SEAD: Unsupervised Ensemble of Streaming Anomaly Detectors

Saumya Gaurang Shah, Abishek Sankararaman,
Balakrishnan Murali Narayanaswamy, Vikramank Singh

Amazon Web Services, Santa Clara, CA
ICML 2025

“First model selection algorithm for streaming, unsupervised AD”

The Problem

The Problem

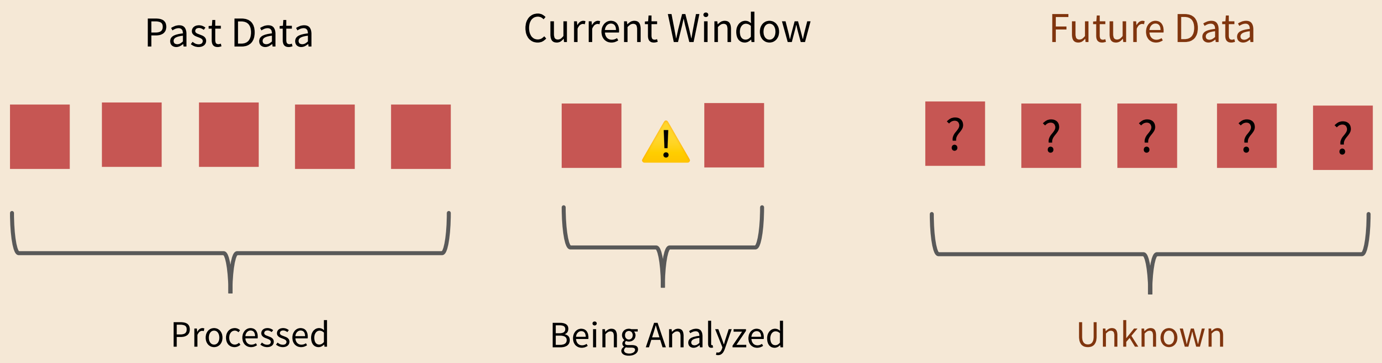

- Continuous stream of data points

The Problem

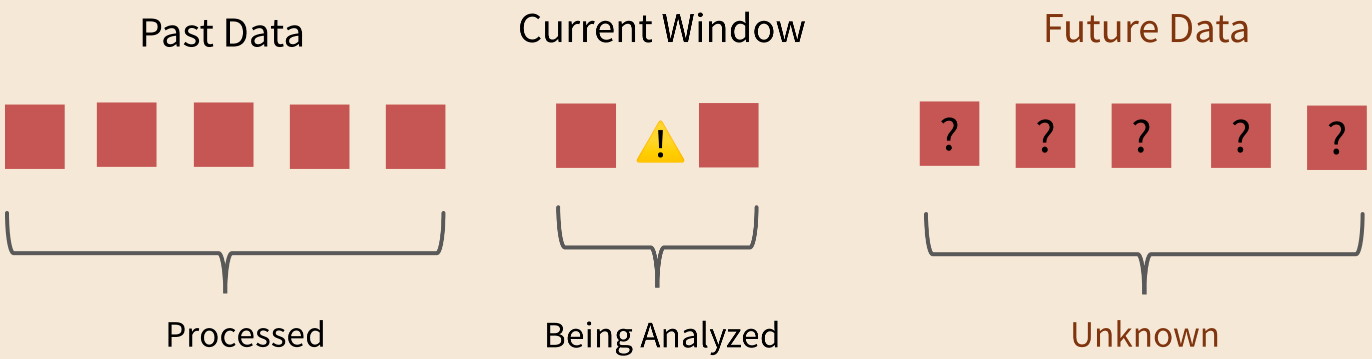

- Continuous stream of data points
- Real-time detections – must process each point in constant time

The Problem

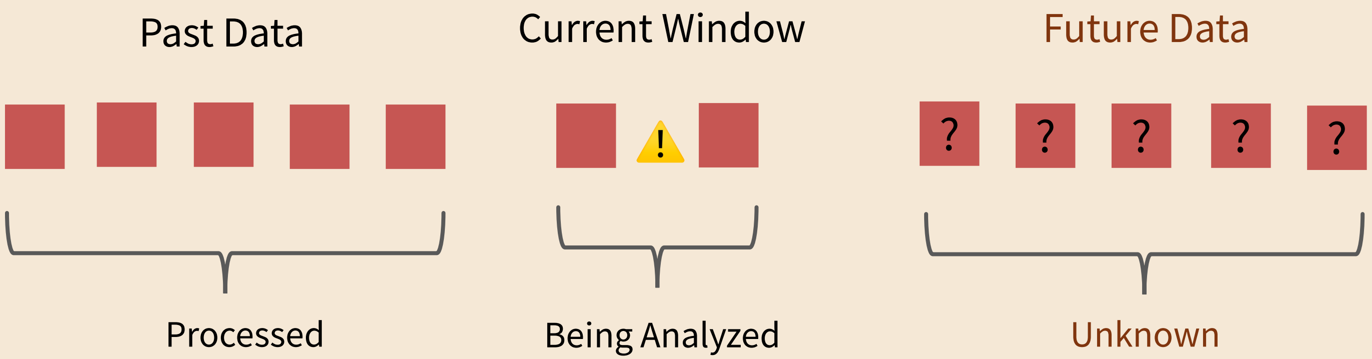

- Continuous stream of data points
- Real-time detections – must process each point in constant time
- Unsupervised - No feedback on whether predictions were correct

Challenges

Challenges

- No single algorithm works well across datasets

Challenges

- No single algorithm works well across datasets
- Toy example

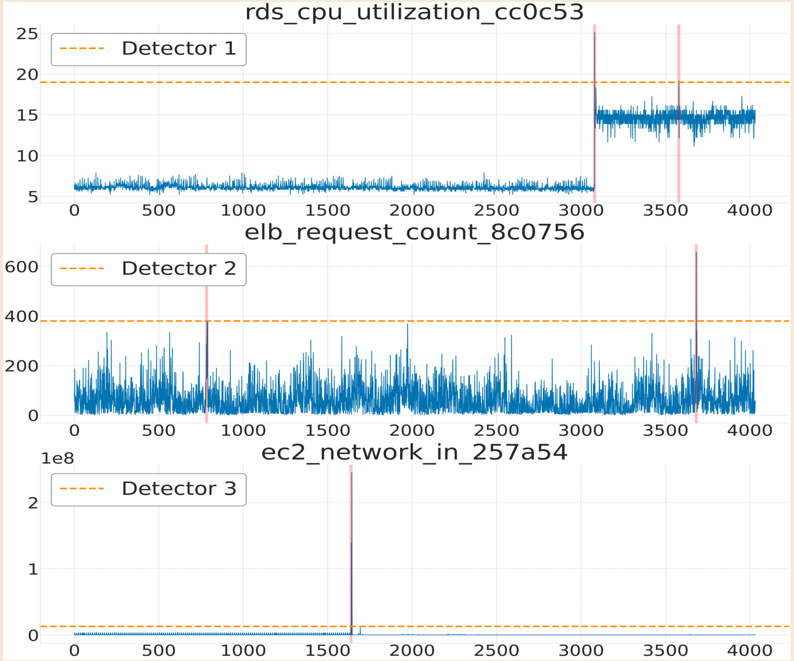

Accuracy

Dataset	Det 1	Det 2	Det 3
RDS	100%	0%	0%
ELB	0%	100%	0%
EC2	0%	0%	100%

Challenges

- No single algorithm works well across datasets
    - IForestASD (Ding & Fei, ‘13) performs the best on Pendigits dataset
    - RRCF (Guha et. al., ‘16) performs the best on Letter dataset
    - xStream (Manzoor et. al., ‘18) performs the best on INSECTS dataset
    - Rule based method (Shewhart, ‘31) performs the best on an internal telemetry dataset

Challenges

- No single algorithm works well across datasets
  - IForestASD (Ding & Fei, ‘13) performs the best on Pendigits dataset
  - RRCF (Guha et. al., ‘16) performs the best on Letter dataset
  - xStream (Manzoor et. al., ‘18) performs the best on INSECTS dataset
  - Rule based method (Shewhart, ‘31) performs the best on an internal telemetry dataset
- Data distributions change over time – must adapt to non stationarity over time

Key Insight

- Anomalies by definition are inherently rare

Key Insight

- Anomalies by definition are inherently rare
- Good anomaly detectors should output small scores most of the time

Key Insight

- Anomalies by definition are inherently rare
- Good anomaly detectors should output small scores most of the time
- Maintain weights for each detector – detectors with consistently lower scores have higher weight and vice versa

SEAD Architecture

SEAD Architecture

Streaming data points

SEAD Architecture

Streaming data points

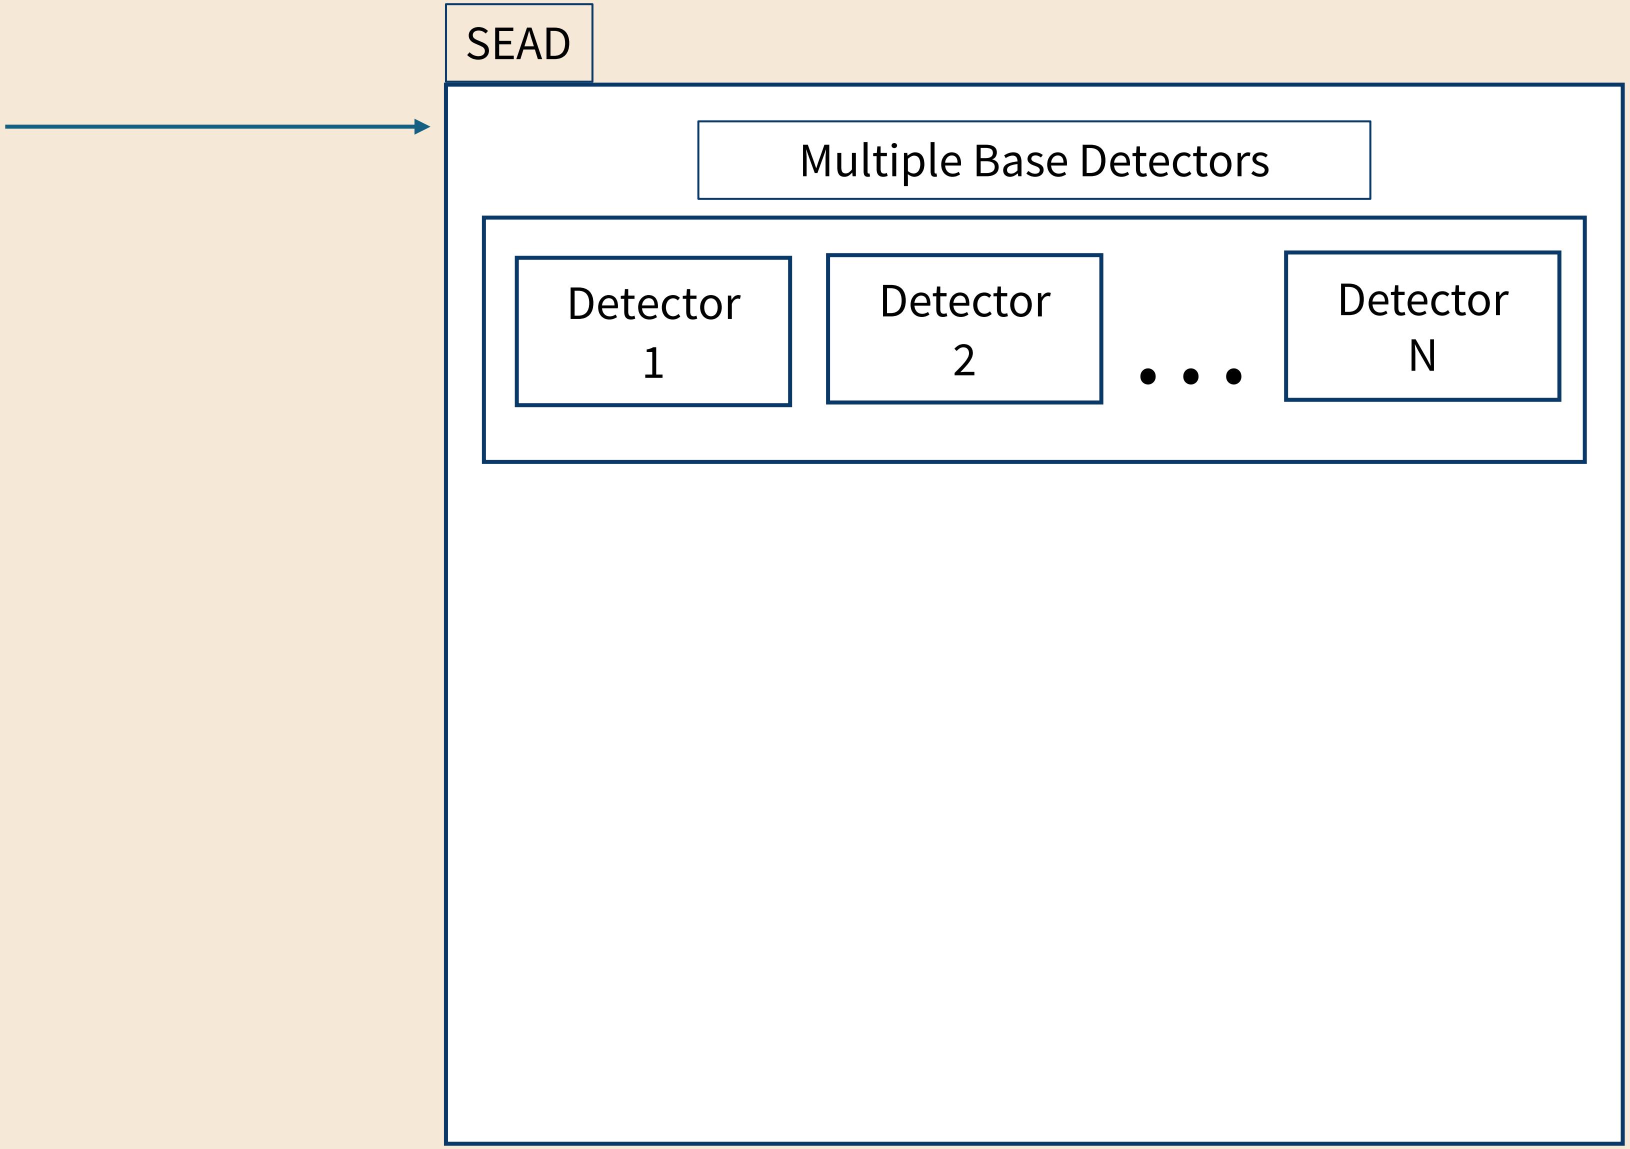

SEAD Architecture

Streaming data points

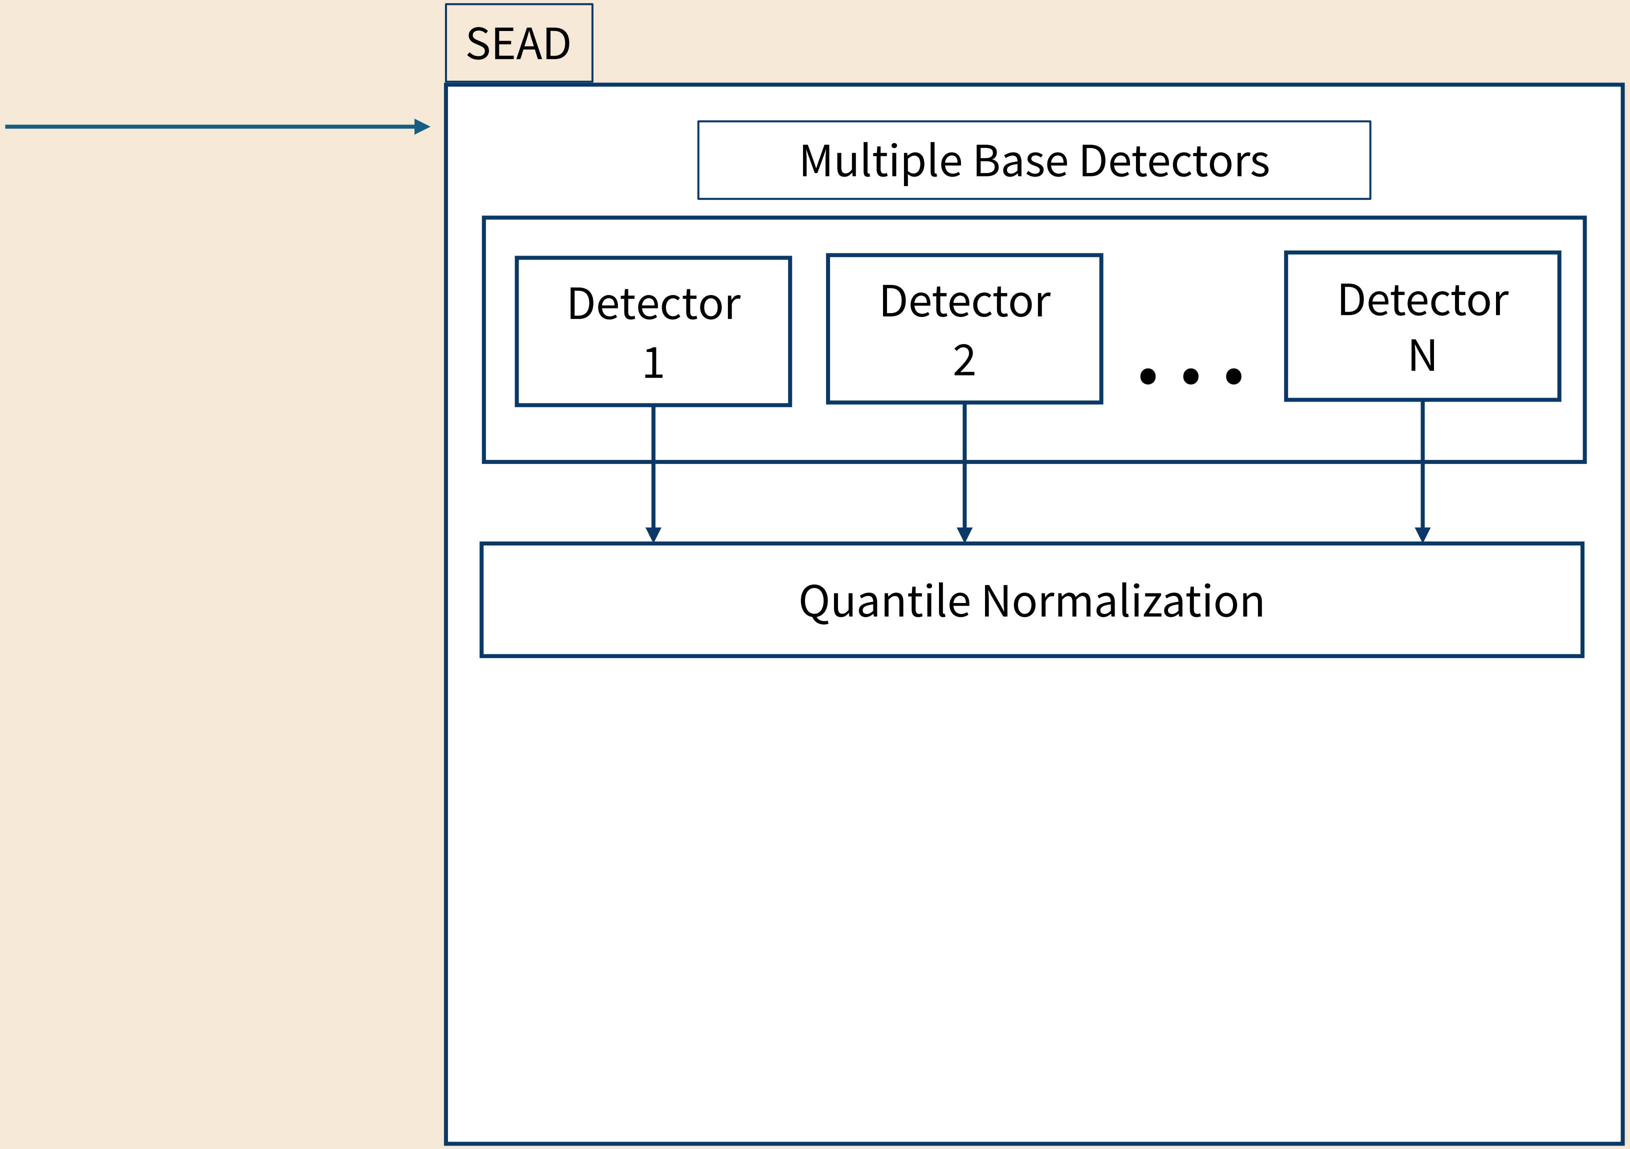

SEAD Architecture

Streaming data points

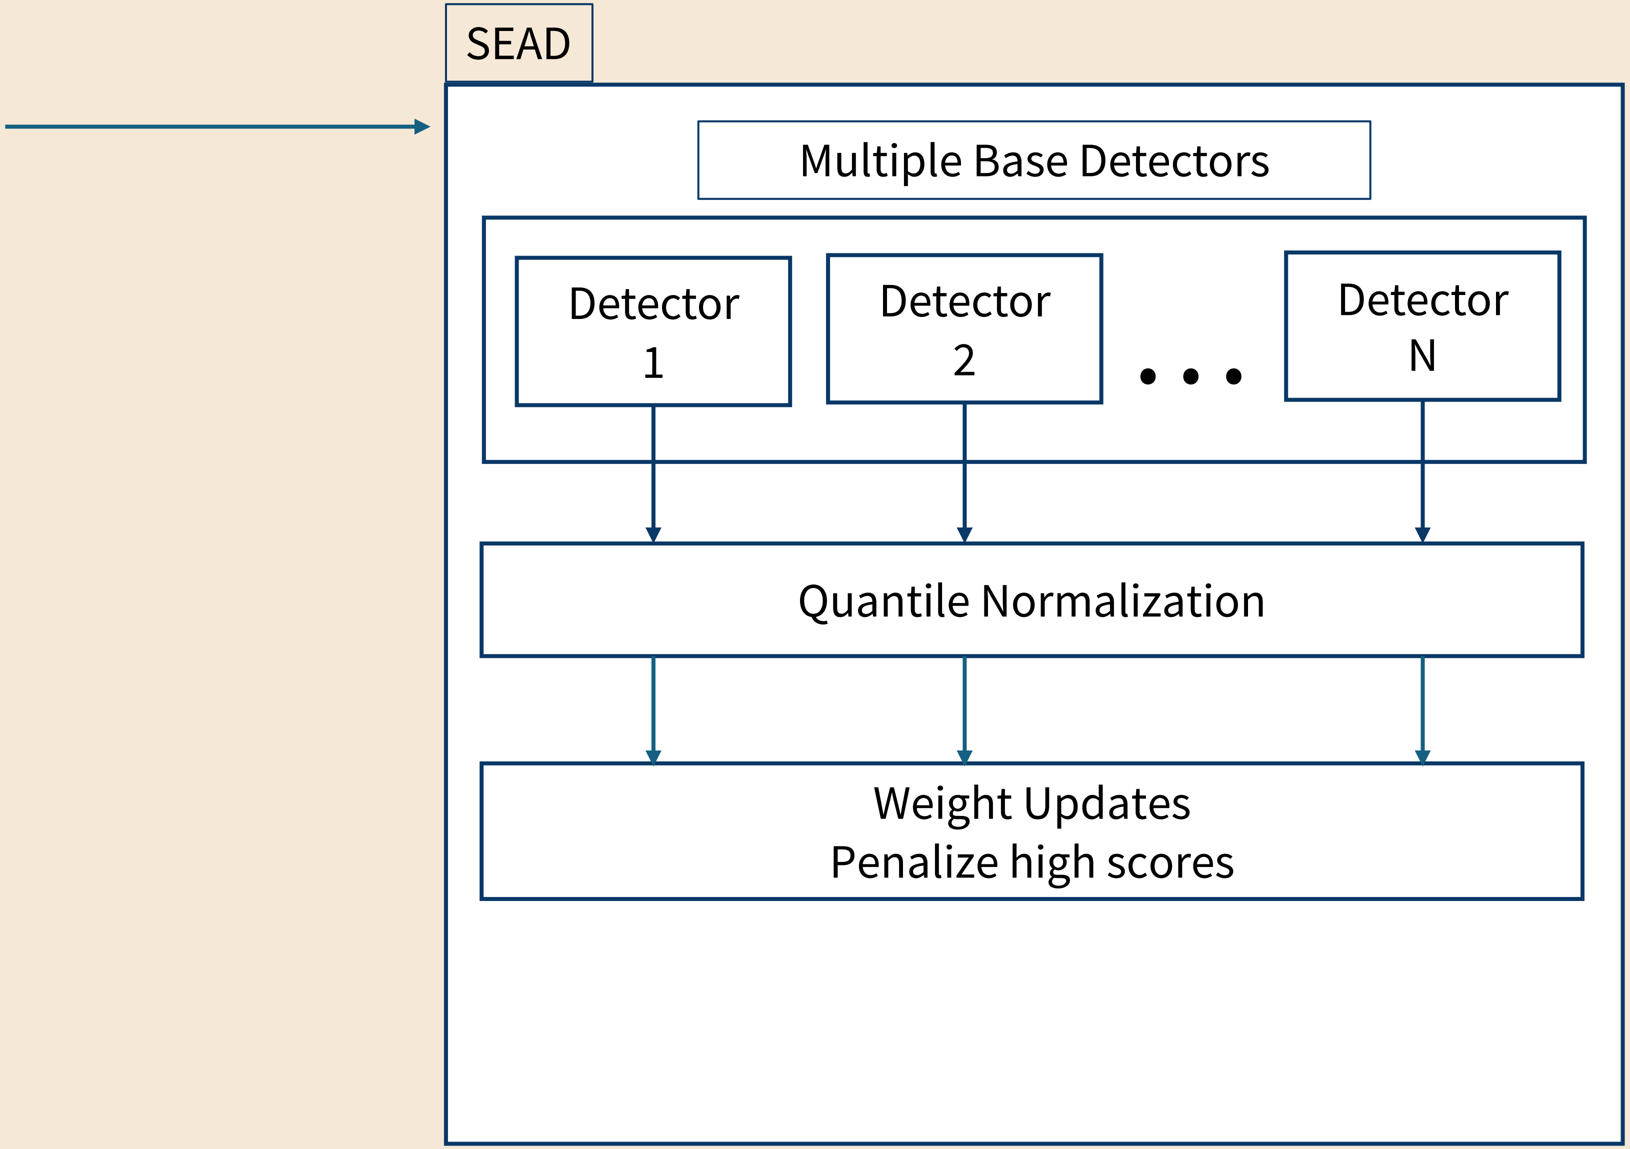

SEAD Architecture

Streaming data points

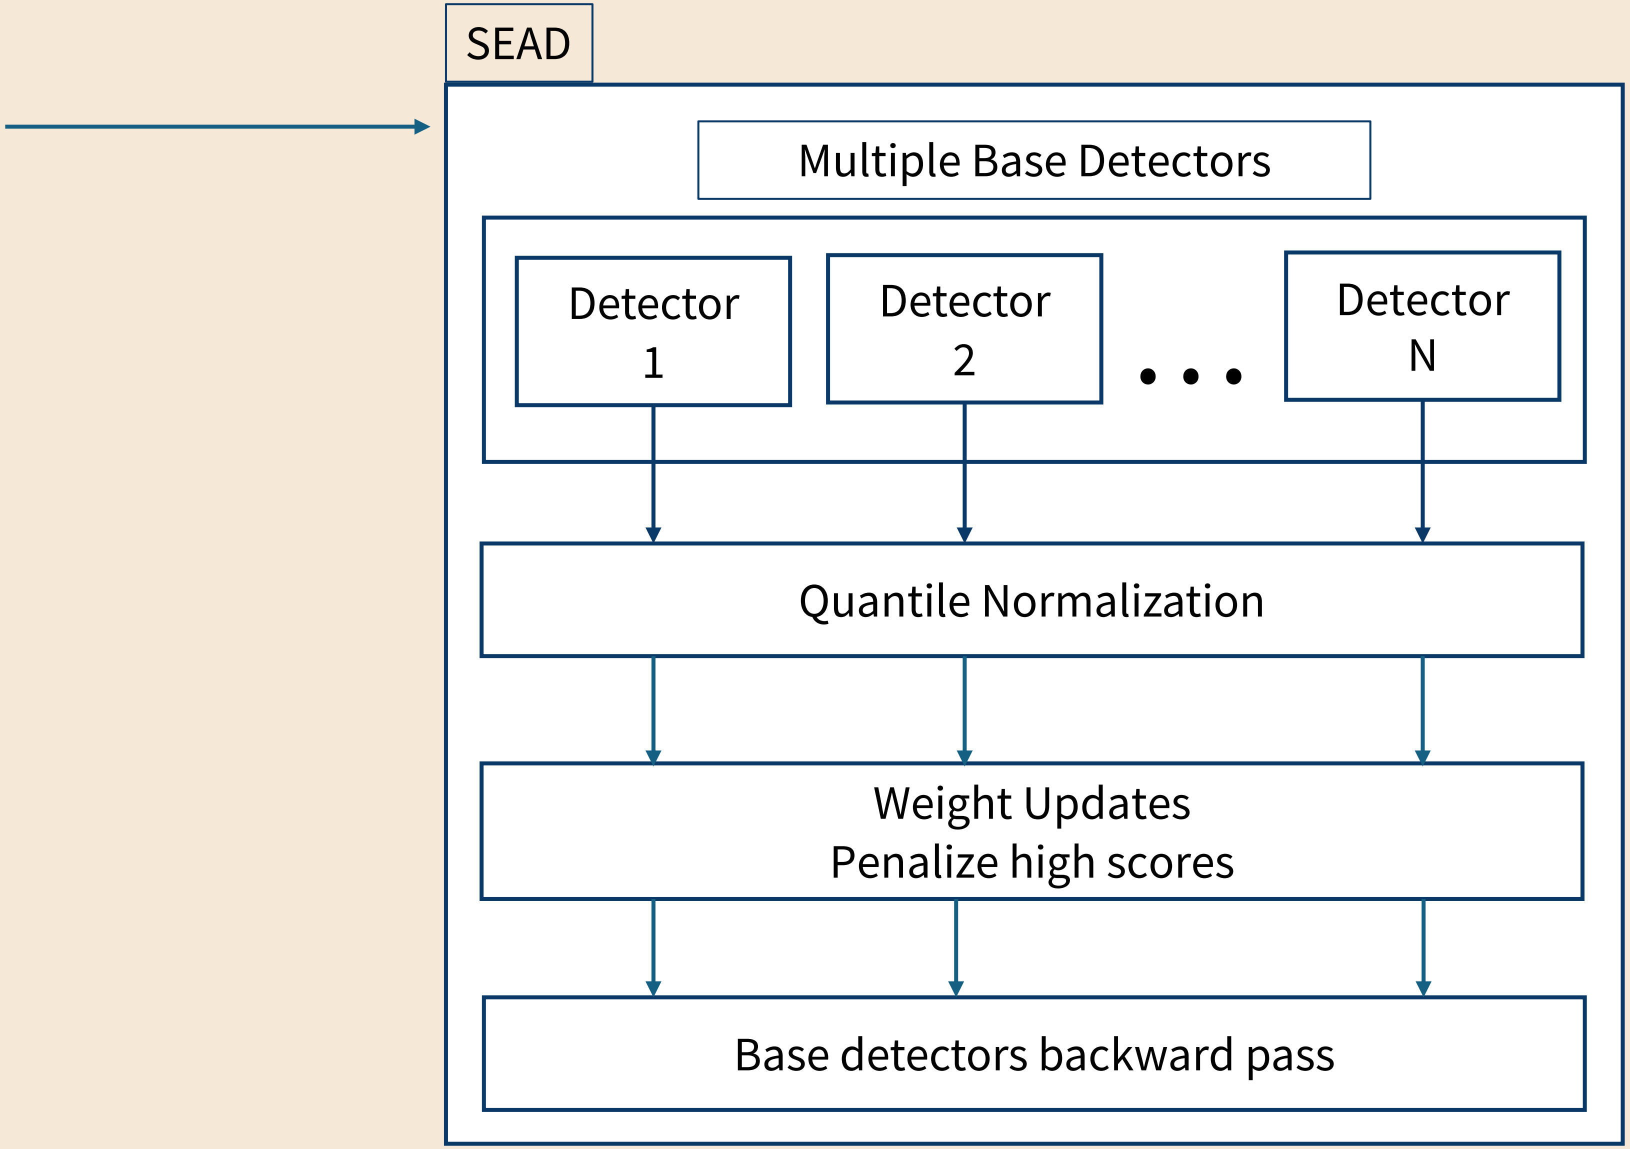

SEAD Architecture

Streaming data points

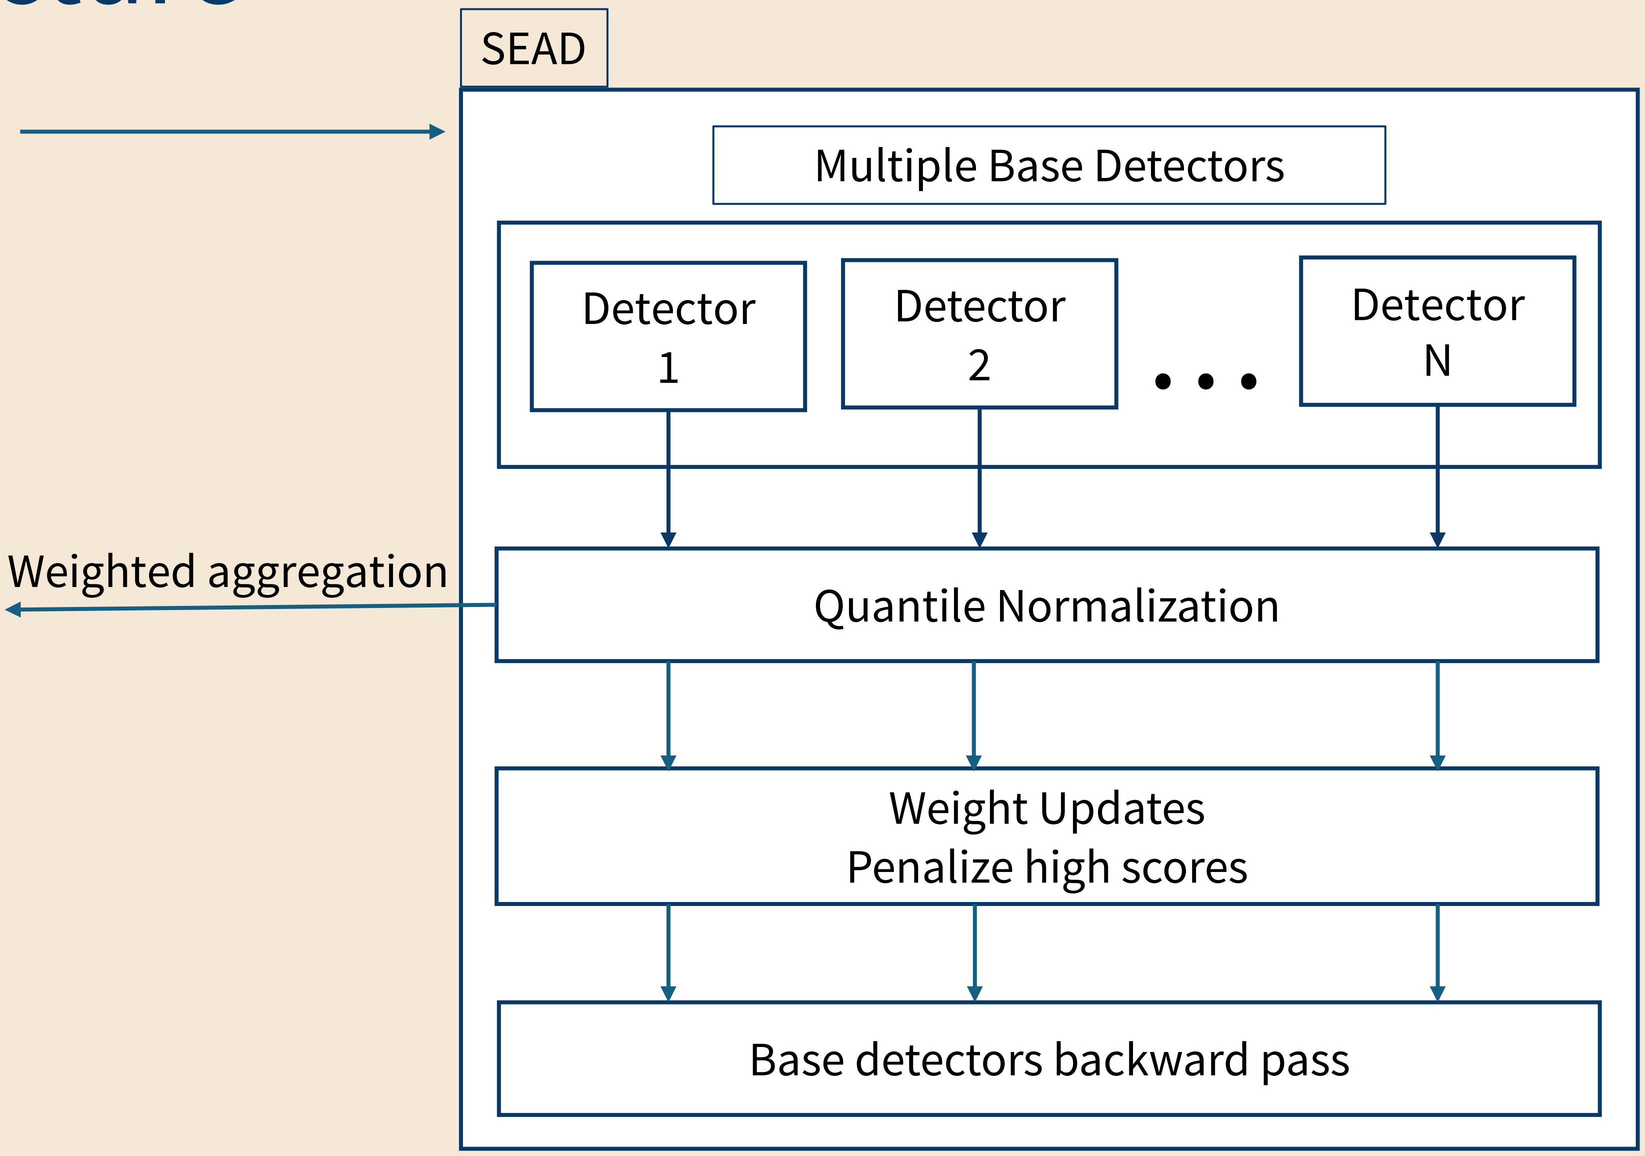

SEAD Architecture

Streaming data points

Final anomaly score

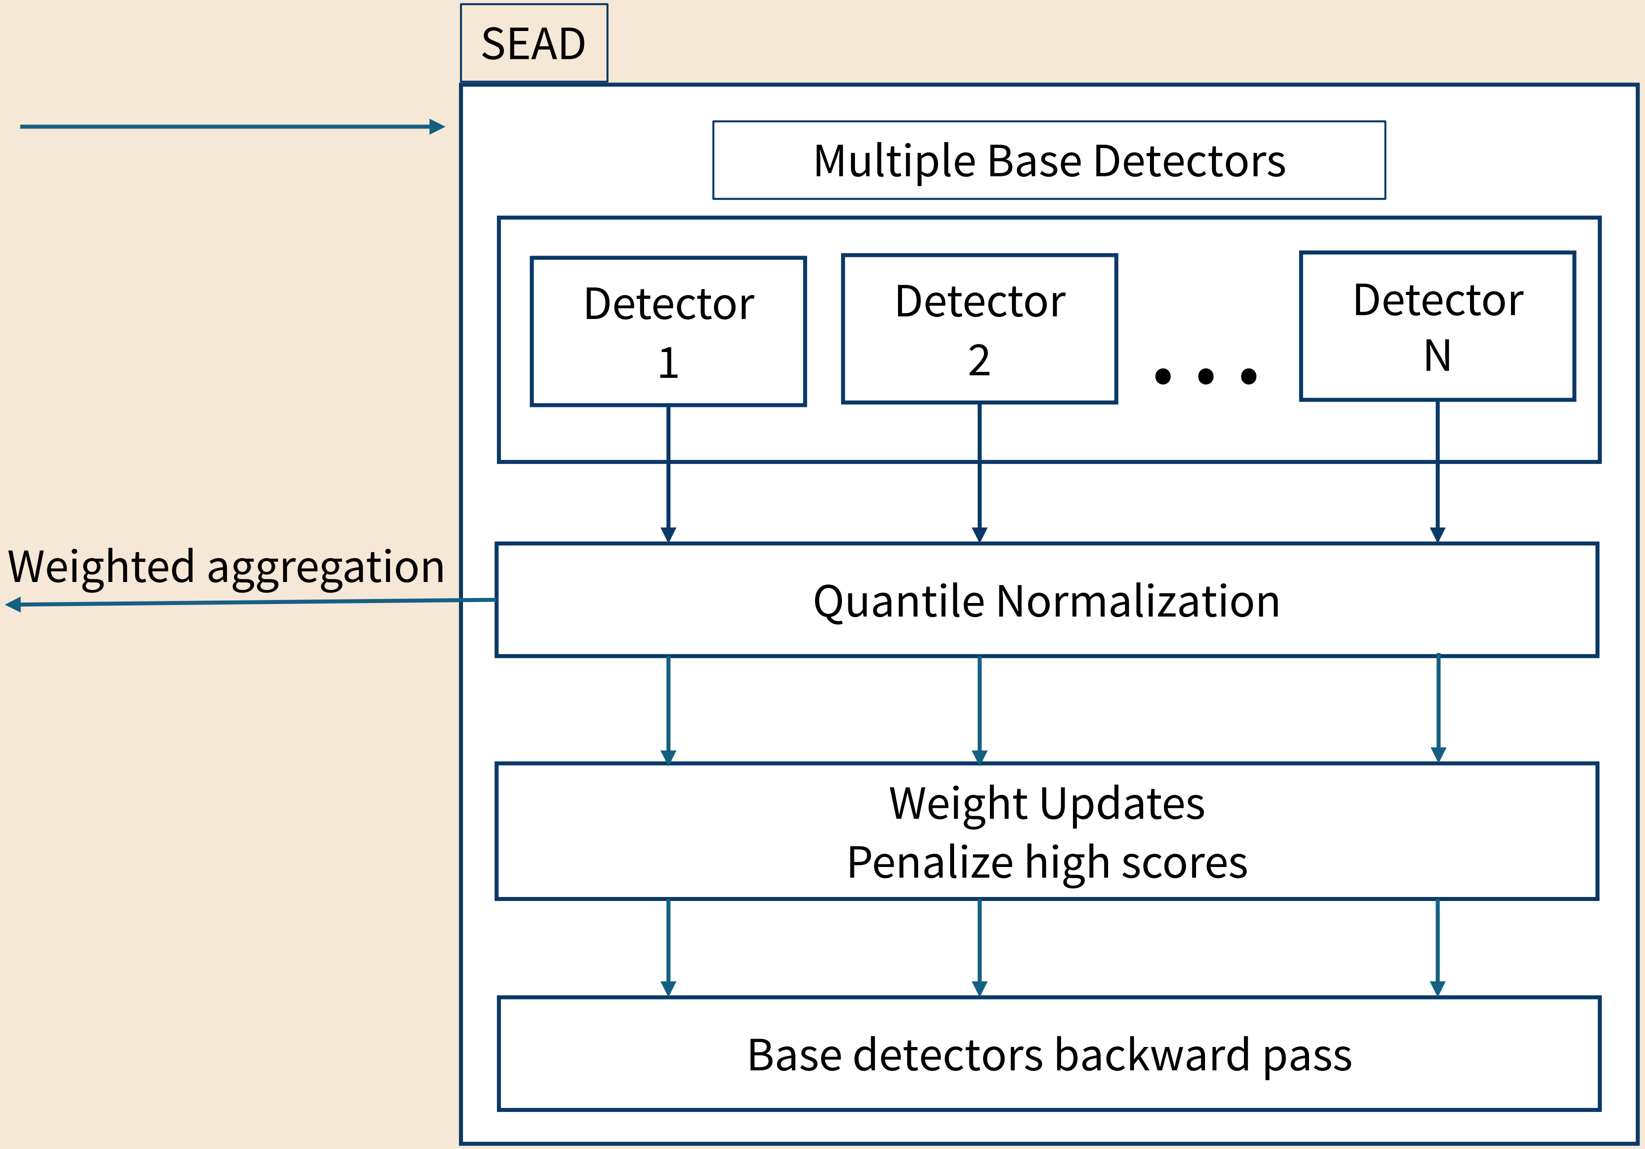

SEAD Architecture

Streaming data points

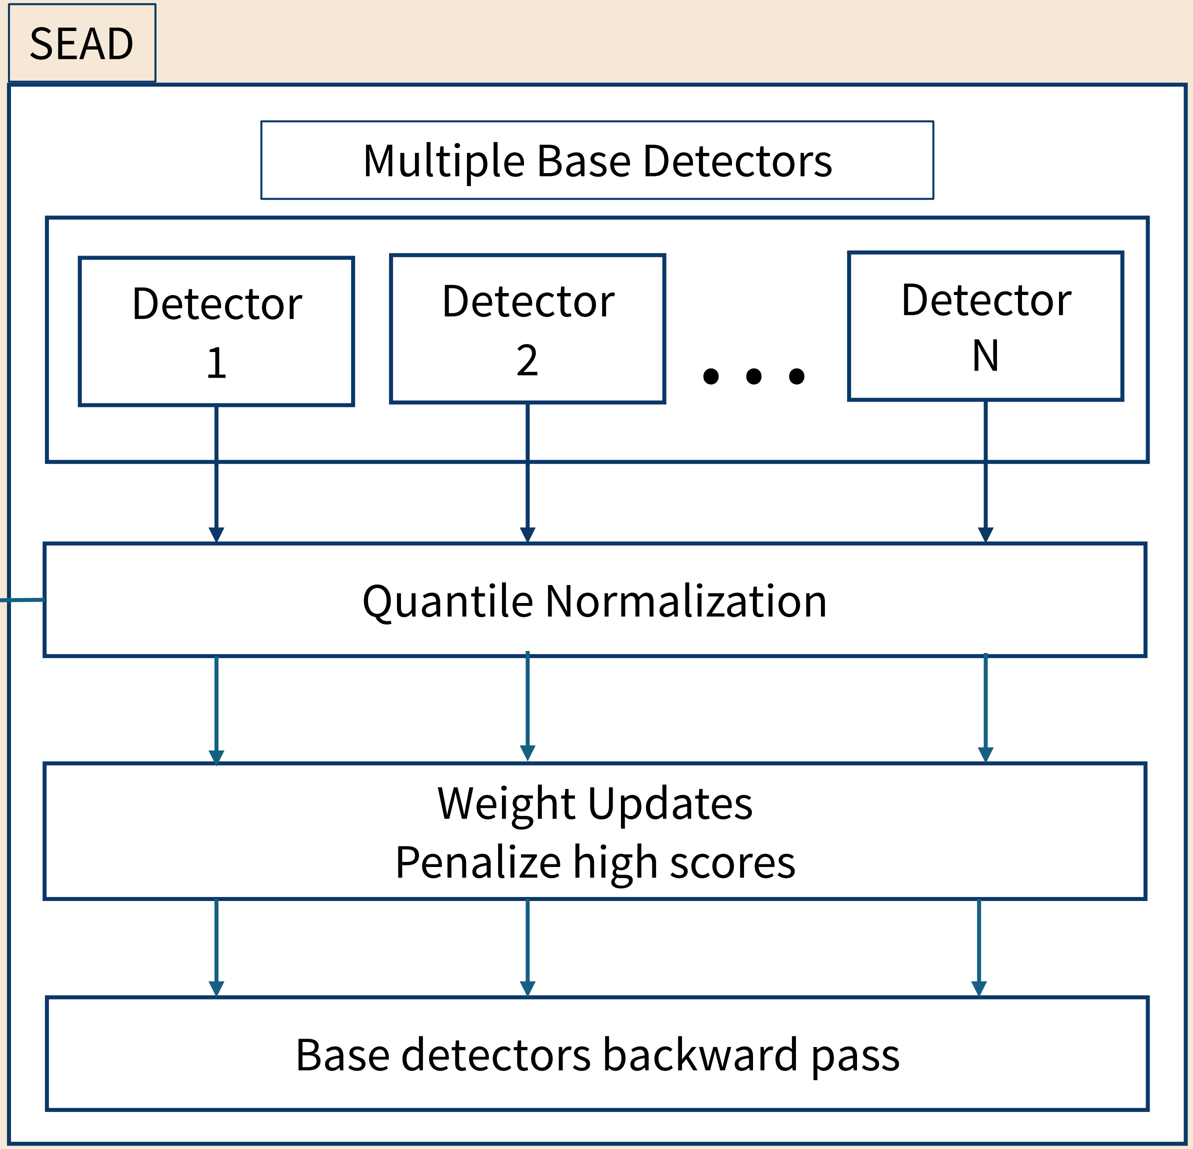

Final anomaly score

Weighted aggregation

- Unsupervised
- O(1) time and space
- Adaptive to distribution shifts
- Agnostic to choice of base detectors

Qualitative Example

Qualitative Example

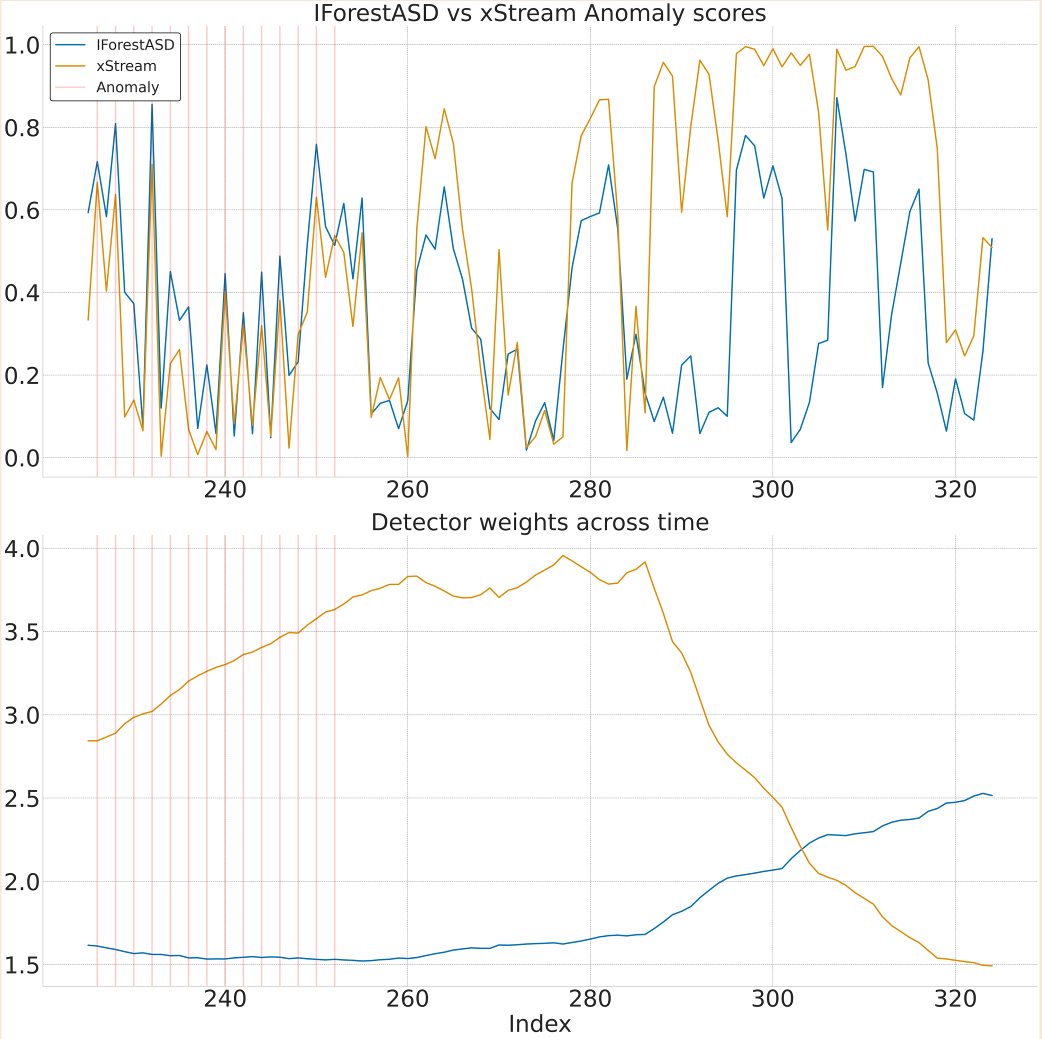

SEAD reassigns detector weights away from misfiring xStream detector

Experimental Evaluation

Experimental Evaluation

- 35x faster than offline methods with comparable detection performance

Experimental Evaluation

- 35x faster than offline methods with comparable detection performance
- Best average rank among all base methods and simple aggregators mean, max and min

Experimental Evaluation

- 35x faster than offline methods with comparable detection performance
- Best average rank among all base methods and simple aggregators mean, max and min
  - 13 base methods with different parameter configurations of IForestASD, xStream and RRCF and a single rule-based method
  - Comparison on 15 datasets including non-stationary INSECTS datasets

Experimental Evaluation

- 35x faster than offline methods with comparable detection performance
- Best average rank among all base methods and simple aggregators mean, max and min
  - 13 base methods with different parameter configurations of IForestASD, xStream and RRCF and a single rule-based method
  - Comparison on 15 datasets including non-stationary INSECTS datasets
- SEAD++ optimization has detection performance comparable to simple aggregators with ~2x speedup in runtime

Conclusions and Future Work

Conclusions and Future Work

- We propose SEAD – the first unsupervised online model selection algorithm for anomaly detection

Conclusions and Future Work

- We propose SEAD – the first unsupervised online model selection algorithm for anomaly detection
- Initializing SEAD weights using existing offline datasets is interesting future work

Conclusions and Future Work

- We propose SEAD – the first unsupervised online model selection algorithm for anomaly detection
- Initializing SEAD weights using existing offline datasets is interesting future work
- Future work can also investigate open regret guarantees on SEAD that holds under non-stationarity

Thank you!