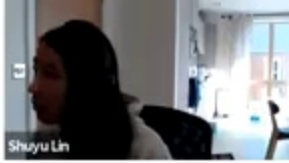

Introduction to Representation learning: Approaches, Challenges and Applications

Shuyu Lin
2016 AIMS cohort,
Cyber Physical Systems Group,
Department of Computer Science,
University of Oxford

What is Representation Learning?

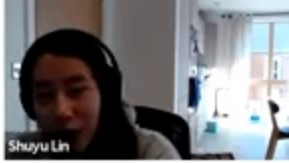

What is Representation Learning?

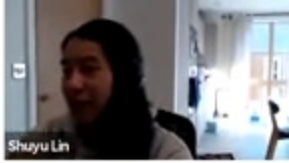

- What is the representation of a cat?

What is Representation Learning?

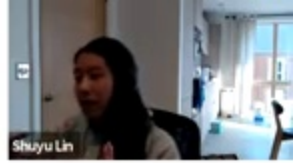

- Representation is a summary of data, which
  - Omits the unnecessary details,
  - And preserves important content.

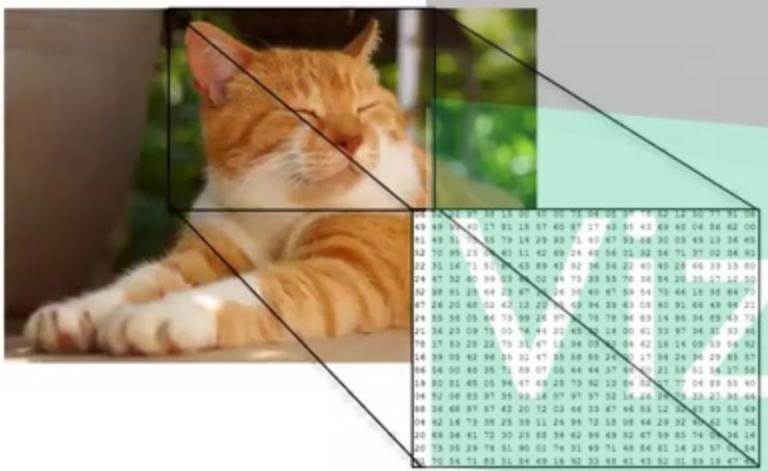

Evolution of Representation Learning

Classifier on raw data

Task

Classifier

or

Evolution of Representation Learning (co

End-to-end training of NN on task

Task

Deep Model

Learning transferable features

Task 1 Task 2 Task 3

Features

Learnt from data

Deep Model

Good

Ideal


or

?


or

?

What makes a Good Representation?

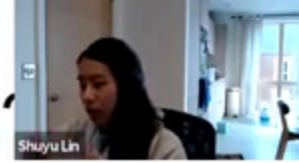

- Low dimensional
- Reusable across tasks
- Smooth and spatially coherent
- Disentangled
- Hierarchical and meaningful

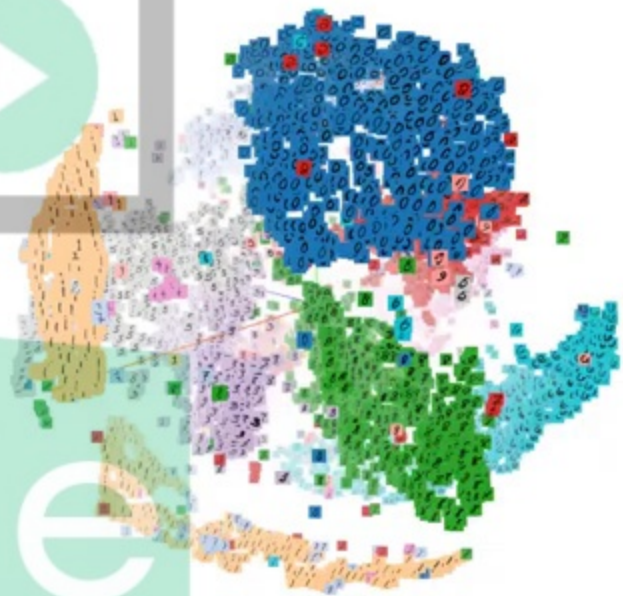

How to learn a representation?

- Generative models

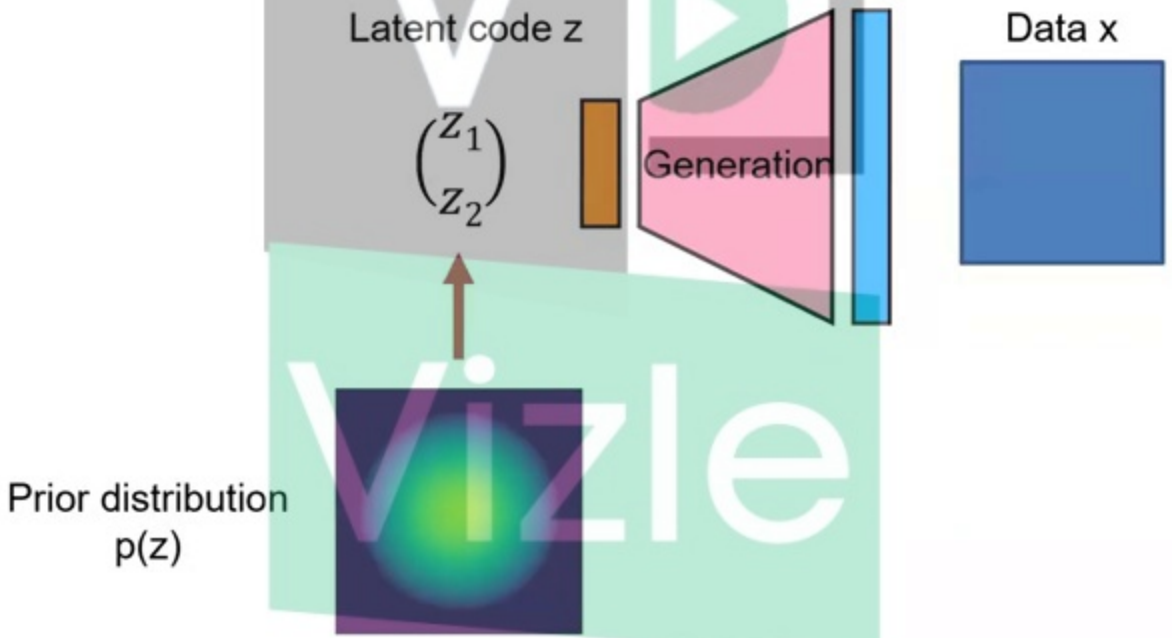

How to learn a representation?

- Variational autoencoders (VAEs)

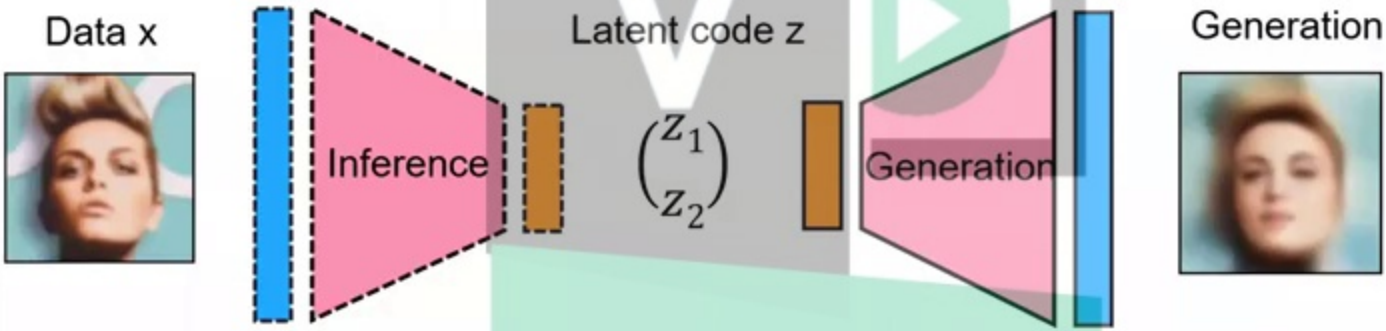

- What's the difference between VAEs and autoencoders?
  - Probabilistic formulation

How to use a learnt representation?

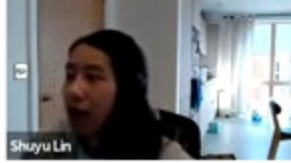

- Anomaly detection in time series
- Latent space interpolation

Anomaly detection in time series

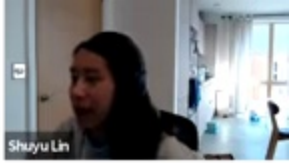

Example 1:
Microsoft Share Price

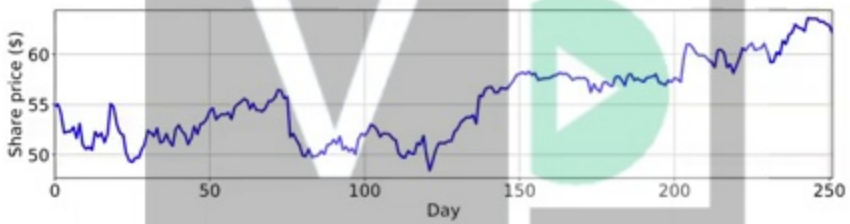

Avoid trading as the market is unusual?

Example 2:
Amazon Server Utilization

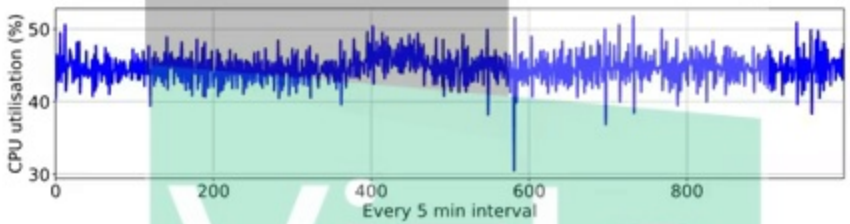

Is there any attack to my service?

Example 3:
A Patient EEG readings

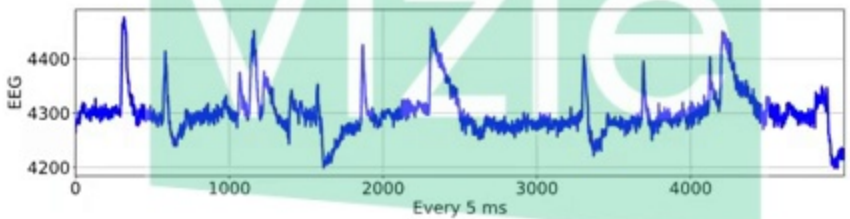

Is my patent in a critical condition and needs intensive care?

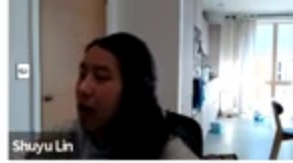

Anomaly detection is very difficult

Because:

1. Expertise and domain knowledge is often required to correctly identify an anomaly.

   → Labels are rare and expensive.

2. Anomalous events are rare and behaviours of anomalous events vary significantly from time to time.

   → Positive events are rare and hard to generalize.

As a result, anomaly detection for time series is practically an **unsupervised learning** task.

A high-level sketch for anomaly detecti

To respect the unsupervised learning nature of an anomaly detection task, we

Step 1: Aim to model the normal behaviours very well during training.

Training data: **no anomaly labels**.

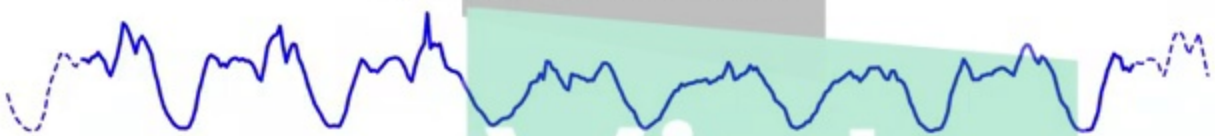

How to use learnt representation for anomaly detec

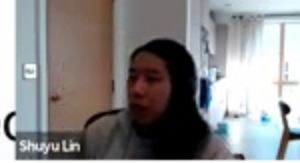

Anomaly detection with a variational autoencoder (VAE) [1]

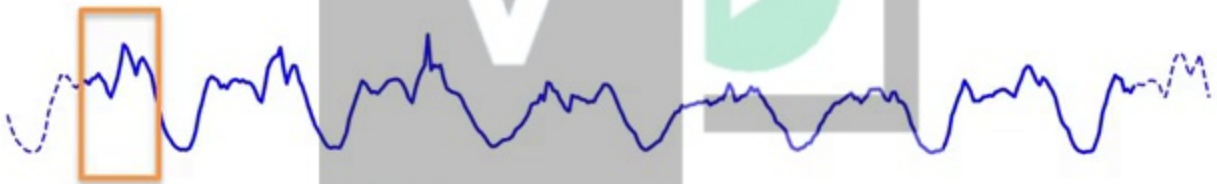

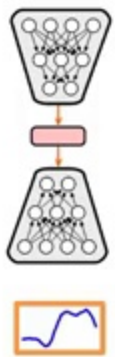

[1] Jinwon An and Sungzoon Cho, "Variational autoencoder based anomaly detection using reconstruction probability," 2015

Anomaly detection with VAE-LSTM hybrid mo

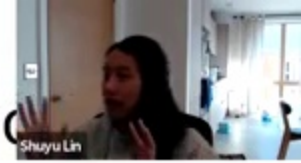

First, we train a VAE model to extract local information of a short window into a low-dimensional embedding.

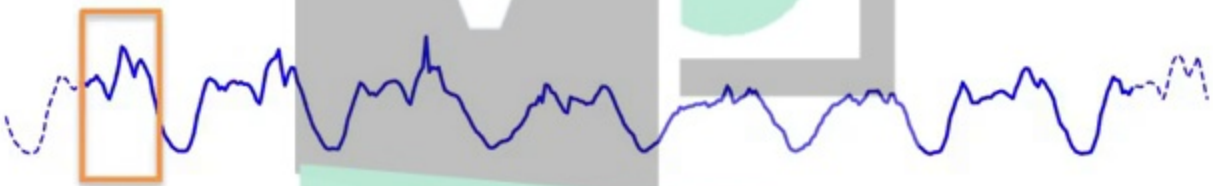

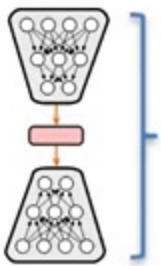

a VAE model

[2]Anomaly Detection for Time Series Using VAE-LSTM Hybrid Model

Anomaly detection with VAE-LSTM hybrid mo

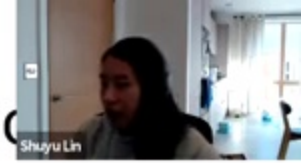

Then, we use a LSTM model, which acts on the low-dimensional embeddings produced by the VAE model, to manage the sequential patterns over longer term.

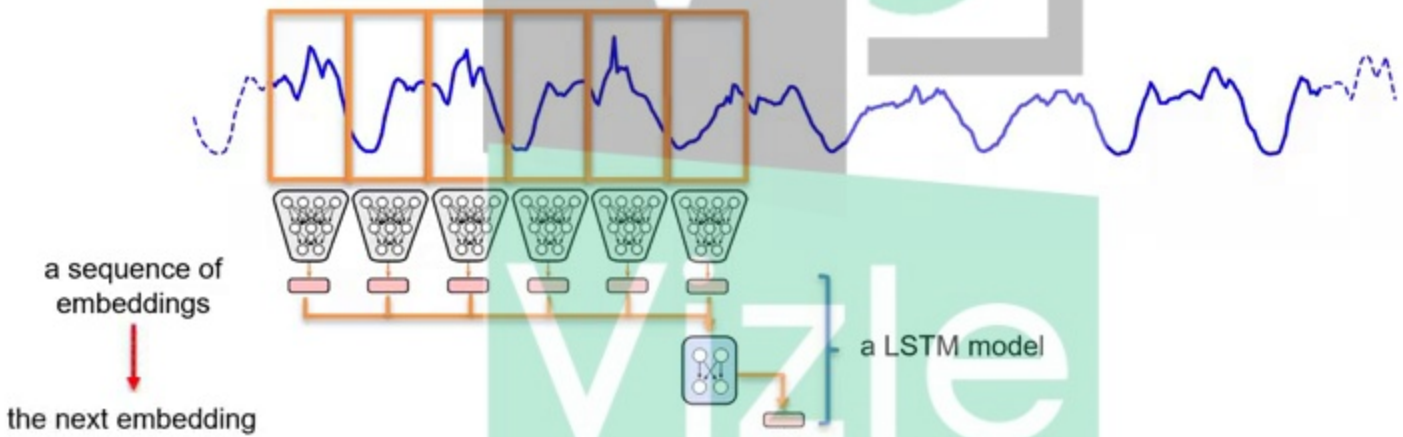

[2]Anomaly Detection for Time Series Using VAE-LSTM Hybrid Model

Anomaly detection with VAE-LSTM hybrid mo

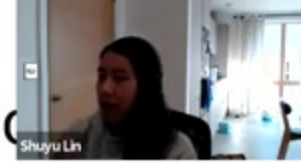

After both the VAE and the LSTM model are trained, we can detect anomalies using the reconstruction error of the next window.

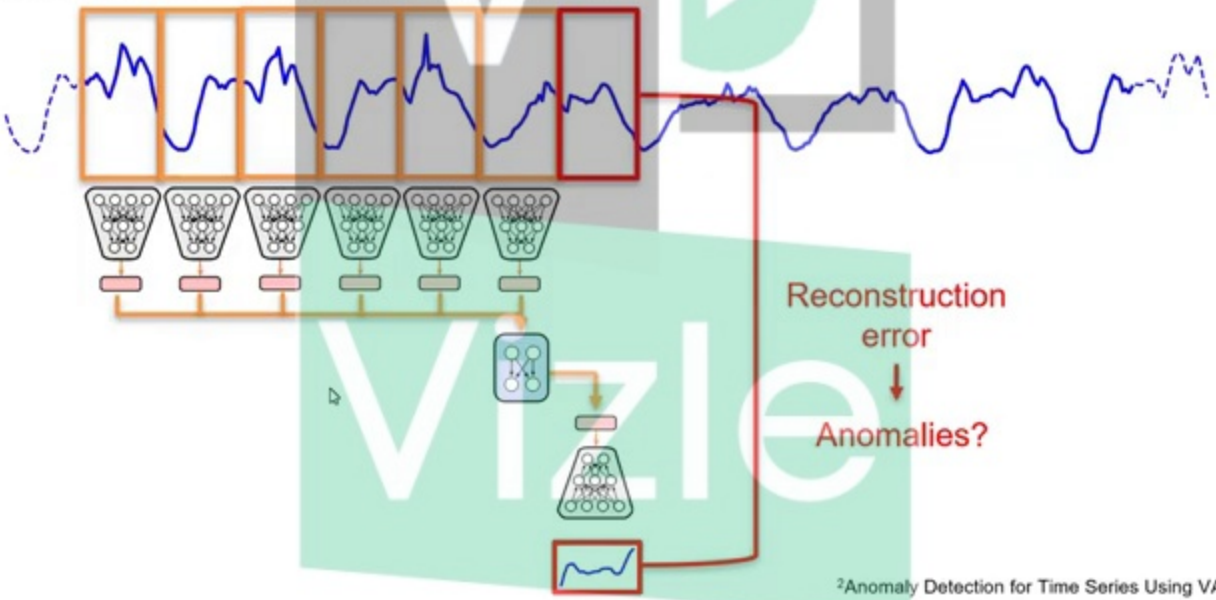

[2]Anomaly Detection for Time Series Using VAE-LSTM Hybrid Model

Anomaly detection with VAE-LSTM hybrid mo

The benefits of our design are clear:

- Our VAE-LSTM hybrid model can extract information beyond a short local window.
- The VAE module forms **robust local features**.
- The LSTM module estimates the **long-term correlations** in the sequence.
- As a result, our detection algorithm can identify anomalies that might span over **multiple time scales**.

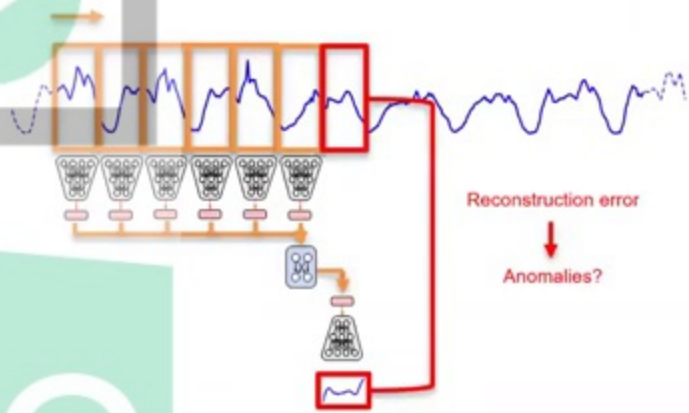

[2]Anomaly Detection for Time Series Using VAE-LSTM Hybrid Model

Detection results

Our VAE-LSTM hybrid model

One false positive.
Good for safety-critical situations.

- Detect all anomalies
- 1 false positive, but it is a case worth further inspection

Detection results

Our VAE-LSTM hybrid model

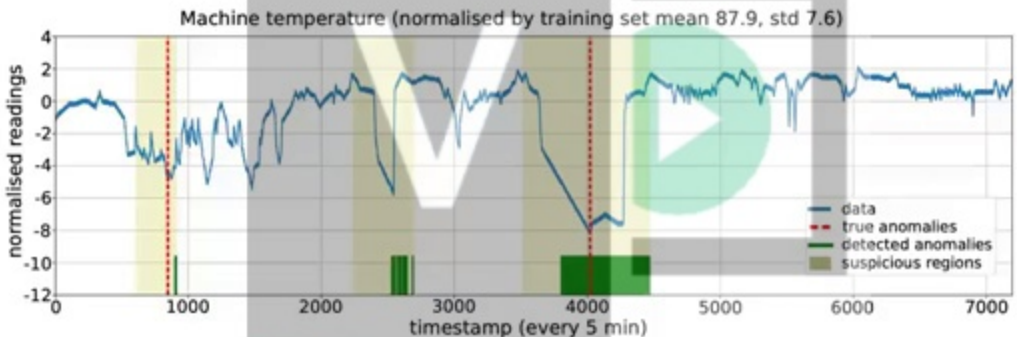

- Detect all anomalies
- 1 false positive, but it is a case worth further inspection

VAE model

- Detect all anomalies
- Too many false positives