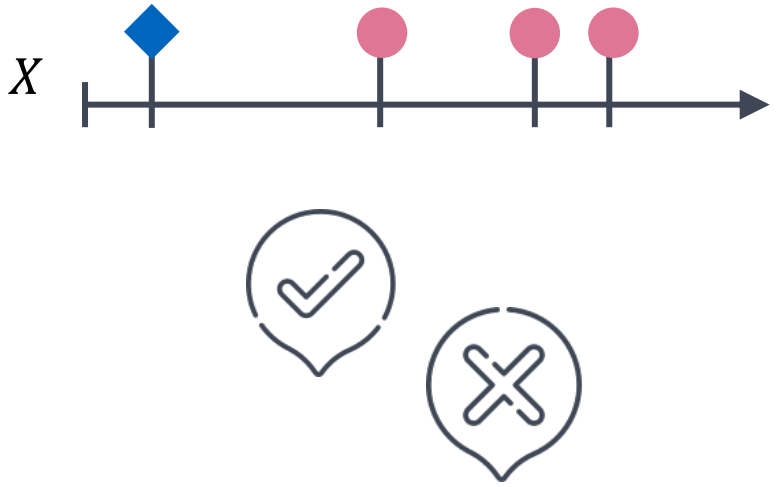

# Out-of-distribution (OoD) detection

Given: $\{X_1, \ldots, X_M\} \overset{\text{i.i.d.}}{\sim} \mathbb{P}_{\text{data}}$

Question:

$H_0: X \sim \mathbb{P}_{\text{data}}$

$H_1: X \sim \mathbb{Q}$ for some $\mathbb{Q} \neq \mathbb{P}_{\text{data}}$

# Goodness-of-fit (GoF) testing

Given: known distribution $\mathbb{P}_{\text{model}}$

Question:

$H_0: X \sim \mathbb{P}_{\text{model}}$

$H_1: X \sim \mathbb{Q}$ for some $\mathbb{Q} \neq \mathbb{P}_{\text{model}}$

The two problems are similar ⇒ We can use tools for GoF testing to perform OoD detection

# GoF hypothesis test

$$H_0: X \sim \mathbb{P}_{\text{model}} \quad \text{or} \quad H_1: X \sim \mathbb{Q} \text{ for } \mathbb{Q} \neq \mathbb{P}_{\text{model}}$$

1. Pick a test statistic $s: \mathcal{X} \to \mathbb{R}$
2. Compute the p-value for the given realization $x$ of $X$

$$p_s(x) = 2 \times \min\{\Pr(s(X) \leq s(x)|H_0), 1 - \Pr(s(X) \leq s(x)|H_0)\}$$

3. Reject $H_0$ if p-value is below a threshold $\alpha$

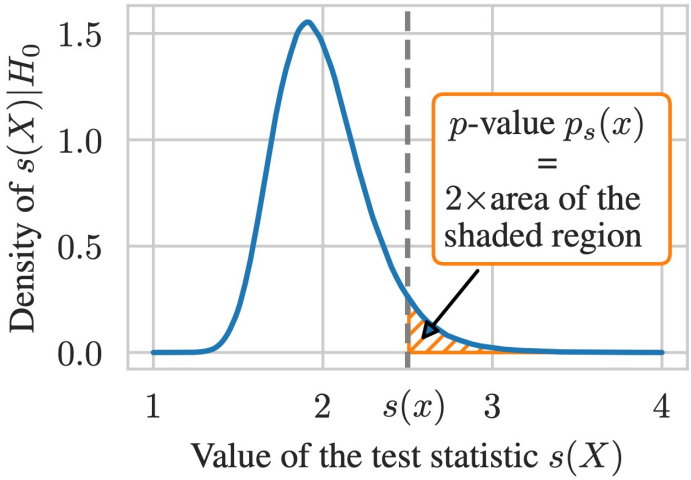

# Back to OoD detection

- OoD detection hypothesis test

$$H_0: X \sim \mathbb{P}_{\text{data}} \quad \text{or} \quad H_1: X \sim \mathbb{Q} \text{ for } \mathbb{Q} \neq \mathbb{P}_{\text{data}}$$

given $\mathcal{D}_{\text{train}} = \{X_1, \ldots, X_M\}$ drawn i.i.d. from $\mathbb{P}_{\text{data}}$

- Problem How can we compute the p-value?

$$p_s(x) = 2 \times \min\{\Pr(s(X) \leq s(x)|H_0), 1 - \Pr(s(X) \leq s(x)|H_0)\}$$

- No assumptions about $\mathbb{P}_{\text{data}} \Rightarrow$ cannot compute CDF of $s(X)|H_0$ analytically

- Solution: Use the empirical distribution (EDF) of the statistic $s(X)$ on $\mathcal{D}_{\text{train}}$

# How to define a test statistic for TPPs?

- A TPP is uniquely defined by its compensator $\Lambda^*$
- The compensator converts any TPP realization
$$X = (t_1, \dots, t_N)$$
into a sample from the standard Poisson process
$$Z = \left(\Lambda^*(t_1), \dots, \Lambda^*(t_N)\right)$$
- Use existing test statistics for Poisson processes on $Z$ to define a statistic $s(X)$ for an arbitrary TPP
  - e.g., Kolmogorov–Smirnov, $\chi$-squared, sum-of-squared-spacings

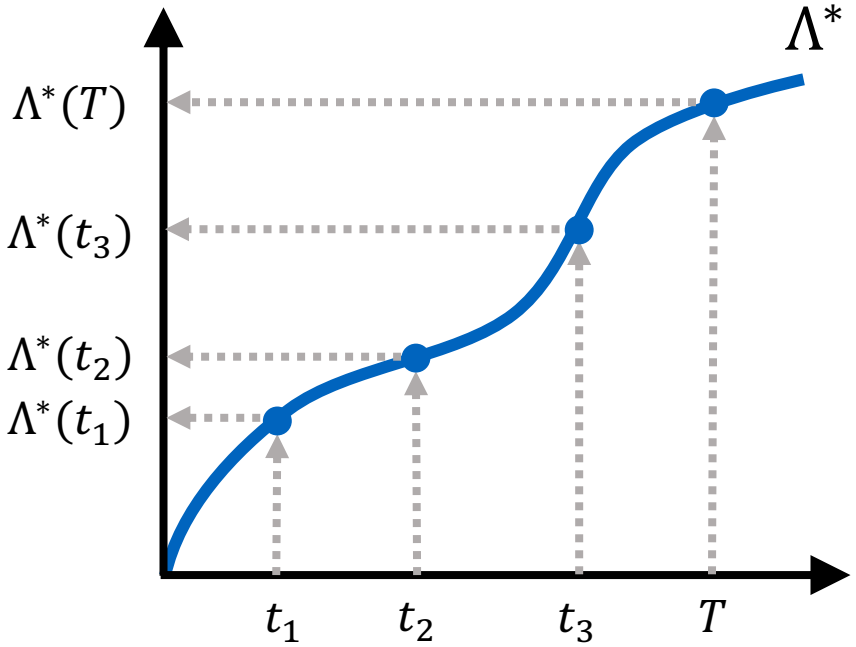

# Putting everything together

1. Fit a TPP model on $\mathcal{D}_{\text{train}} = \{X_1, \ldots, X_M\}$
2. Define statistic $s(X)$ based on the compensator $\Lambda^*$ of the learned TPP
3. Approximate CDF of $s(X)|X \sim \mathbb{P}_{\text{data}}$ with the EDF on $\mathcal{D}_{\text{train}}$
4. For a test sequence $X$, compute the p-value using the EDF
5. Label $X$ as anomalous if p-value is less than chosen $\alpha$

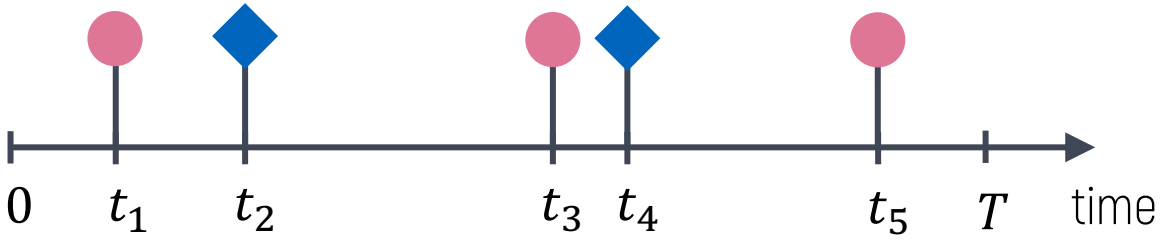

# Experiments

- Accurate detection of OoD sequences in simulated and real-world data
- Also works for marked sequences with multiple event types

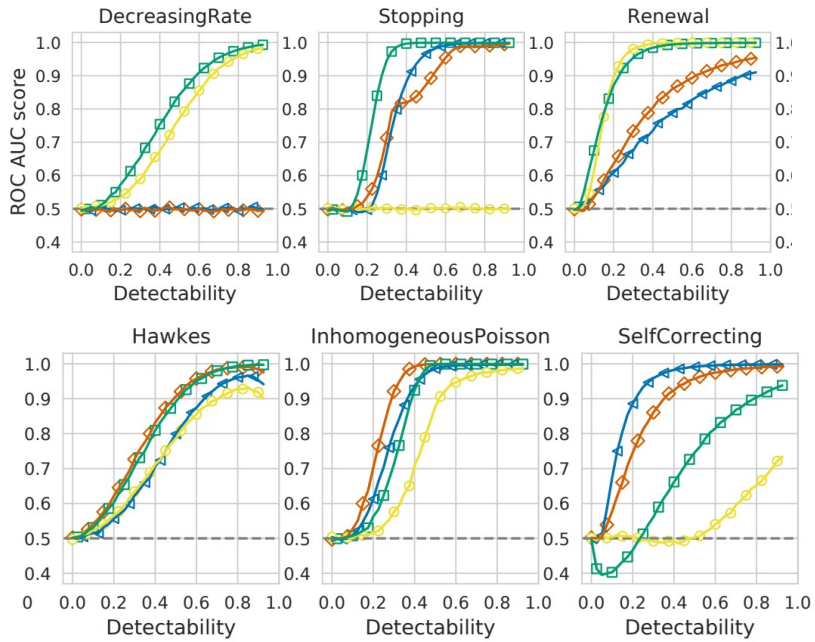


| | KS arrival | KS inter-event | Chi-squared | Log-likelihood | 3S statistic |
|---|---|---|---|---|---|
| LOGS — Packet corruption (1%) | 57.4 ± 1.7 | 62.1 ± 0.9 | 66.6 ± 1.8 | 75.9 ± 0.1 | **95.5** ± 0.3 |
| LOGS — Packet corruption (10%) | 59.2 ± 2.3 | 97.8 ± 0.6 | 59.1 ± 2.3 | 99.0 ± 0.0 | **99.4** ± 0.1 |
| LOGS — Packet duplication (1%) | 81.1 ± 5.2 | 82.8 ± 5.0 | 74.6 ± 6.5 | 88.1 ± 0.1 | **90.9** ± 0.3 |
| LOGS — Packet delay (frontend) | 95.6 ± 1.2 | 98.9 ± 0.4 | **99.3** ± 0.1 | 90.9 ± 0.0 | 97.6 ± 0.1 |
| LOGS — Packet delay (all services) | **99.8** ± 0.0 | 94.7 ± 1.1 | **99.8** ± 0.0 | 96.1 ± 0.0 | 99.6 ± 0.1 |
| STEAD — Anchorage, AK | 59.6 ± 0.2 | 79.7 ± 0.1 | 67.4 ± 0.2 | 88.0 ± 0.1 | **88.3** ± 0.6 |
| STEAD — Aleutian Islands, AK | 53.8 ± 0.5 | 88.8 ± 0.3 | 62.2 ± 0.9 | 97.0 ± 0.0 | **99.8** ± 0.0 |
| STEAD — Helmet, CA | 59.1 ± 0.9 | **98.7** ± 0.0 | 70.0 ± 0.6 | 96.9 ± 0.0 | 92.6 ± 0.3 |

# Summary

- OoD detection and GoF testing are related but different problems
- We can use GoF statistics for TPPs to detect anomalous event sequences
- The general framework extends to other data types, such a time series

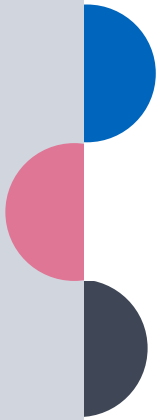

Code:
github.com/shchur/tpp-anomaly-detection