

BacktrackSTL: Ultra-Fast Online Seasonal-Trend Decomposition with Backtrack Technique



Haoyu Wang, Hongke Guo, Zhaoliang Zhu, You Zhang, Yu Zhou, Xudong Zheng Alibaba Cloud, Alibaba Group

INTRODUCTION

Online seasonal-trend decomposition (STD)

Decompose a periodic time series to trend, seasonality, and **residual** components incrementally.

METHODS





Huge data scale => Extremely high time efficiency

Time efficiency has the highest priority, even taking precedence over accuracy.

Algorithm	Trend	Seasonality	Outlier	Online
	Jump	Shift	Tolerance	Complexity
STL	No	No	No	-
TBATS	Yes	No	No	-
STR	No	Yes	Yes	-
SSA	No	No	No	-
RobustSTL	Yes	Yes	Yes	_
OnlineSTL	No	No	No	O(T)
OneShotSTL	Yes	Yes	No	O(I)
BacktrackSTL	Yes	Yes	Yes	O(1)

MOTIVATION

DEPLOYMENT

Analysis on L1-norm optimization of RobustSTL

- Time efficiency: bottleneck of the algorithm (96.7%)
- Accuracy: robust to both outlier and trend jump



Insight to improve the time efficiency

Combining various low-complexity methods may be more effective in addressing complex requirements than using a single high-complexity method.

Deployment in Alibaba Cloud ECS

- MaxCompute: a daily job based on UDF to identify the period lengths of all metric series
- Tair: cache of metric period lengths
- Simple Log Service (SLS): data source
- Apache Flink: online decomposition based on BacktrackSTL and downstream tasks (anomaly detection, etc.)





EXPERIMENTS

- Comparable accuracy to SOTA offline RobustSTL
- Decompose a value within 1.6 us, 15x faster than SOTA online OneShotSTL



