

TODS: An Automated Time Series Outlier Detection System

Kwei-Herng Lai^{1*}, Daochen Zha^{1*}, Guanchu Wang¹, Junjie Xu¹, Yue Zhao², Devesh Kumar¹, Yile Chen¹, Purav Zumkhawaka¹, Minyang Wan¹, Diego Martinez¹, Xia Hu¹

¹Department of Computer Science and Engineering, Texas A&M University

²H. John Heinz III College, Carnegie Mellon University

{khlai, daochen.zha, hegsns, devesh, yilechen17, puravzum, w1996, dmartinez05, xiahu}@tamu.edu, xu.jj724@gmail.com, zhaoy@cmu.edu

Abstract

We present TODS, an automated Time Series Outlier Detection System for research and industrial applications. TODS is a highly modular system that supports easy pipeline construction. The basic building block of TODS is primitive, which is an implementation of a function with hyperparameters. TODS currently supports 70 primitives, including data processing, time series processing, feature analysis, detection algorithms, and a reinforcement module. Users can freely construct a pipeline using these primitives and perform end-to-end outlier detection with the constructed pipeline. TODS provides a Graphical User Interface (GUI), where users can flexibly design a pipeline with drag-and-drop. Moreover, a data-driven searcher is provided to automatically discover the most suitable pipelines given a dataset. TODS is released under Apache 2.0 license at <https://github.com/datamllab/tods>. A video is available on YouTube¹.

Introduction

Outlier detection aims to identify unexpected or rare instances in data. As one of the most important tasks of data analysis, outlier detection has various real-world applications. For example, Yahoo (Laptev, Amizadeh, and Flint 2015) and Microsoft (Ren et al. 2019) have built their own time-series outlier detection services to monitor their business data and trigger alerts for outliers.

While machine learning models have shown promise in time series outlier detection, building an effective time series outlier detection pipeline heavily relies on human expertise. Due to the unique data characteristics of different application scenarios, one often needs expensive laboring trials to implement and identify the most suitable processing modules, detection algorithms, and hyperparameters for a specific task, which significantly impedes the real-world application of time series outlier detection.

To bridge this gap, we aim to save human efforts to build effective pipelines for time series outlier detection. Specifically, we implement commonly used time series analysis modules and detection algorithms under unified interfaces

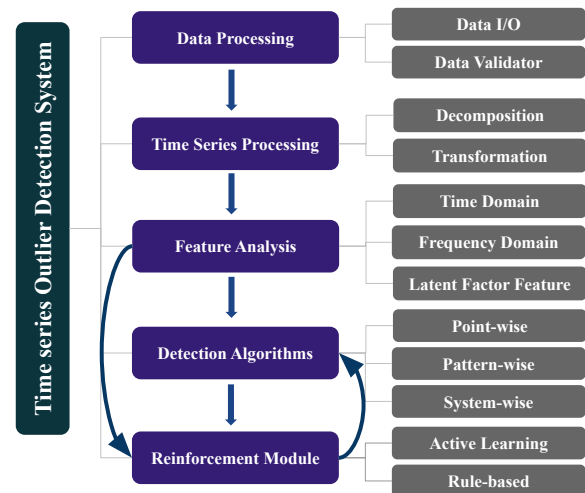


Figure 1: System overview. TODS provides end-to-end pipeline construction from data processing to detection algorithms. It also provides a reinforcement module to incorporate human knowledge into predictions.

so that users can easily and flexibly try different pipelines. In addition, we provide data-driven searchers to help the users to automatically discover suitable pipelines. While our previous work has delivered an open-source toolbox for outlier detection (Zhao, Nasrullah, and Li 2019) and explored automated machine learning for outlier detection (Li et al. 2020b,a), they provide limited automation for pipeline construction and are not designed for time series data. Whereas, we provide a highly modular system with a wide range of algorithm choices, specifically designed for various time series outlier detection scenarios.

We present an automated Time series Outlier Detection System called **TODS**. It is designed under D3M infrastructure (Milutinovic et al. 2020) for data-driven model discovery via automated machine learning. It currently supports 70 primitives for data processing, time series processing, feature analysis, outlier detection, and incorporating human knowledge. It can be applied to various application scenarios, including point-wise, pattern-wise and system-wise detection. The goal of TODS is providing an end-to-end solution for real-world time series outlier detection.

*Those authors contribute equally to this project

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://youtu.be/H0bBXuDUe7s>

Module	Number	Examples
Data Processing	7	TimestampValidation
Time Series Processing	9	SeasonalityDecomposition
Feature Analysis	30	AutoCorrelation, NMF
Detection Algorithms	23	LSTM, IForest
Reinforcement Module	1	RuleBasedFilter
Total	70	-

Table 1: The number of primitives implemented in each module of TODS. We only list some example algorithms in each module due to space limitation.

```

from tods import generate_dataset
from tods import load_pipeline
from tods import evaluate_pipeline

# Read data and generate dataset
df = pd.read_csv(table_path)
dataset = generate_dataset(df, target_index)

# Load pipeline
pipeline = load_pipeline(pipeline_path)

# Run the pipeline
evaluate_pipeline(dataset, pipeline, metric)

```

Example 1: Code Snippet of function calls.

TODS System

An overview of TODS is shown in Figure 1. The basic building block in TODS is *primitive*, which is an implementation of a function with some hyperparameters (e.g., Auto-Encoder algorithm). A *pipeline* is defined as a Directed Acyclic Graph (DAG), where each step represents a primitive. A typical pipeline has four primitive steps, including data processing, time series processing, feature analysis, and detection algorithms. An optional reinforcement primitive can be added to the end of a pipeline to incorporate predefined rules. The primitives developed in each module are shown in Table 1. Various pipelines can be constructed based on these primitives in different application scenarios. To enable easy pipeline construction, we provide a Graphical User Interface (GUI) to build and evaluate pipelines with drag-and-drop. In addition, a data-driven searcher is provided to automatically discover the suitable pipelines for a given task.

Interface Design

The interface is built upon Axolotl², which is a high-level abstraction of D3M developed by our team. The pipeline is described with pipeline language. More details about pipeline language can be found in (Milutinovic et al. 2020). To run a pipeline, we take a CSV file, a target index, and a pipeline description file as input, where the CSV file consists of timestamps, several features, and a target. A minimum example of evaluating a pipeline is shown in Example 1, where `pipeline_path` is the path of pipeline description,

²<https://gitlab.com/axolotl1/axolotl>

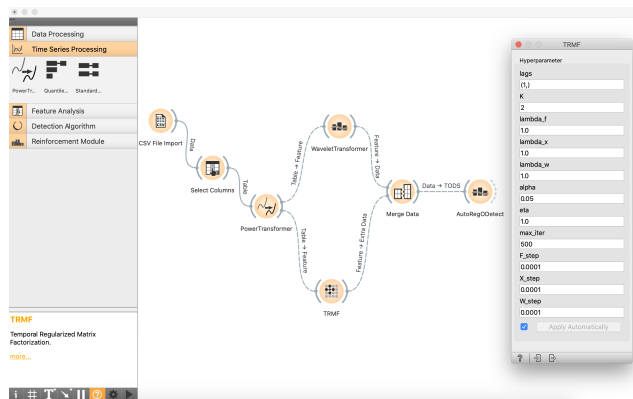


Figure 2: An illustration of the GUI. Users can easily build and evaluate a pipeline with drag-and-drop.

`table_path` is the path of CSV file, and `target_index` describes which column is the target.

Live and Interactive Part

In the demo, we will show the Graphical User Interface (GUI) of TODS. The GUI is implemented based on Orange (Demšar et al. 2004). Figure 2 shows an example of how users can build a pipeline with GUI. The primitives in TODS are categorized into several tabs. Each primitive will have a corresponding icon and a text description. Users can move the primitives of interest to the canvas and connect them with lines. Users can also easily configure the hyperparameters of each primitive by double-clicking the icon, as shown in the pop-up window in the figure. TODS will automatically generate the pipeline language based on the structure in GUI. The generated pipeline will then be passed to the backend to perform time series outlier detection.

Automated Pipeline Discovery

Beyond manual pipeline construction, TODS provides data-driven searchers to automatically find good pipelines for a given task. In addition to the detection pipeline, TODS provides a data preparation pipeline for data splitting, such as k-fold splitting, and a scoring pipeline to evaluate the performance based on a metric, such as F1 score. The data-driven searcher takes a dataset as input to automatically discover a detection pipeline that can achieve the best performance on a given dataset. Then the discovered pipeline can be outputted for real-world deployment.

Conclusions and Future Work

TODS is a highly modular open-source system for automated time series outlier detection. We build TODS upon our previous research and open-source efforts in automated machine learning and outlier detection, hoping that it can enable end-users to quickly deploy a pipeline for a specific application scenario. In the future, we will add more primitives and design more efficient searchers for pipeline discovery. We will also improve the reinforcement module with learning-based active learning (Zha et al. 2020).

References

- Demšar, J.; Zupan, B.; Leban, G.; and Curk, T. 2004. Orange: From experimental machine learning to interactive data mining. In *PKDD*, 537–539. Springer.
- Laptev, N.; Amizadeh, S.; and Flint, I. 2015. Generic and scalable framework for automated time-series anomaly detection. In *KDD*.
- Li, Y.; Chen, Z.; Zha, D.; Zhou, K.; Jin, H.; Chen, H.; and Hu, X. 2020a. AutoOD: Automated Outlier Detection via Curiosity-guided Search and Self-imitation Learning. *arXiv preprint arXiv:2006.11321*.
- Li, Y.; Zha, D.; Venugopal, P.; Zou, N.; and Hu, X. 2020b. PyODDS: An End-to-end Outlier Detection System with Automated Machine Learning. In *WWW*.
- Milutinovic, M.; Schoenfeld, B.; Martinez-Garcia, D.; Ray, S.; Shah, S.; and Yan, D. 2020. On Evaluation of AutoML Systems. In *ICML Workshop*.
- Ren, H.; Xu, B.; Wang, Y.; Yi, C.; Huang, C.; Kou, X.; Xing, T.; Yang, M.; Tong, J.; and Zhang, Q. 2019. Time-Series Anomaly Detection Service at Microsoft. In *KDD*.
- Zha, D.; Lai, K.-H.; Wan, M.; and Hu, X. 2020. Meta-AAD: Active Anomaly Detection with Deep Reinforcement Learning. *arXiv preprint arXiv:2009.07415*.
- Zhao, Y.; Nasrullah, Z.; and Li, Z. 2019. PyOD: A python toolbox for scalable outlier detection. *JMLR*.