

NASA Jet Propulsion Laboratory, March 25th, 2020

Unsupervised anomaly detection in high-dimensional flight data using Convolutional Variational Auto-Encoder


Milad Memarzadeh, Ph.D.

with: Bryan Matthews, Ilya Avrekh, and Daniel Weckler

Data Sciences Group, NASA Ames Research Center

Motivation for unsupervised anomaly detection

Accident rate of commercial flights has been **cut in half** since 1999.




Source: National Transportation Safety Board

Motivation for unsupervised anomaly detection

Accident rate of commercial flights has been **cut in half** since 1999.

Passenger load factor has increased to 82.5% in 2017.



Source: National Transportation Safety Board

Motivation for unsupervised anomaly detection

Accident rate of commercial flights has been **cut in half** since 1999.

Passenger load factor has increased to 82.5% in 2017.

Creating labels for data in the aviation domain:

- requires huge effort from subject-matter experts.
- is largely expensive and impractical.

Hence, **unsupervised machine learning** is the only feasible choice.


Aviation anomaly detection literature

Exceedance detection:

Comparing against the **pre-defined thresholds**, which are fixed by subject-matter experts.

Cons:

- complete reliance on domain knowledge.
- can only identify known anomalies.




Aviation anomaly detection literature

Distance-based anomaly detection:

Anomalies are defined as a point in the feature space whose **nearest neighbors are far from it.**

Cons:

- poor performance when dealing with high-dimensional data
- subject to error depending on percentage of anomalous patterns existing in the training data




Aviation anomaly detection literature

Kernel-based anomaly detection:

One-Class SVM: Finds a maximal gap hyperplane that separates data from the origin (as the only data point of the none-existent class).

Cons:

- computational complexity of the kernel building step.
- poor performance in dealing with high-dimensional data.



Building upon recent advancements in machine learning


Variational Auto-encoder:

Finding a trade-off between reconstruction quality and disentanglement of the latent space (i.e., **balancing the bias and variance of the model**):

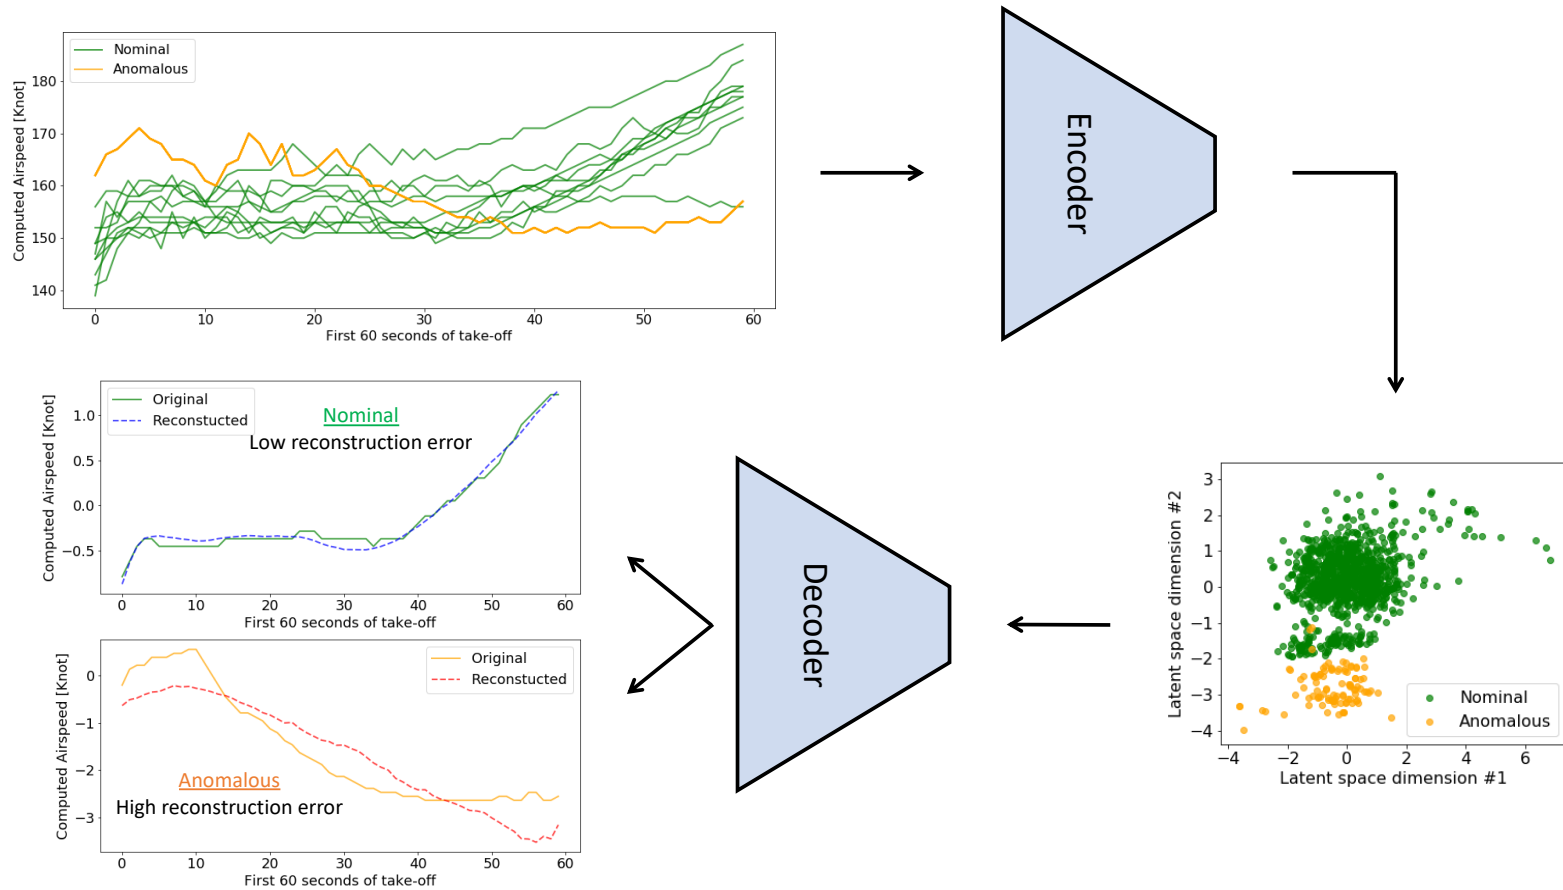
$$L(\theta, \phi; x) = \mathbb{E}_{q_{\phi}(z|x)}[\log p_{\theta}(x|z)] - \beta \text{KL}(q_{\phi}(z|x) \parallel p(z))$$

represents the reconstruction error

how close the posterior is to prior



The visual illustration of the method

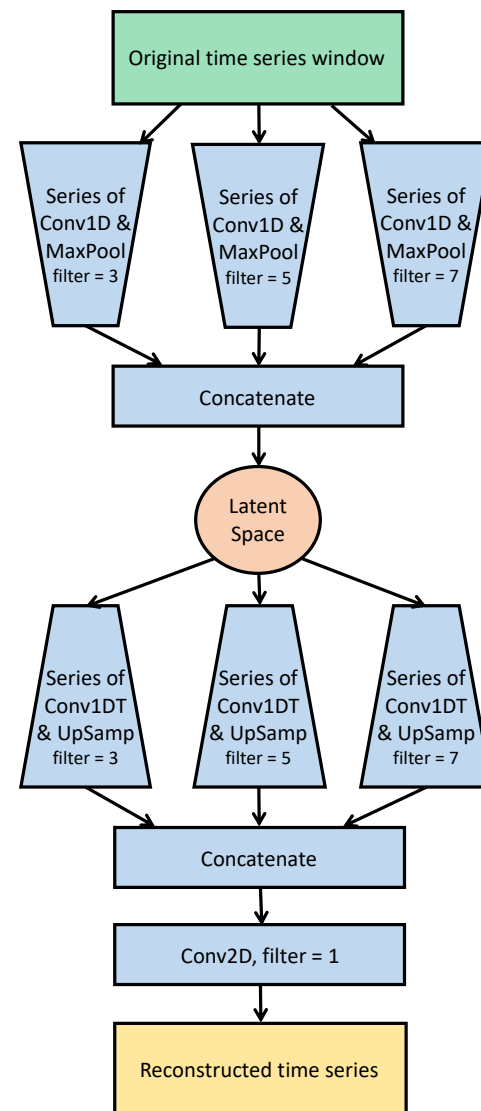


A convolutional network for anomaly detection

Convolutional Variational Auto-encoder:





Using convolutional layers (instead of recurrent) to **speed up the training process.**

Using multiple filter sizes to capture **local and global temporal dependence** in the time series.



Validation on Yahoo!'s benchmark dataset

Data set consists of real (A1) and synthetic (A2) univariate time series as well as synthetic multivariate time series without (A3) and with (A4) change points.



Validation on Yahoo!'s benchmark dataset


Data set consists of real (A1) and synthetic (A2) univariate time series as well as synthetic multivariate time series without (A3) and with (A4) change points.

We compare performance to:

- **ADOPT**: which is a supervised learning approach using recurrent neural network
- **Kmeans ++**
- **One-Class SVM (OC-SVM)**

On average:

- **CVAE outperforms Kmeans++ and OC-SVM** with 62% higher precision and 30% higher recall.
- CVAE has 10% lower precision and 17% lower recall compared to ADOPT.



Flight operational quality assurance (FOQA) data


Drop in airspeed case study:

- Our goal is to identify anomalies in the **first 60 seconds of commercial flight's take-off**.
- Data consists of **30,000 nominal** take-offs and **1000 anomalous** ones.
- Each time series consists of **17 variables** measuring the roll attitude, altitude information, pitch attitude, speed information and yaw attitude.

Flight operational quality assurance (FOQA) data

Drop in airspeed case study:

- Our goal is to identify anomalies in the **first 60 seconds of commercial flight's take-off.**
- Data consists of **30,000 nominal** take-offs and **1000 anomalous** ones.
- Surrogate labels for anomalies by subject matter experts: if drop in air speed is **more than 20 knots**




Anomaly detection performance on the FOQA data

CVAE outperforms Kmeans++ and OC-SVM with **24% higher precision and 26% higher recall**.

CVAE-Semi improves the performance of CVAE by 25%:


- is the semi-supervised approach of CVAE, where **the training data is only comprised of nominal data**.



Interpretability and explain-ability of the findings

Effect of hyper-parameter β :

Higher values of β result in a model that is more robust to existing of anomalies in the training data, but decreases the interpretability of the latent space.



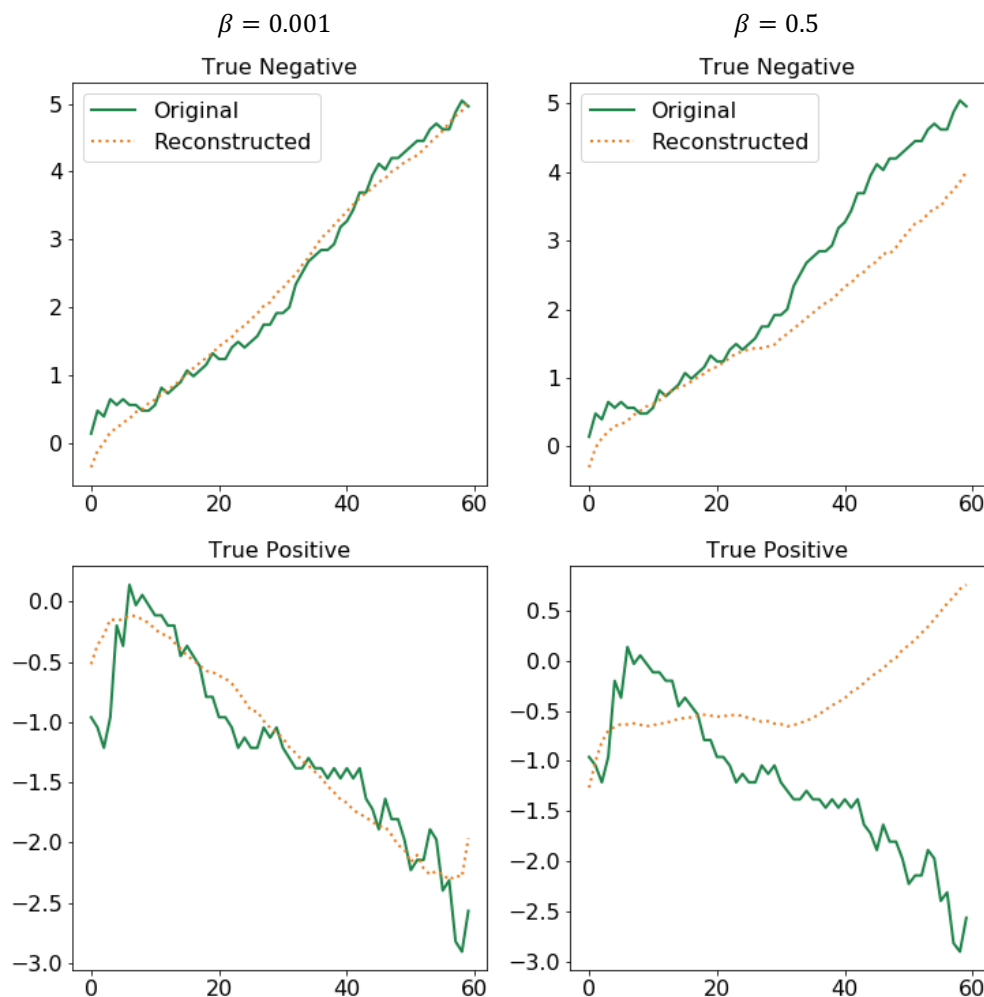
Interpretability and explain-ability of the findings

Effect of hyper-parameter β :

Higher values of β result in a model that is **more robust to existing of anomalies in the training data**, but **decreases the interpretability of the latent space**.

Higher values of β also result in **higher reconstruction error**.

This means that a model with very low β can **over-fit to reconstruct the anomalous patterns**.



Concluding remarks and future work

We developed an approach to detect anomalies in an **unsupervised fashion** for **high-dimensional heterogenous time-series** data.

Our approach **significantly outperforms** clustering/distance/kernel-based methods that are common in the domain.

Future works:

Develop an architecture to accommodate the binary channels of FOQA data:

- **State-based anomaly detection** by building the state space based on binary variables.

Improve the **explain-ability** of the findings in the latent space.

Test and validate the approach on other case studies in **aviation anomaly detection**.