Anomaly Detection on Streaming Data using Hierarchical Temporal Memory (and LSTM)

Jaime Coello de Portugal

Many thanks to Jochem Snuverink

"An observation which deviates so much from other observations as to arouse suspicions that it was generated by a different mechanism."

– Hawkins, 1980

- Try to find observation or sequences that deviate from the "normal behaviour".
- Experts would recognize these anomalous patterns easily, but cannot be monitoring the huge amount of data some systems produce.
- E.g: Credit card fraud detection, intrusion detection in cybersecurity, or **fault diagnosis in industry**.
- Specific e.g: At HIPA the MHB7R:ILOG:2 temperature detector broke down without anyone noticing.

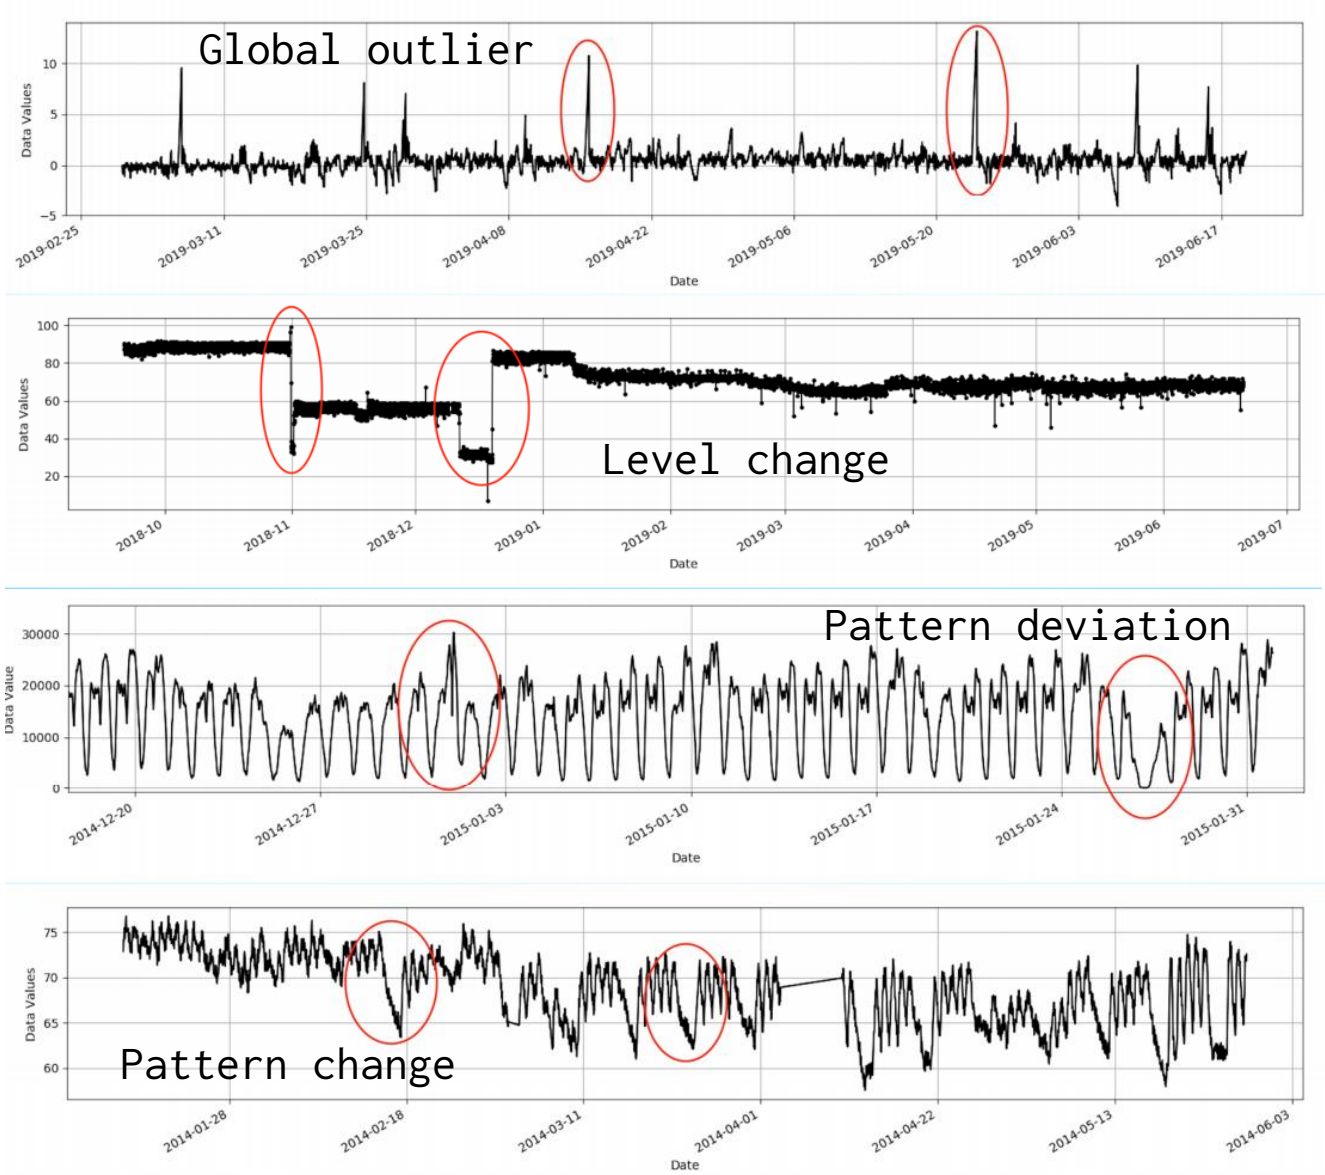

Plots from: Ted D'Ottavio et al, ***Experience Using NuPIC to Detect Anomalies in Controls Data****, ICALEPCS 2019*

- Model of the organization of “pyramidal neurons” in the neocortex of mammals.
- Tries to explain how neuronal structure remember sequences.

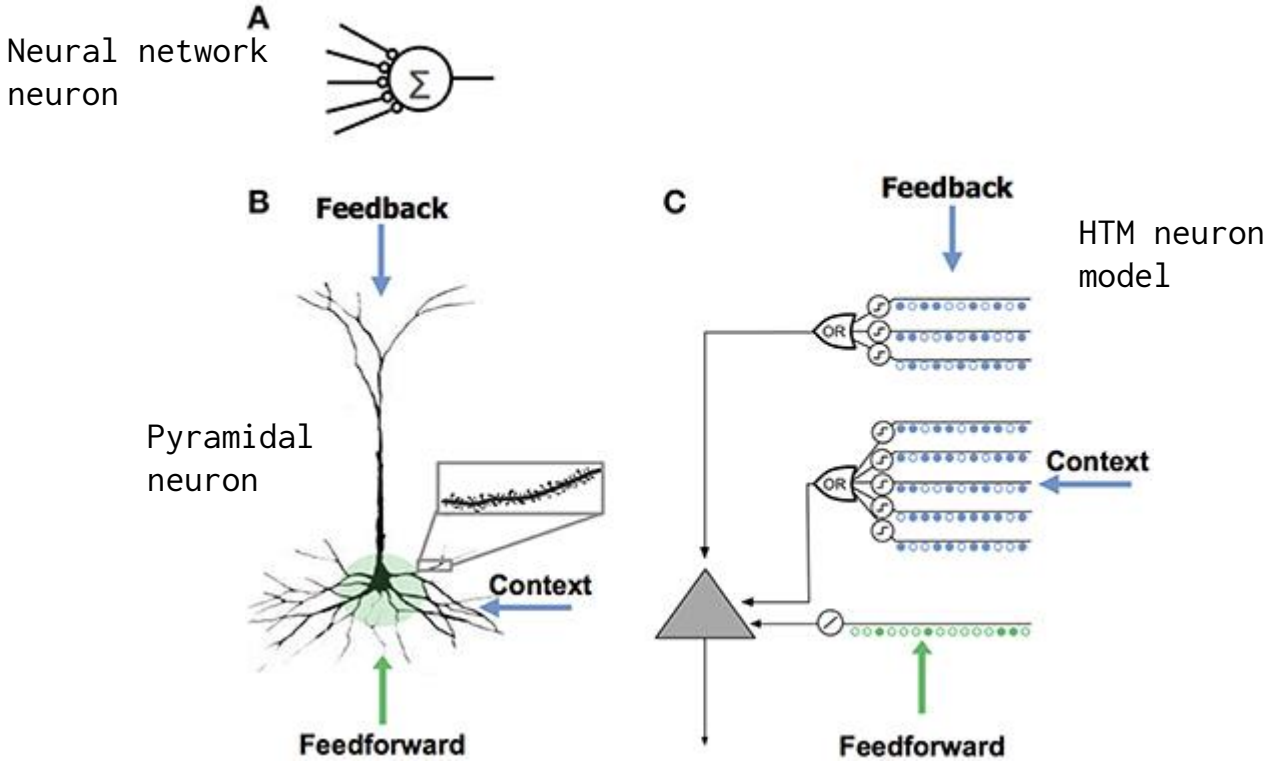

Pyramidal neurons connect forming “columns” that share input and output.

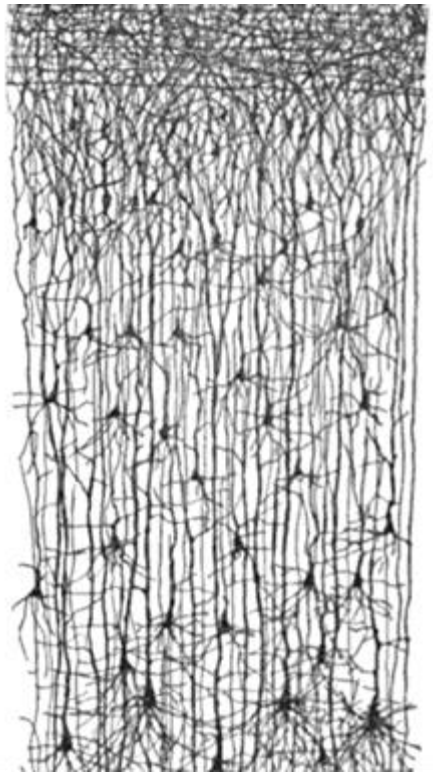

J. Hawkins and S. Ahmad, ***Why Neurons Have Thousands of Synapses, a Theory of Sequence Memory in Neocortex***
https://doi.org/10.3389/fncir.2016.00023

Hierarchical Temporal Memory: Sparse Distributed Representations (SDR)

- Scalar variables:

1:	1	1	1	0	0	0	0	0	0	0	0	0	0	0
2:	0	1	1	1	0	0	0	0	0	0	0	0	0	0
3:	0	0	1	1	1	0	0	0	0	0	0	0	0	0

……

- Similar values overlap ->
  - Encodes “proximity” between values.

- Produces some resistance to noise.
  Discrepancies due to noise will get the same or close representation.

- Categorical variables:

	1	1	1	0	0	0	0	0	0	0	0	0	0	0

0	0	0	1	1	1	0	0	0	0	0	0	0	0	0

0	0	0	0	0	0	1	1	1	0	0	0	0	0	

No overlap between different categories

- Several variables can share an SDR.
- Can also encode periodic proximity.
- E.g. power consumption in a gym:

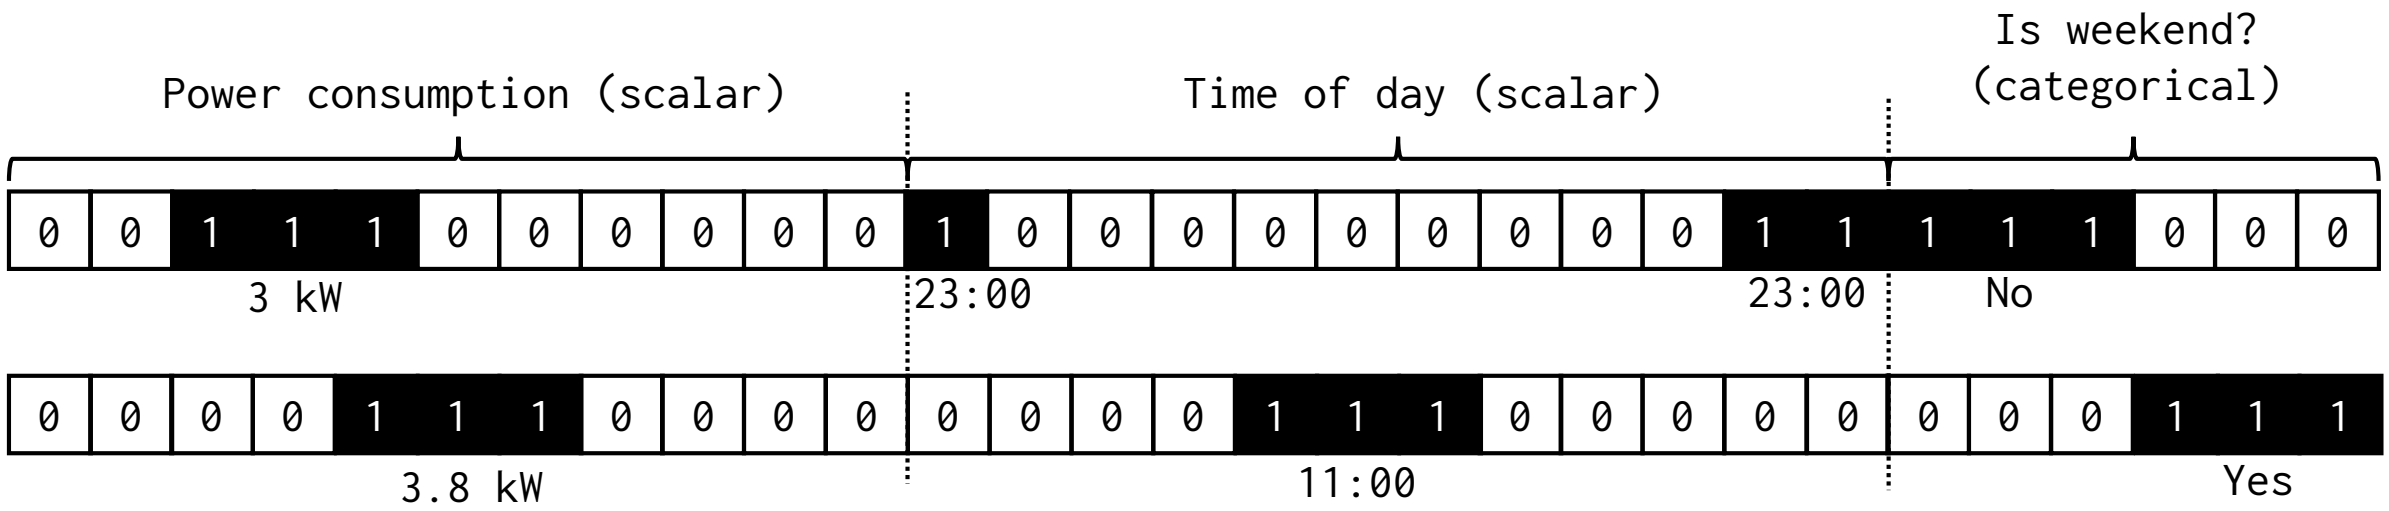

- The input to the HTM network is one of these SDR per time unit.

- Each cell on the SP gets random weighted connections to the input space.
- If the weight is over some threshold the cell gets connected to a bit of the input space.
- The N cells that have the most connection overlap to the input space become “active”.
- Target:

Maintain the semantic information from the input space with a fixed (~2%) sparsity.

- **Active** cells **learn** by updating their weights.

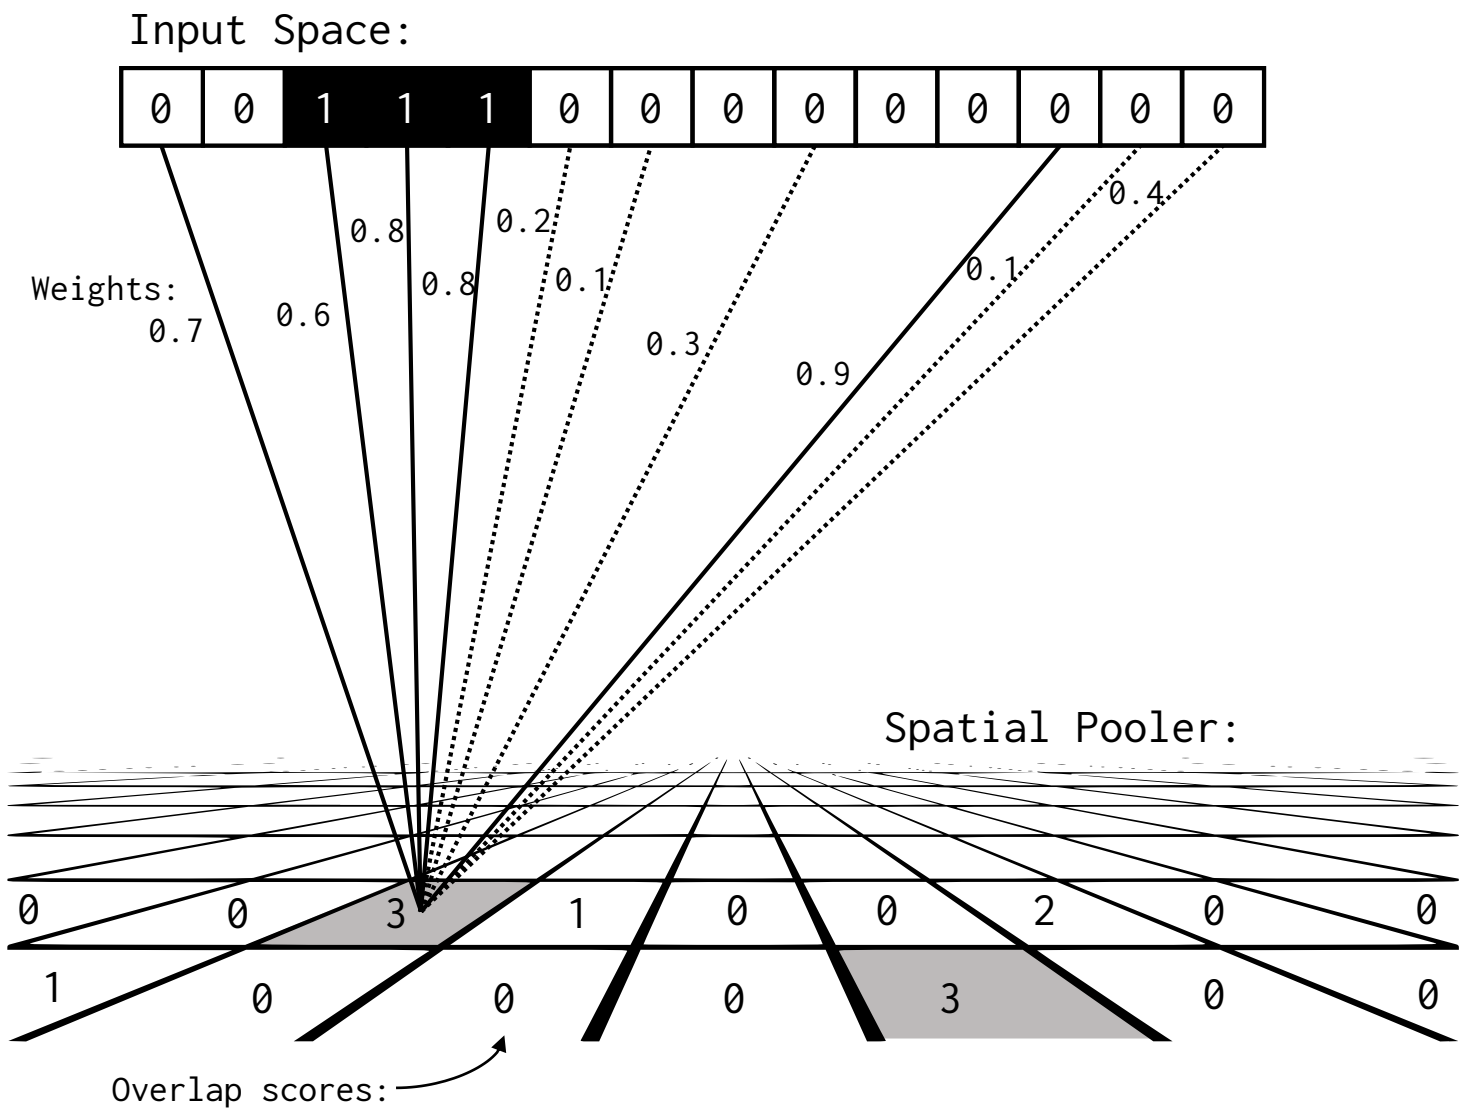

- Each cell of the SP is actually made of several cells that respond to the same input, forming “mini-columns”.
- These cells have connections with cells in other columns.
- If enough of these connections are active, the cell goes into “predictive” state.

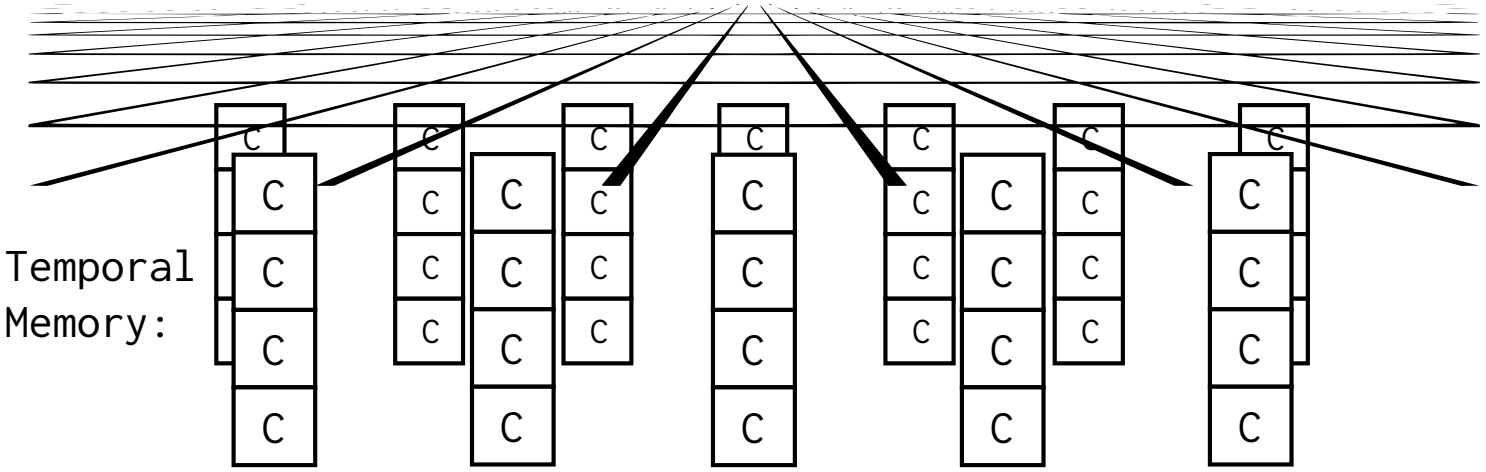

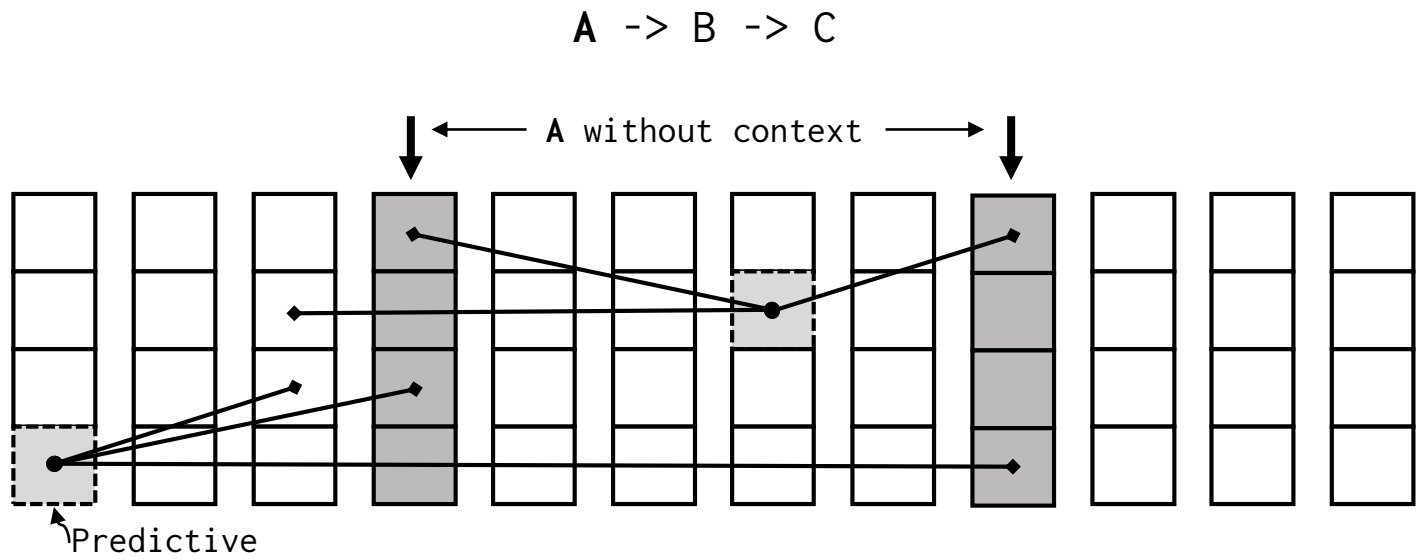

- On the next input, if a column containing a predictive cell becomes active only the predictive cell becomes active.

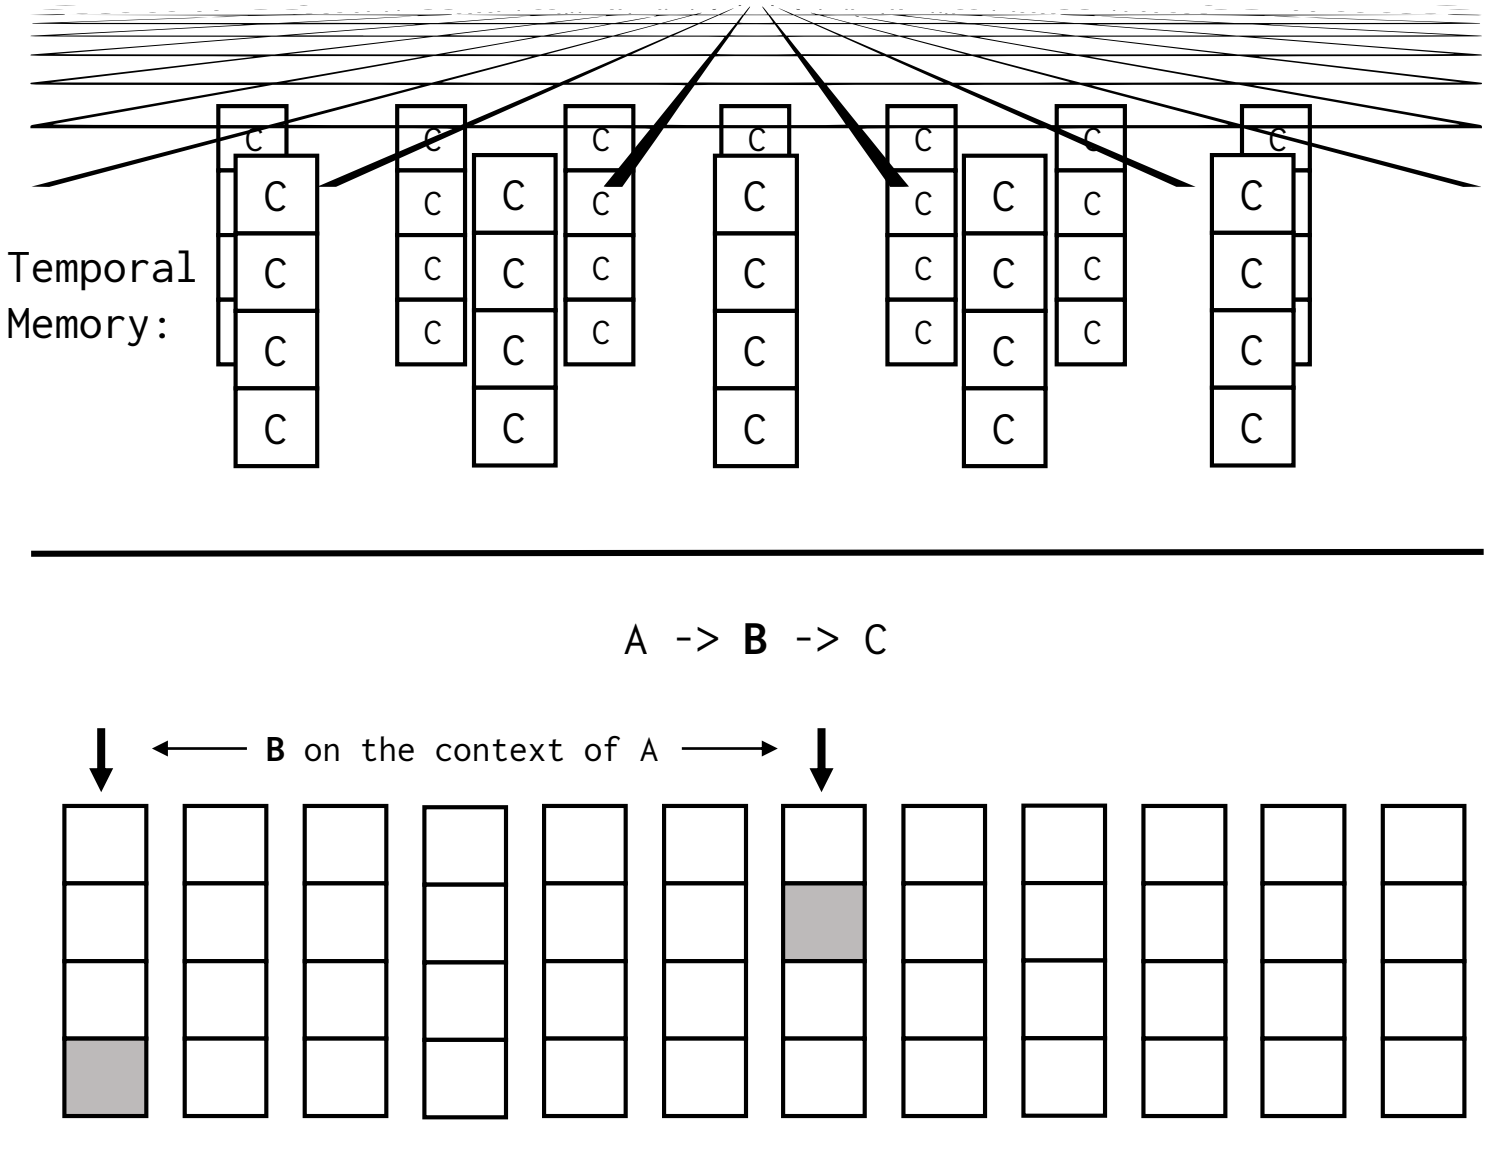

- When an unexpected input happens, the whole column activates.
- One of the cells of the column is then chosen as the “winner” and grows connections to the previous state.
- Old unused connections are slowly forgotten.

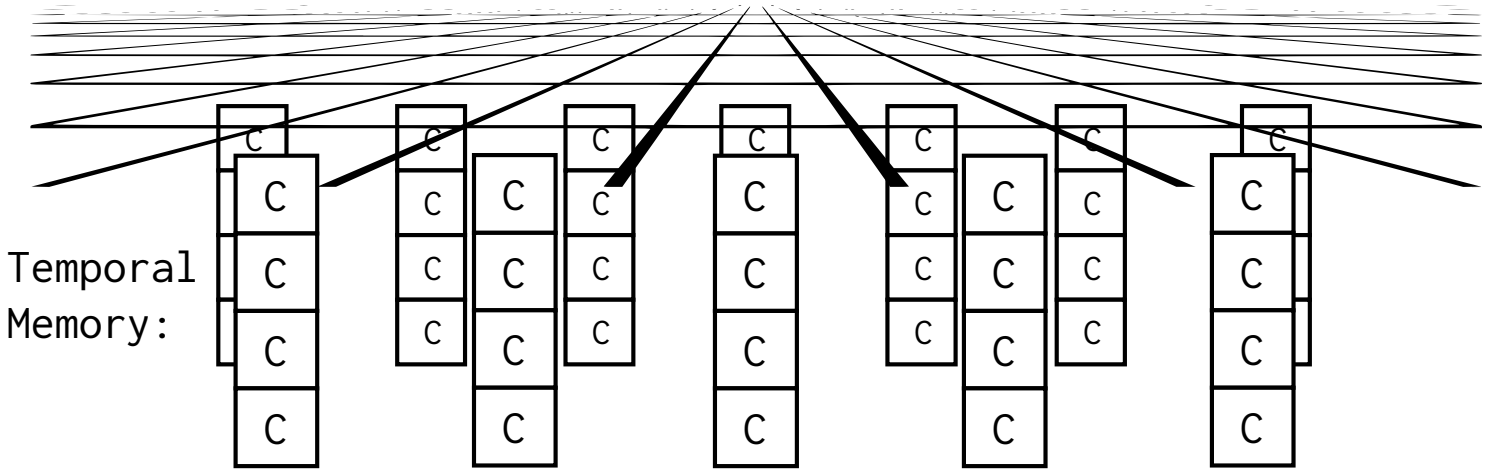

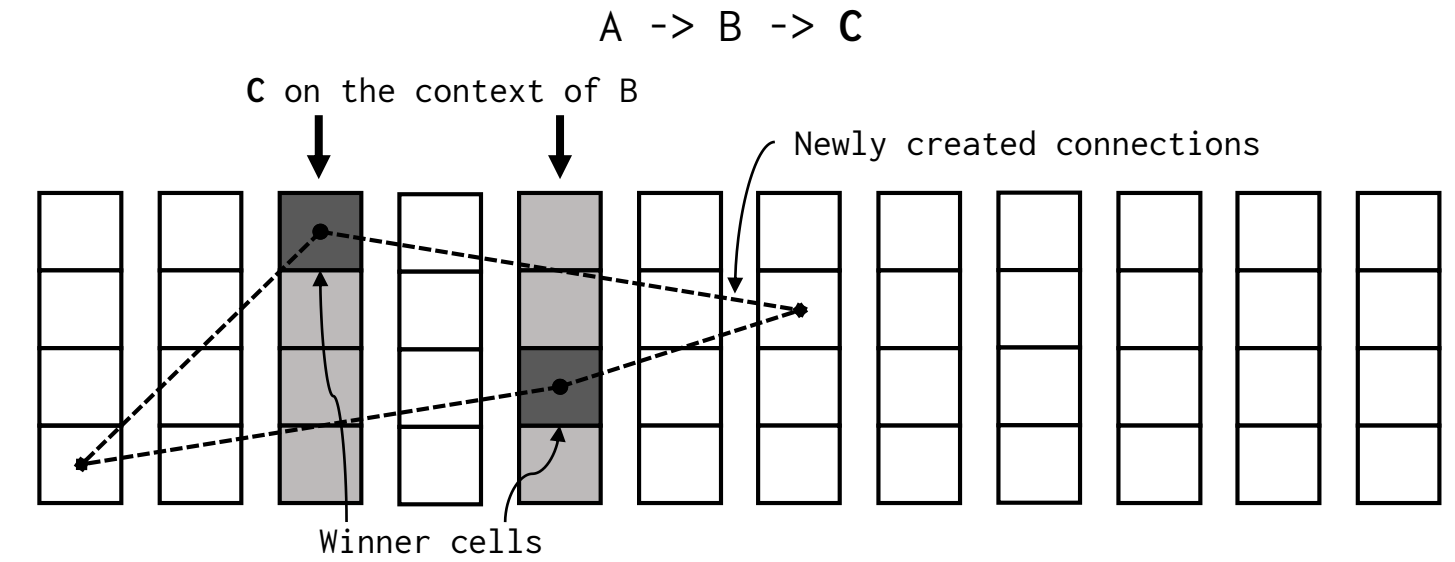

HTM Anomaly Score:

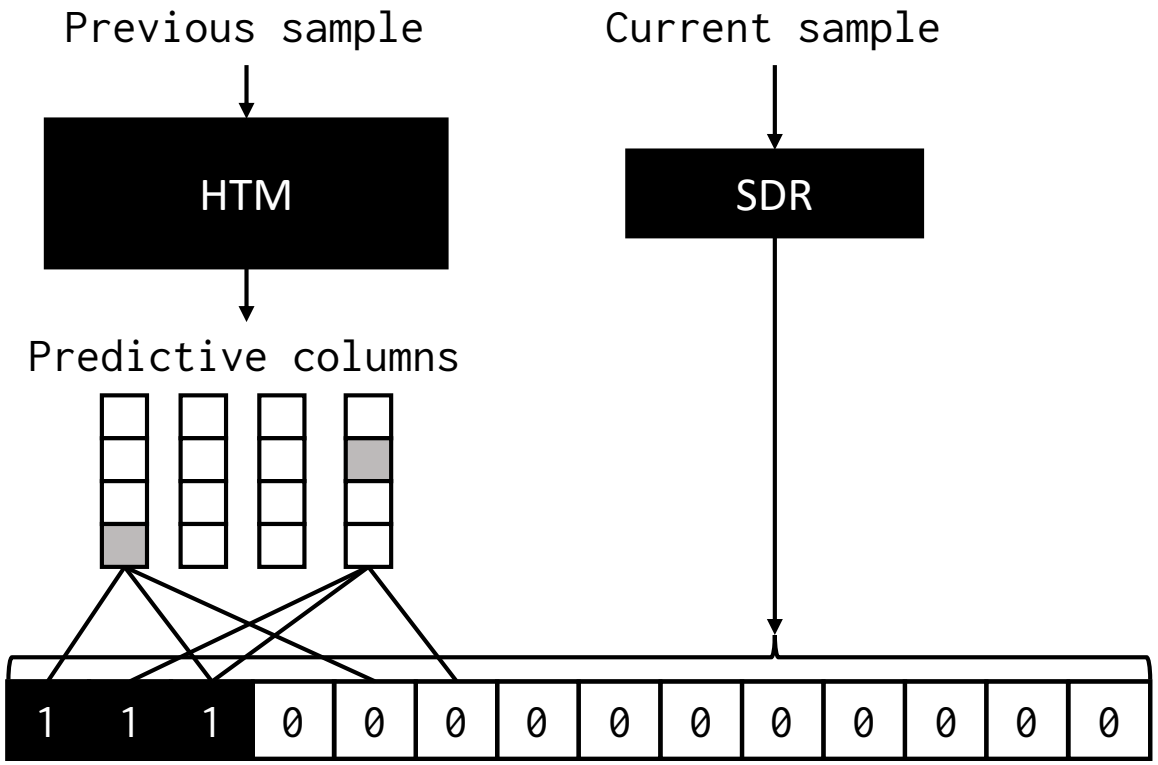

$$\text{Anomaly Score} = \frac{\text{Correctly predicted inputs}}{\text{Number of connections}}$$

LSTM Anomaly Score:

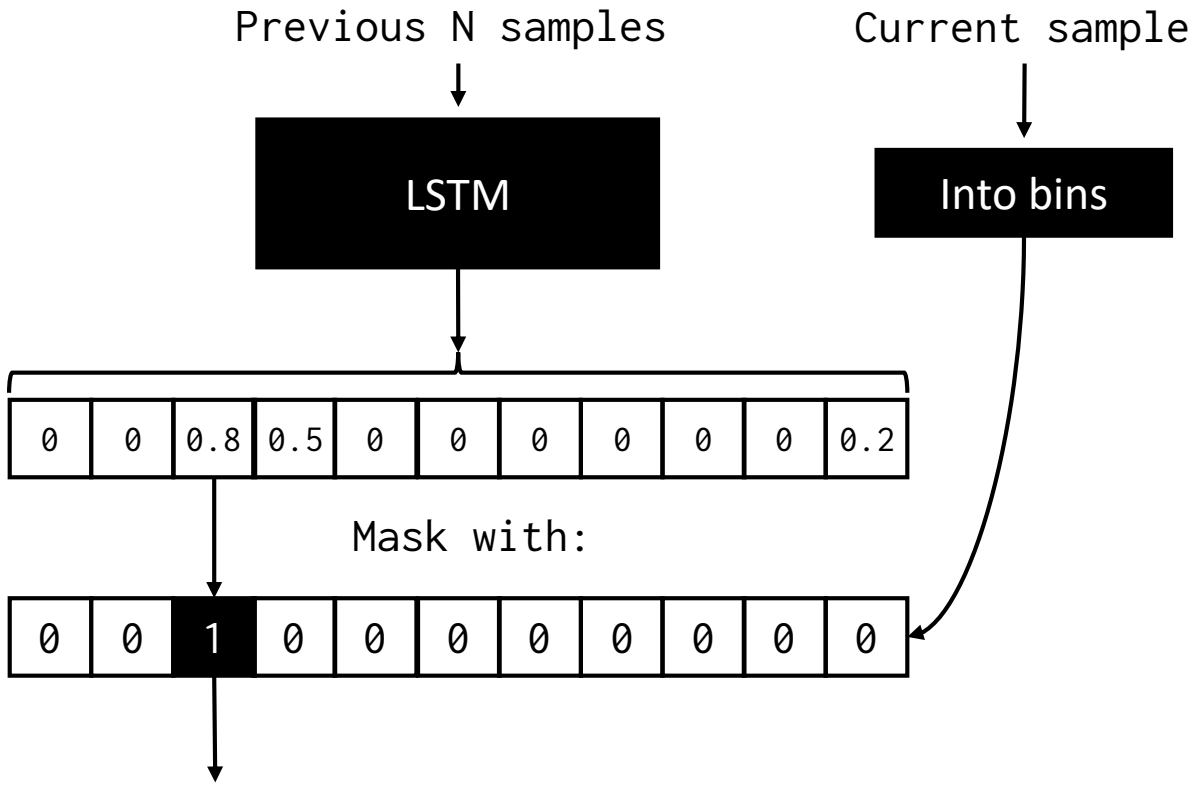

Training penalty = Anomaly Score + $\gamma \sum |prediction|$

Anomaly detection implementation

NUPIC is a Python 2 (wrapping C++) implementation of the HTM algorithm by Numenta (numenta.com).

They provide a set of hyperparameters that works good for anomaly detection out of the box

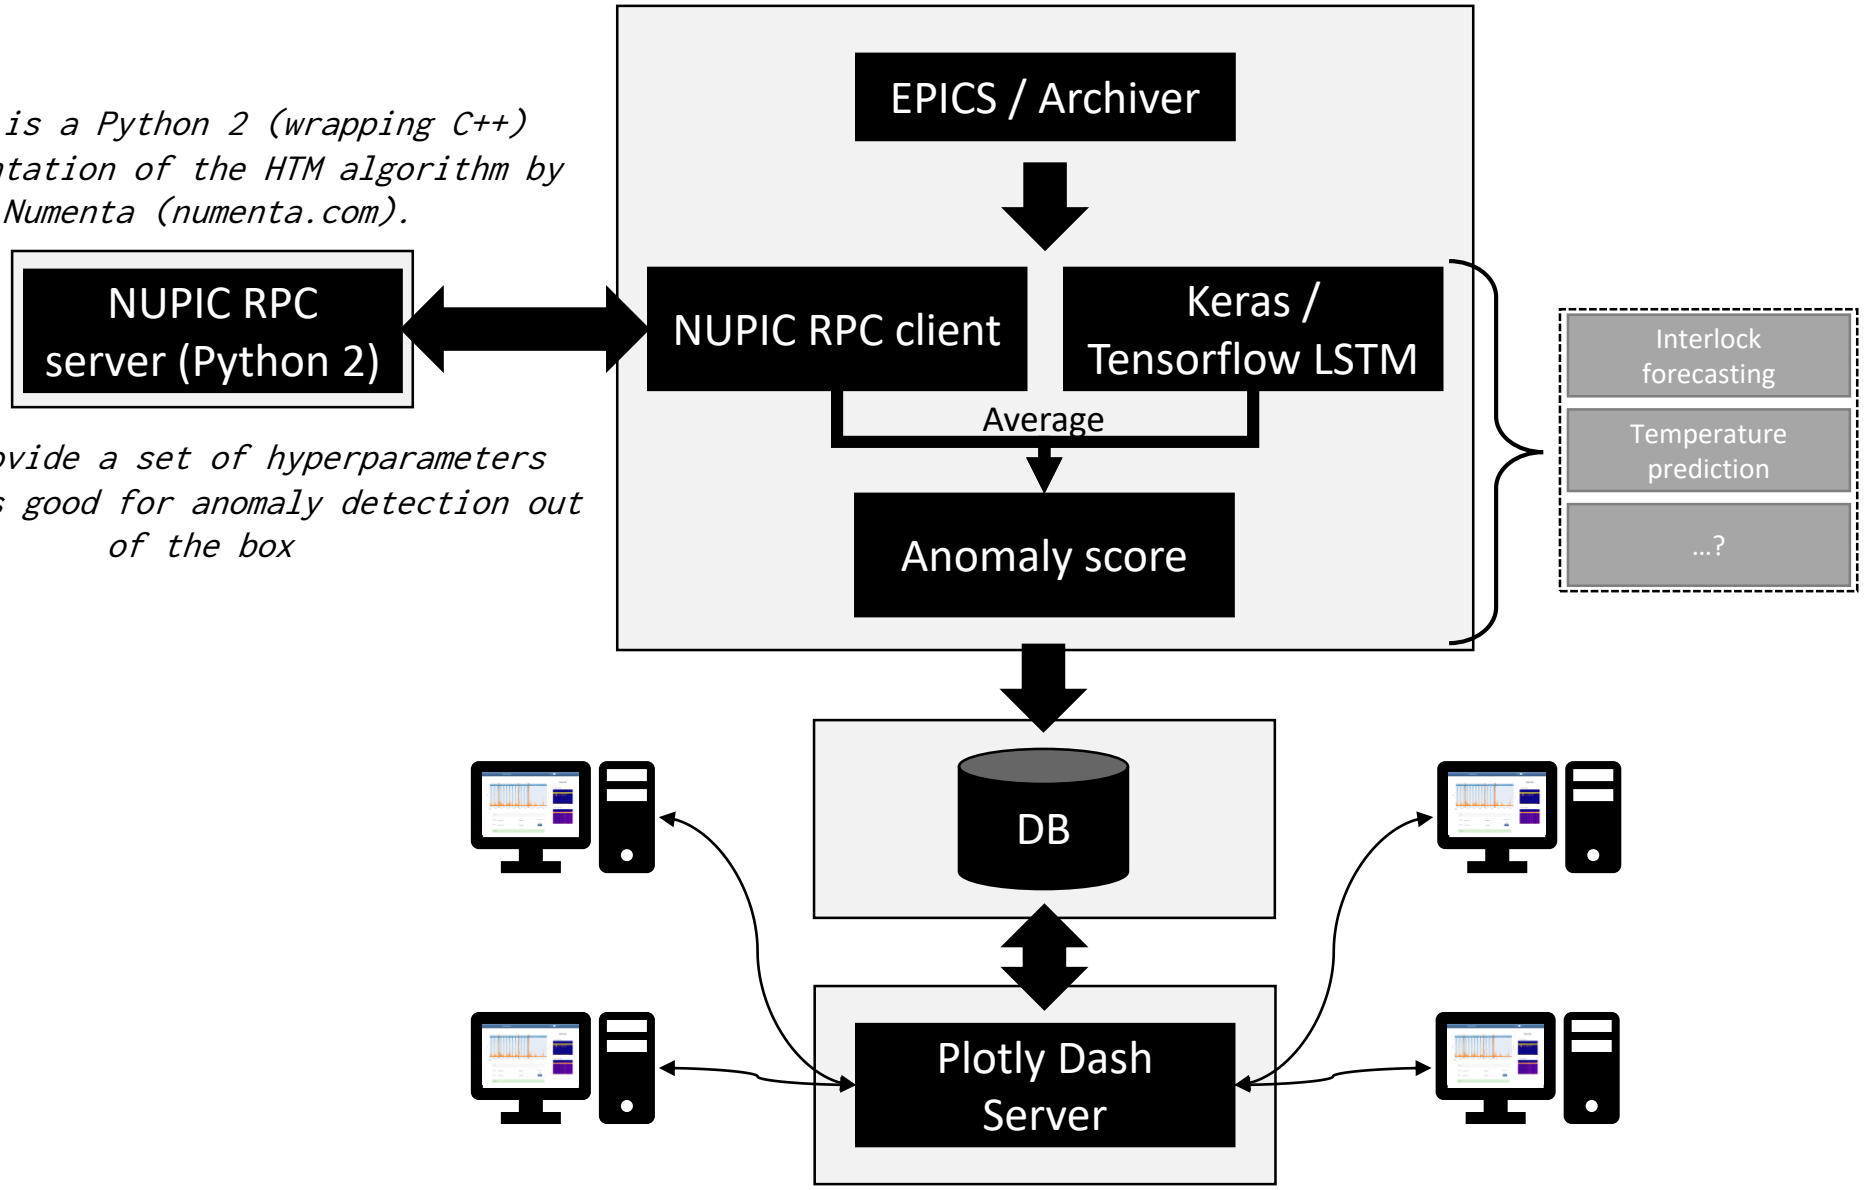

Anomaly detection in HIPA

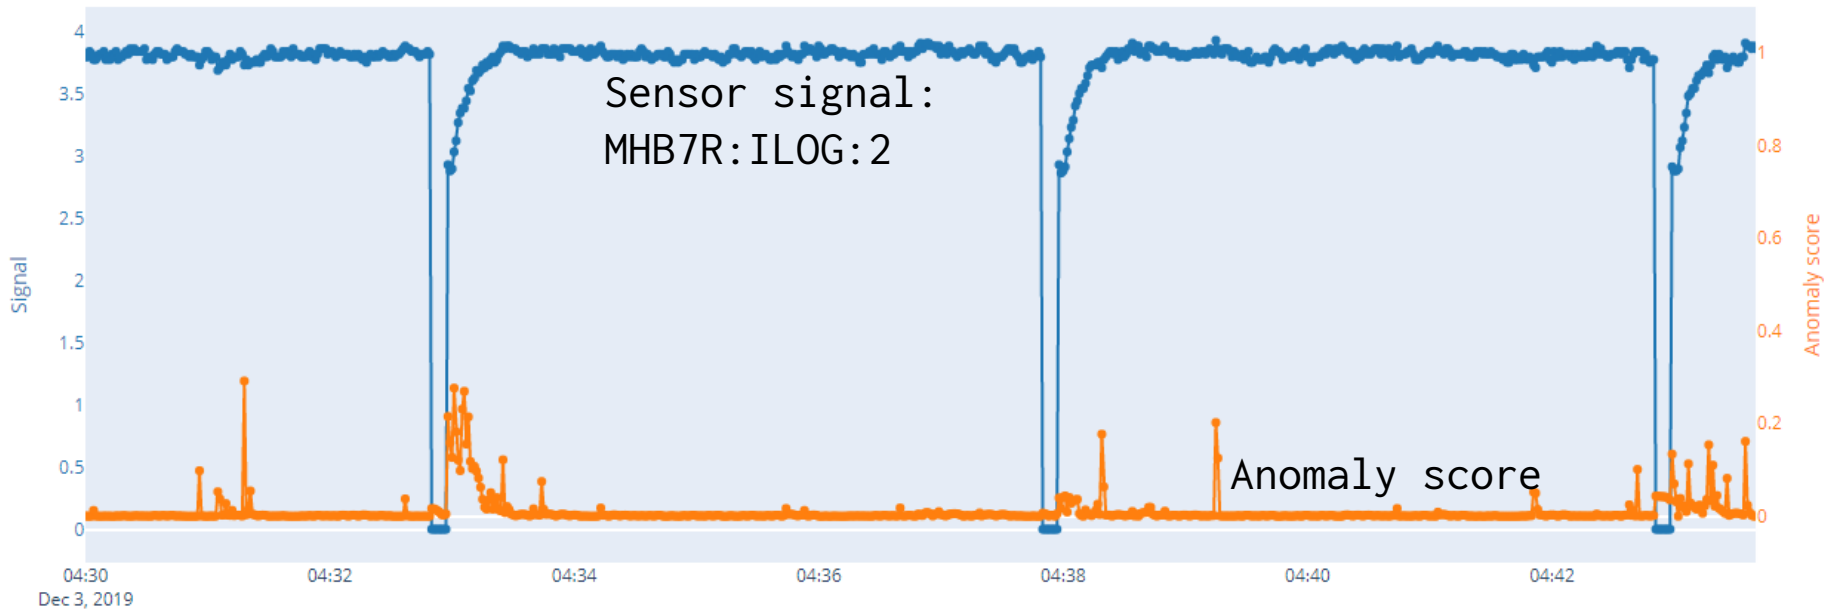

Normal

LSTM probabilities

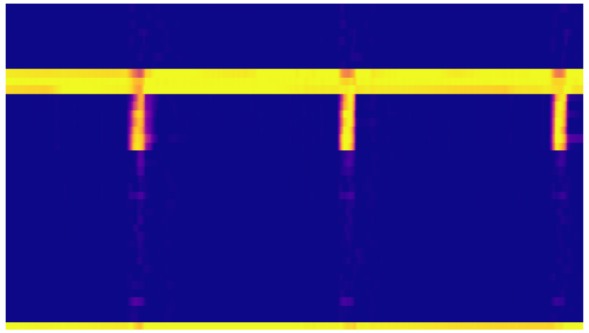

NUPIC probabilities

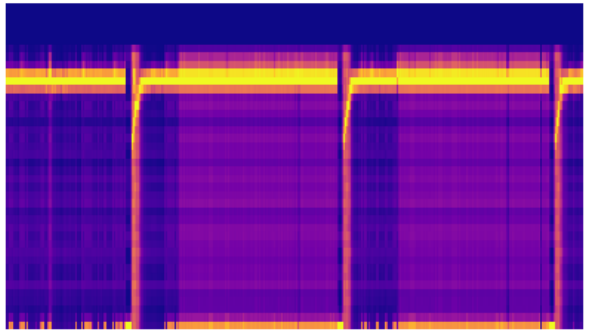

History

Start 2019-12-03 04:30:00

End 2019-12-03 06:00:00

Replay from here?

Refresh

No alarms

Anomaly detection in HIPA

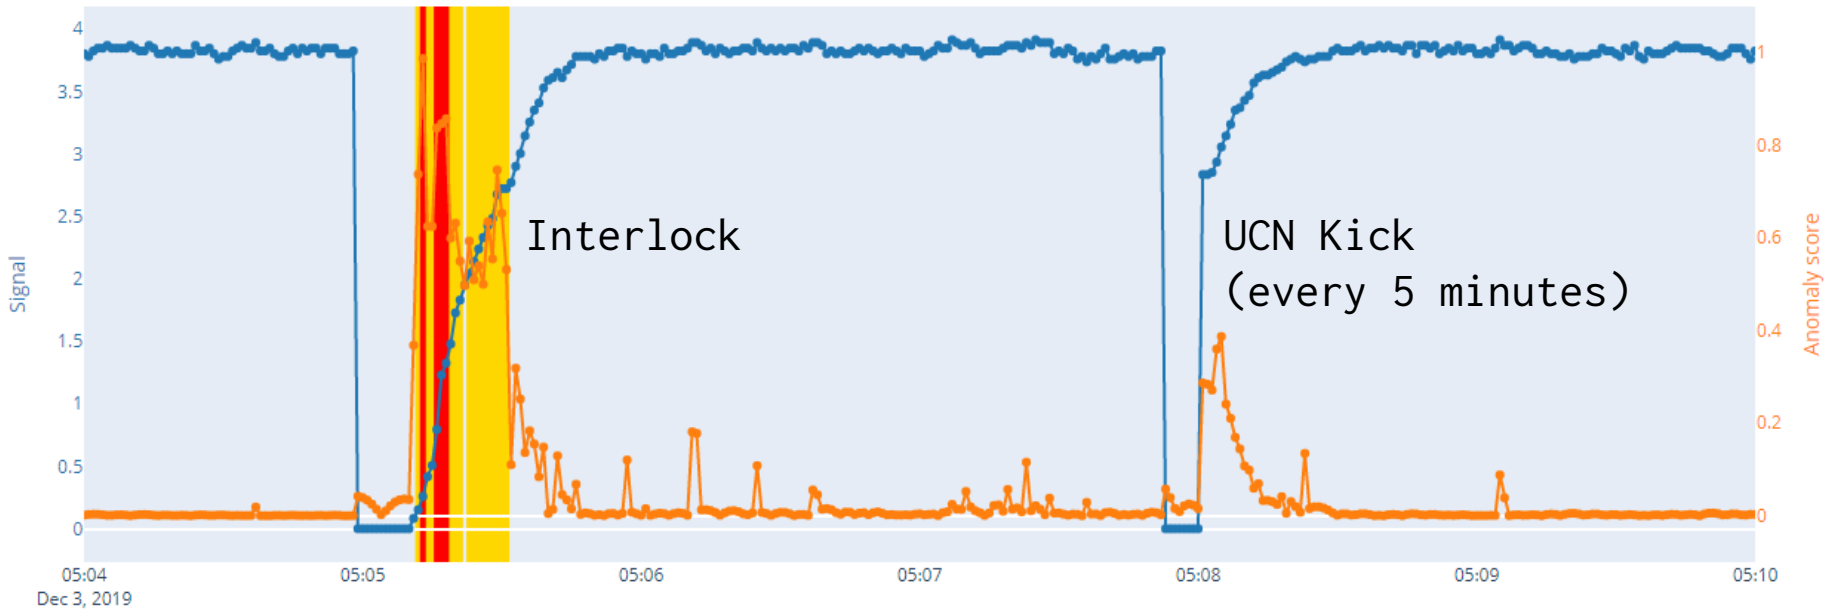

Normal

LSTM probabilities

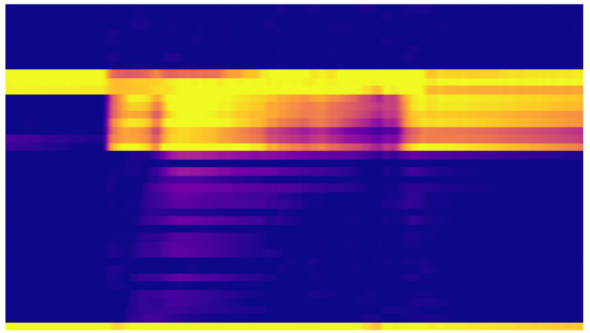

NUPIC probabilities

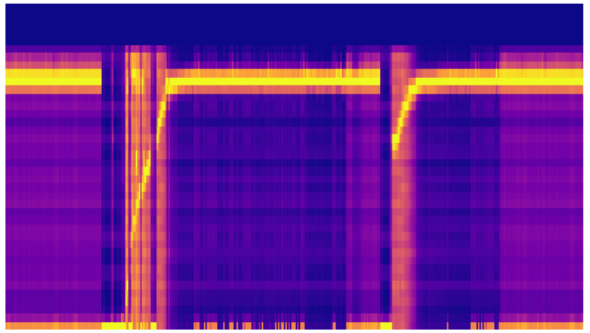

History

Start 2019-12-03 05:04:00

End 2019-12-03 05:10:00

Replay from here?

Refresh

No alarms

Anomaly detection in HIPA

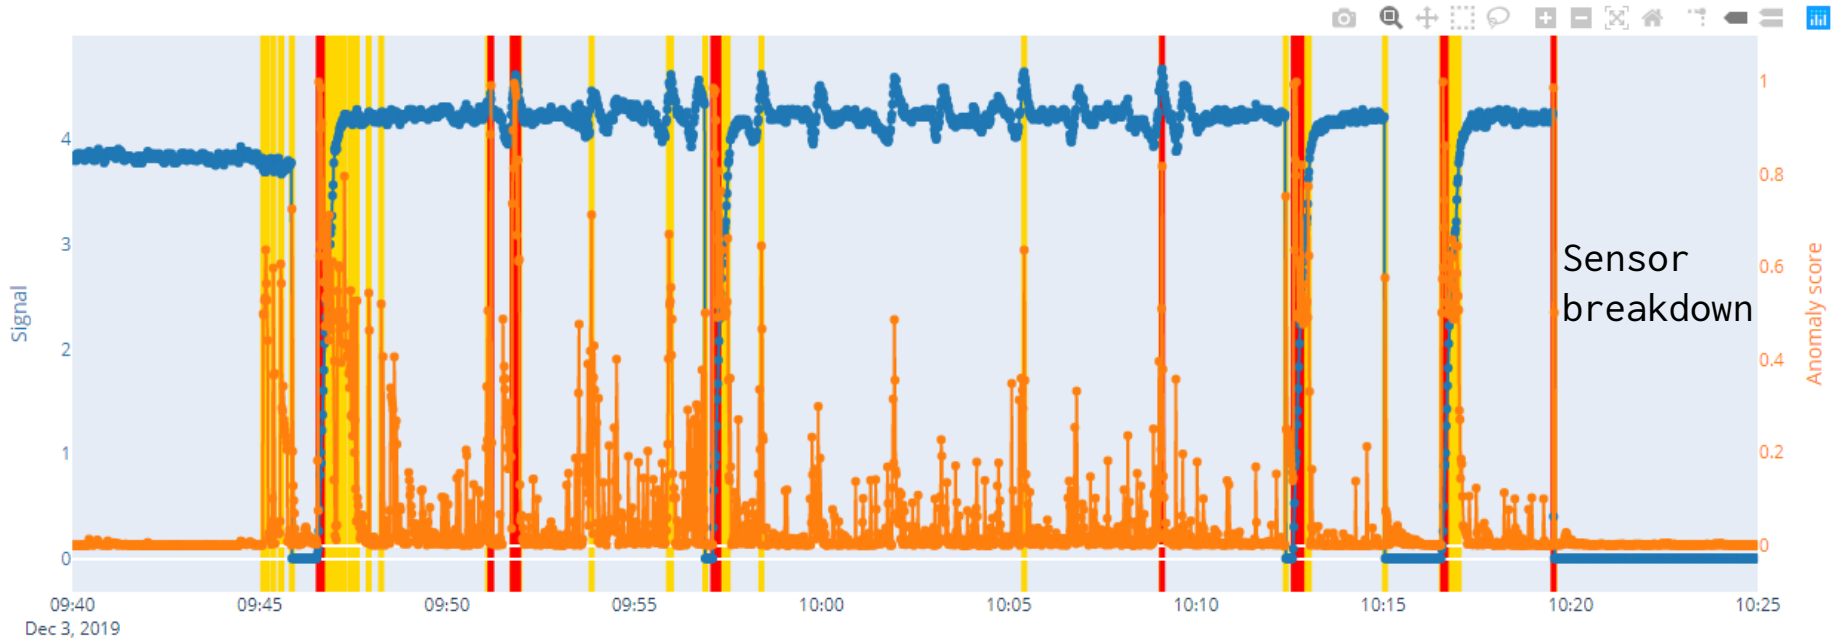

History

Start 2019-12-03 09:40:00

End 2019-12-03 10:25:00

Replay from here?

Refresh

No alarms

Normal

LSTM probabilities

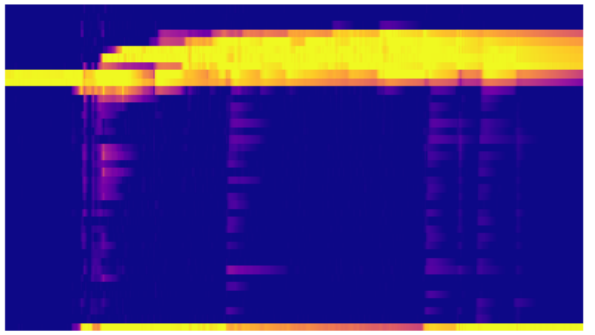

NUPIC probabilities

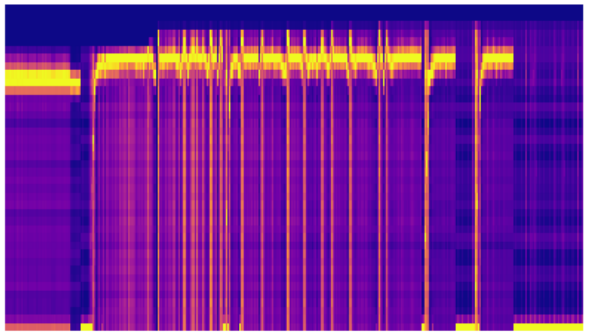

Conclusions:

- The combination of the HTM network and the LSTM seems to provide a good real time anomaly detection system.
- The system works well on archived data, adapts to both stable and periodic behaviour.
- It would have alerted the experts of the failure of the MHB7R:ILOG:2 sensor tenths of minutes in advance.

Further work:

- Real-time test on HIPA (and PROSCAN) are foreseen.
- Tests with many correlated signals may improve the sequence memory.
- Combine or compete against other models? Suggestions?

A nice video series to learn about the HTM:
https://www.youtube.com/playlist?list=PL3yXMgtrZmDqhsFQzwUC9V8MeeVOQ7eZ9