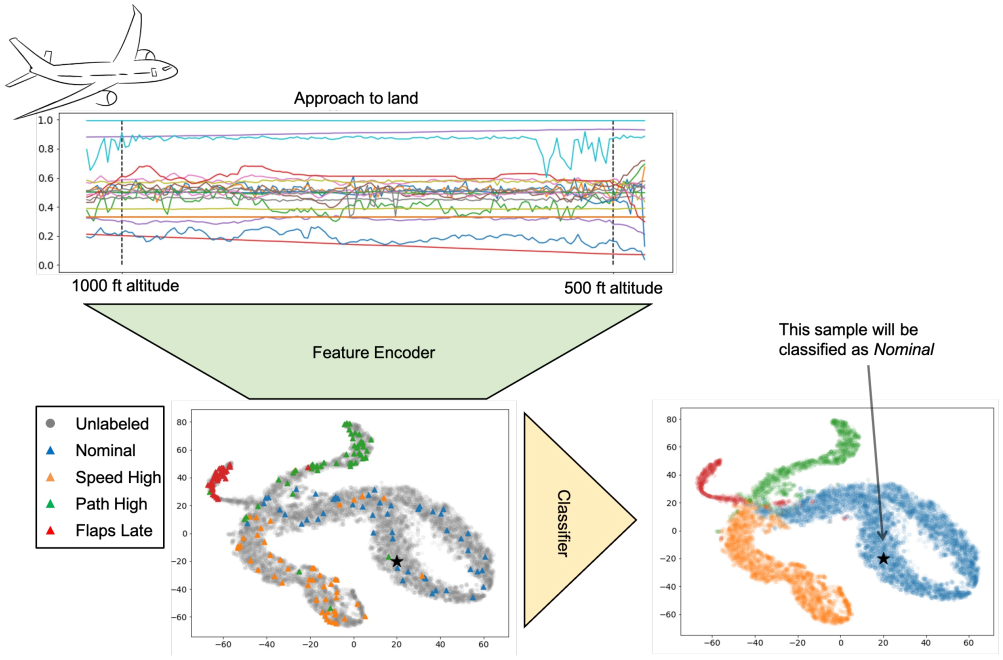

National Aeronautics and Space Administration

Anomaly Detection in Flight Operational Data Using Deep Learning

Milad Memarzadeh, Bryan Matthews, and Daniel Weckler
Data Sciences Group, NASA Ames Research Center

System Wide Safety

In this presentation

Learn about the literature in anomaly detection in flight operational data

Dive deep into two developed deep learning models:

1. Convolutional Variational Auto-Encoder (CVAE): an unsupervised encoder-decoder model for anomaly detection in multivariate time-series
   - Demonstration of CVAE: Fuser Streaming Data Prototype.
2. Robust and Explainable Semi-supervised Anomaly Detection (RESAD): a semi-supervised deep learning architecture capable of detecting multiple anomalies with limited number of labeled data.
   - Demonstration of RESAD: Interactive Data Visualization Walkthrough.

Aviation anomaly detection literature

National Aeronautics and
Space Administration

Exceedance detection:
Comparing against the pre-defined thresholds, which are identified by subject-matter experts.

Example of unstable approach to landing

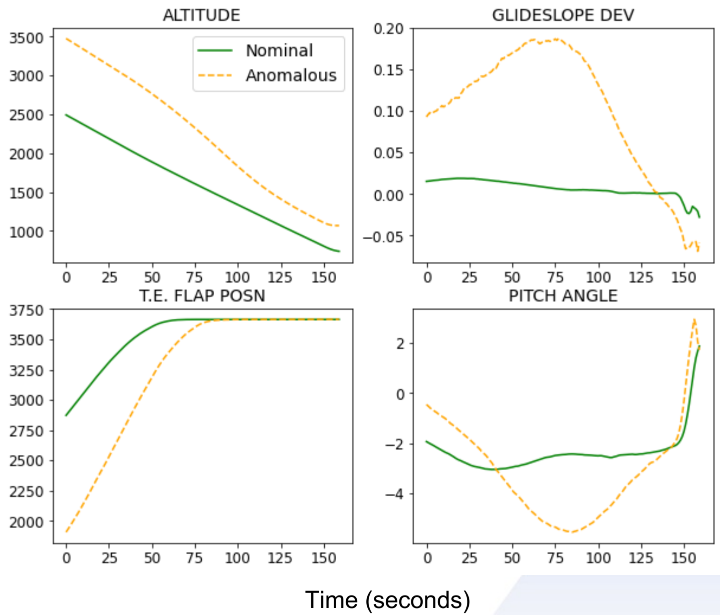

Aviation anomaly detection literature

National Aeronautics and
Space Administration

Exceedance detection:
Comparing against the pre-defined thresholds, which are identified by subject-matter experts.

Cons:

- Complete reliance on domain knowledge.
- Requires extensive reviews of entire data.
- Can only identify known anomalies.

Example of unstable approach to landing

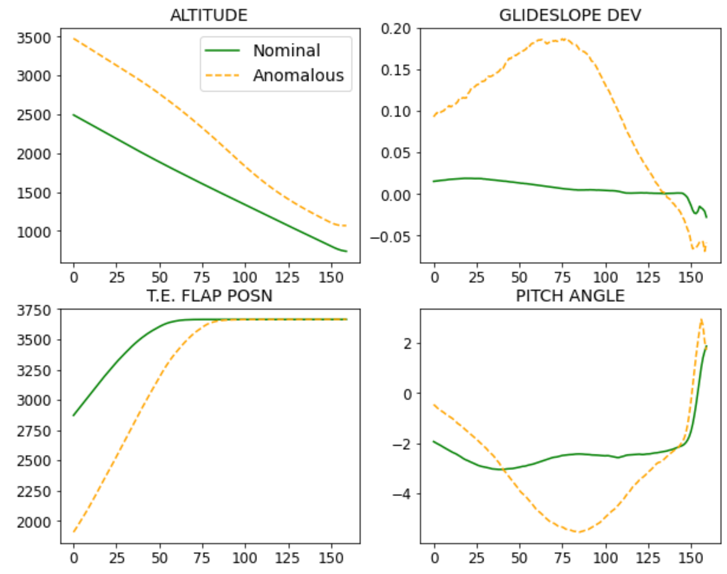

Time (seconds)

Aviation anomaly detection literature

Supervised learning:

Produces inference using only labeled data.

Pros:

- It demonstrates amazing performance when a sufficient number of labeled data is available.

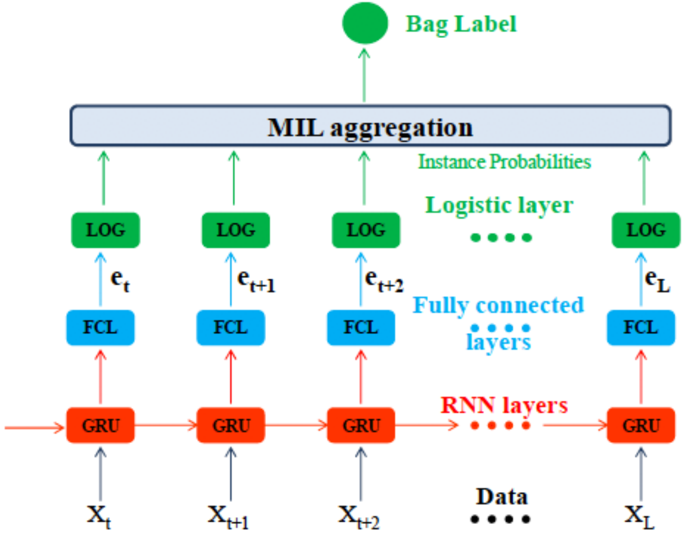

Jankiraman (2018), *KDD*.

Aviation anomaly detection literature

National Aeronautics and Space Administration

Supervised learning:

Produces inference using only labeled data.

Cons:

- Can only identify known anomalies.
- Creating labels for data requires huge effort from subject-matter experts and is largely expensive and impractical.

Hence, unsupervised learning or semi-supervised learning are the only feasible choices.

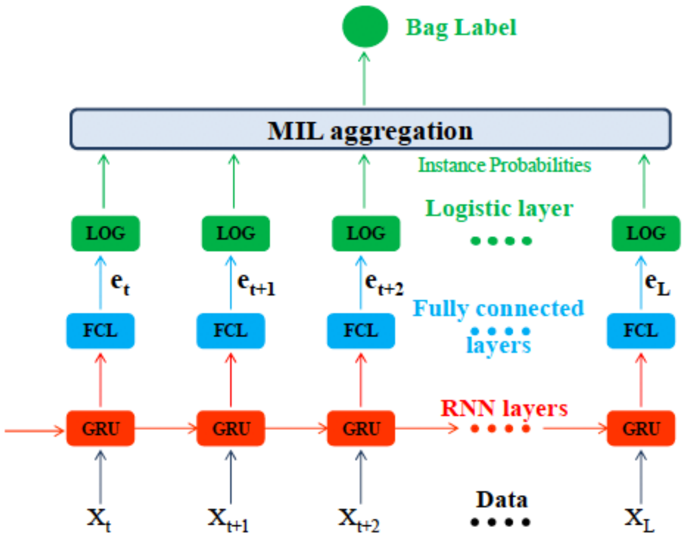

Jankiraman (2018), *KDD.*

Convolutional Variational Auto-Encoder - CVAE

National Aeronautics and Space Administration

Using deep auto-encoders to identify anomalies without the need for labels.

reconstruction fidelity

$$\mathcal{J}_{\text{CVAE}} = \mathbb{E}_{q_\phi(z|x)}[\log p_\theta(x \mid z)] - \beta \mathrm{KL}\left(q_\phi(z \mid x) \,||\, p_\theta(z)\right)$$

distance between posterior and prior

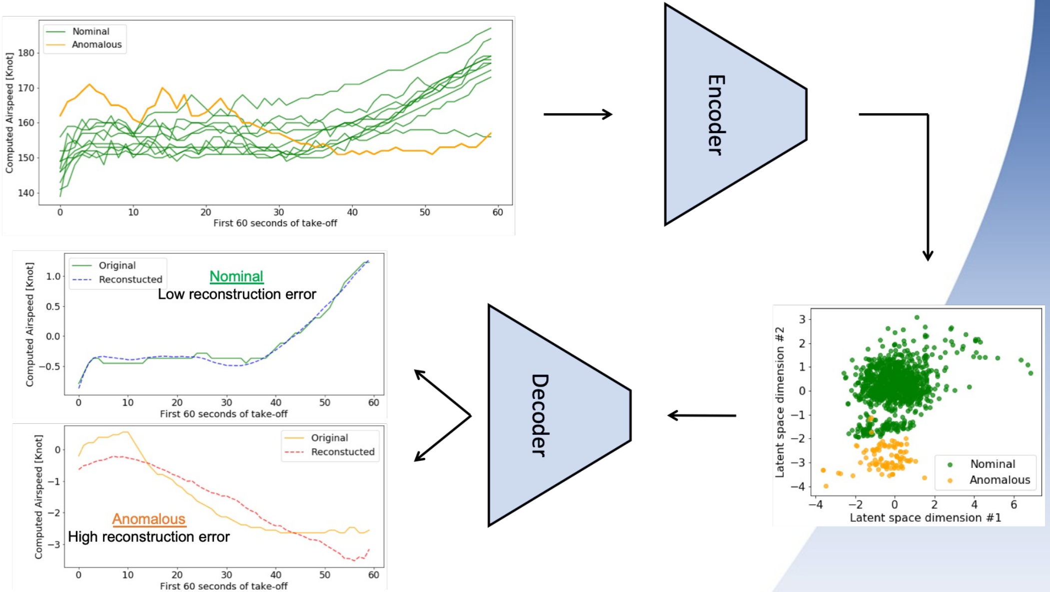

Aviation anomaly detection literature

Using deep auto-encoders to identify anomalies without the need for labels.

reconstruction fidelity

$$\mathcal{J}_{\text{CVAE}} = \mathbb{E}_{q_\phi(z|x)}[\log p_\theta(x \mid z)] - \beta \text{KL}\left(q_\phi(z \mid x) \,||\, p_\theta(z)\right)$$

distance between posterior and prior

Identifying anomalies:

$$\zeta_i = \|x_i - \hat{x}_i\|_2^2, i \in \{1, \ldots, N\}$$

$$thr = \mathbb{E}[\zeta] + \alpha\sigma(\zeta)$$

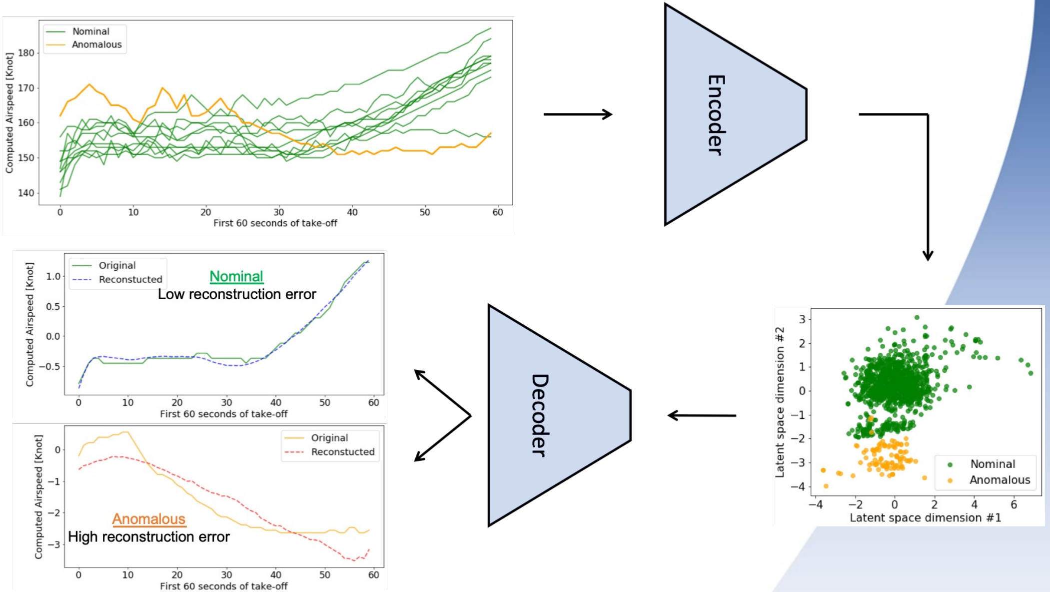

Open-Sourced Repository: https://github.com/nasa/CVAE

Memarzadeh et al. (2020), *Aerospace, 7(8)*, 115.

Performance comparison

National Aeronautics and
Space Administration

<u>Pros:</u>

- Does not require labels to make inference.

<u>Cons:</u>

- Low precision, which means a high number of false positives and low reliability.

- It is not easy to extend to multi-class anomaly detection.

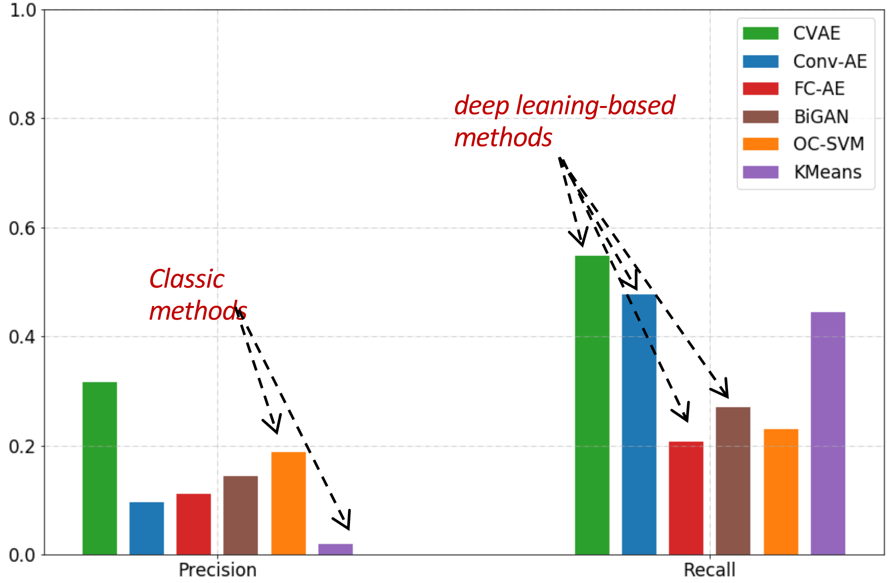

Open-Sourced Repository: https://github.com/nasa/CVAE

Memarzadeh et al. (2020), *Aerospace, 7(8)*, 115.

How to improve the reliability of unsupervised learning

National Aeronautics and Space Administration

Training CVAE (our model) only on nominal data improved the performance significantly:

- 36.8pp higher precision
- 27.3pp higher recall

<u>Takeaway</u>: how to take advantage of minimally labelled data that are available?

Open-Sourced Repository: https://github.com/nasa/CVAE

Memarzadeh et al. (2020), *Aerospace, 7(8)*, 115.

NASA's Digital Information Platform (DIP)

CVAE USE CASE: STREAMING FUSER DATA

Deploying CVAE on DIP Platform

- Data Source: Fuser
- Streaming radar track data:
  - TFM/ASDE-X
  - lat/lon, altitude, ground speed
- Focus on the last 20NM
  before landing at core 30 airports.
- Identify anomalous flight track.
- Generate anomaly report.

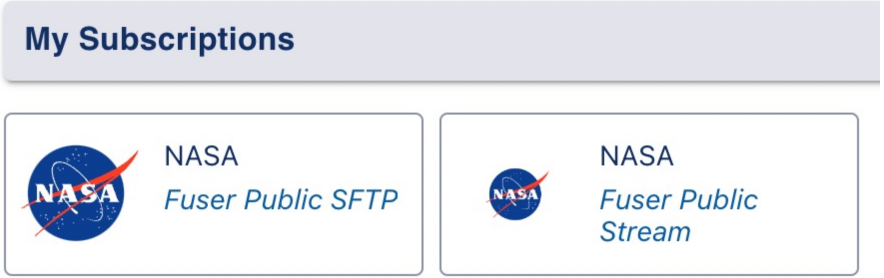

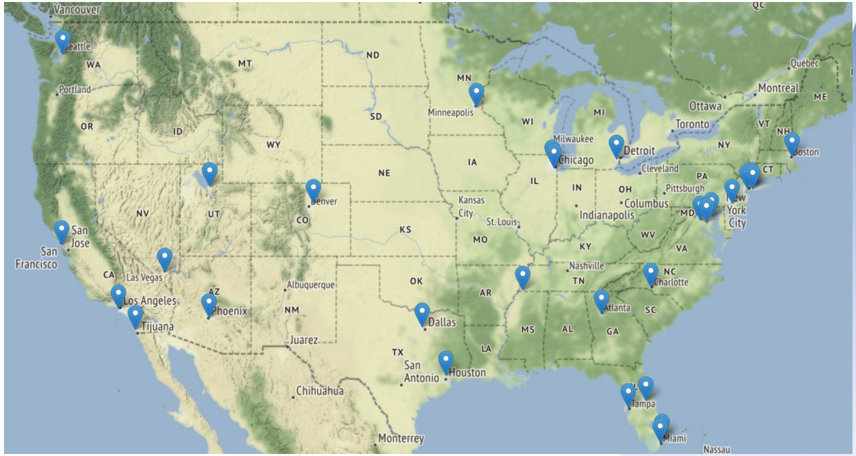

CVAE on Streaming DIP Fuser Data

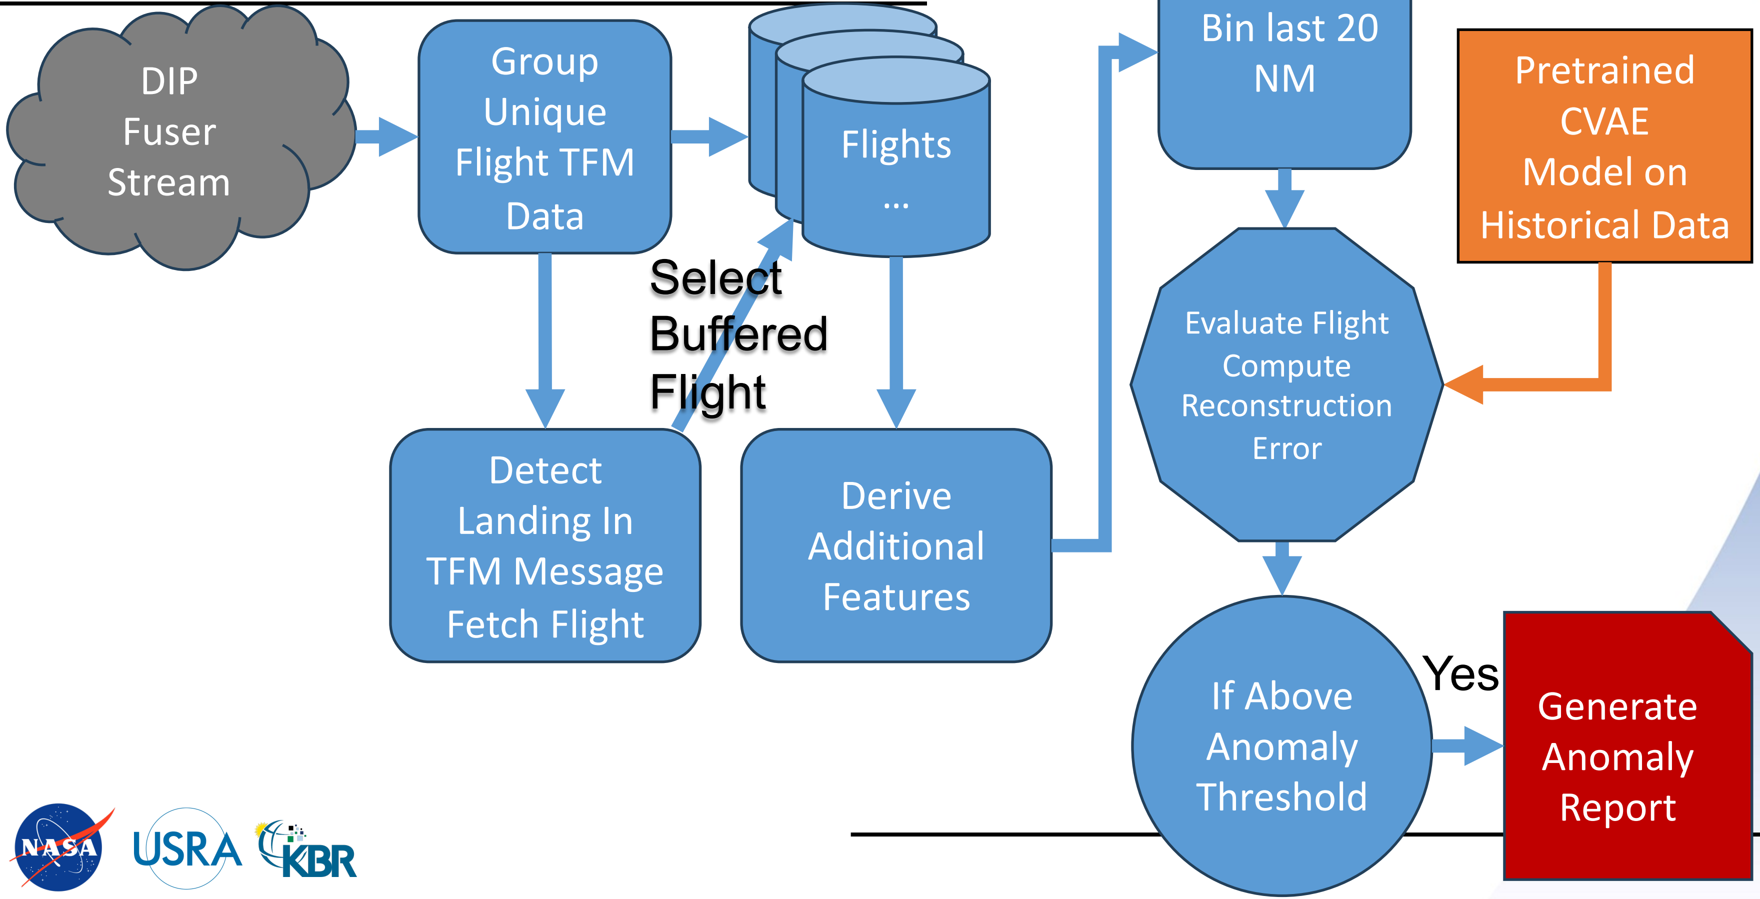

DIP Fuser TFM Anomalies

- Candidate: Short Turn to Final

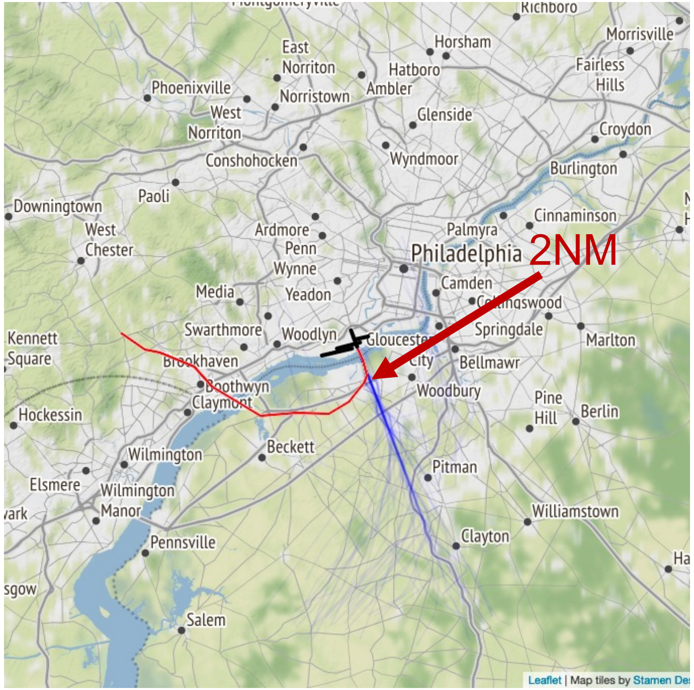

[redacted]807

National Aeronautics and Space Administration

Date: 06/04/23
GUFI: [redacted]807.RST.PHL.230604.1540.0091.TMA
Dest. Airport: KPHL
Runway: 35
AC Type: E55P

Anomaly Score: 340.673

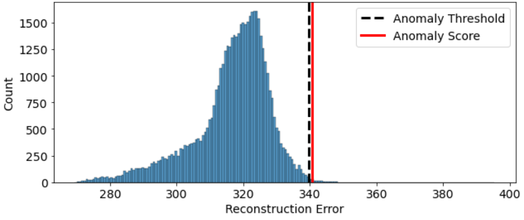

Variables:

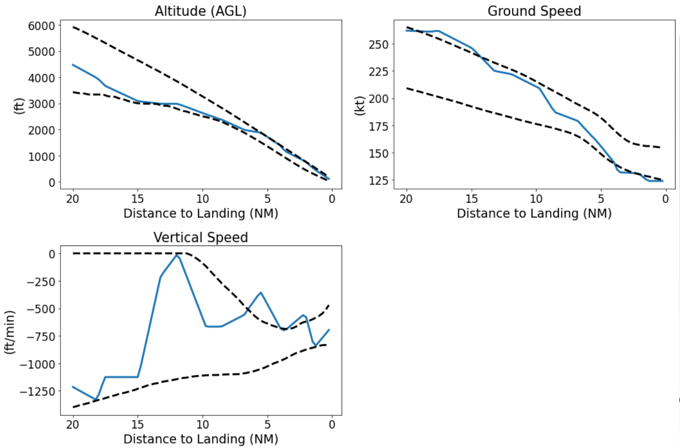

DIP Fuser TFM Anomalies

- Candidate: Potential Misalignment

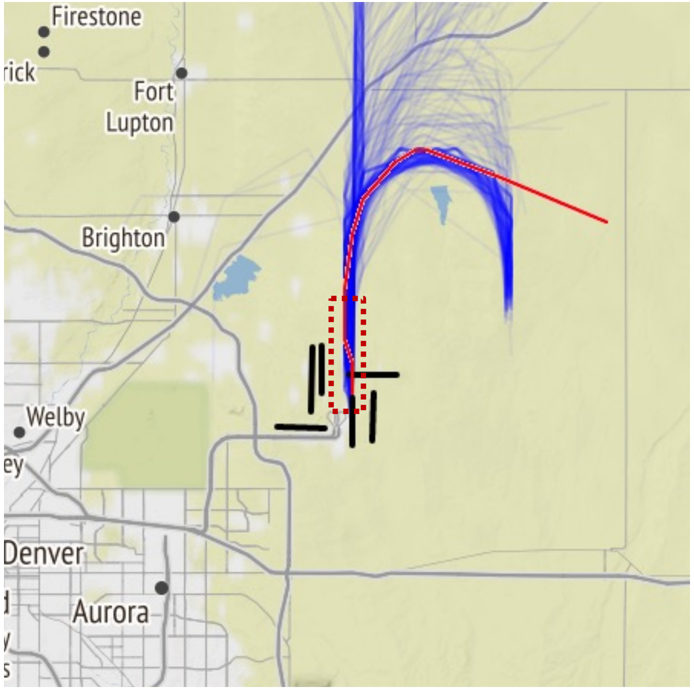

[redacted]736

```
Date:           04/25/23
GUFI:           [redacted]736.ATL.DEN.230424.1045.0031.TFM
Dest. Airport:  KDEN
Runway:         17R
AC Type:        A321
```

Anomaly Score: 344.143

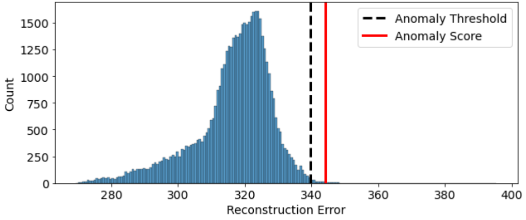

Variables:

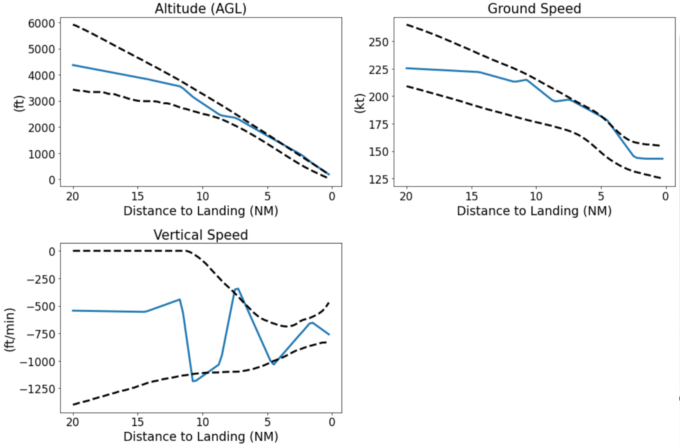

DEMO STREAMING CVAE

Multi-class anomaly detection case study based on real flight data

National Aeronautics and Space Administration

Each data instance is 160-s recording of 19 variables during approach of a commercial aircraft to landing. Attributes cover a variety of systems, including the state and orientation of the aircraft, positions and inputs of the control surfaces, engine parameters, and auto pilot modes and corresponding states.

Training data consists of 18,313 samples falling into four classes:

1. Nominal (66.7%)
2. Speed Anomaly (22.9%)
3. Path Anomaly (7.2%)
4. Control Anomaly (3.2%)

Separate **test data** of 6105 samples is used for evaluating the models.

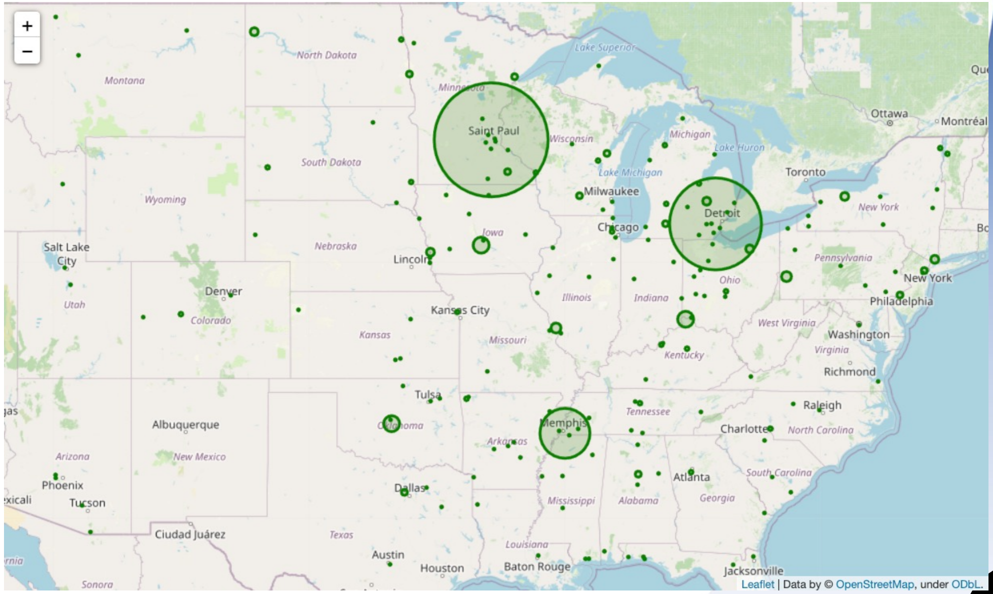

Data from: https://c3.nasa.gov/dashlink/projects/85/

USRA KBR

Robust and Explainable Semi-supervised Anomaly Detection (RESAD)

National Aeronautics and Space Administration

Objective = $w_c \mathbb{E}_{(X_L, y_L)}$ [classification performance] +

$w_e \mathbb{E}_{(x \in (X_L \cup X_U), y_L)}$ [latent space configuration/explainability] +

$w_r \mathbb{E}_{x \in (X_L \cup X_U)}$ [reconstruction fidelity]

Unsupervised learning ignores y_l, while supervised learning ignores X_U.

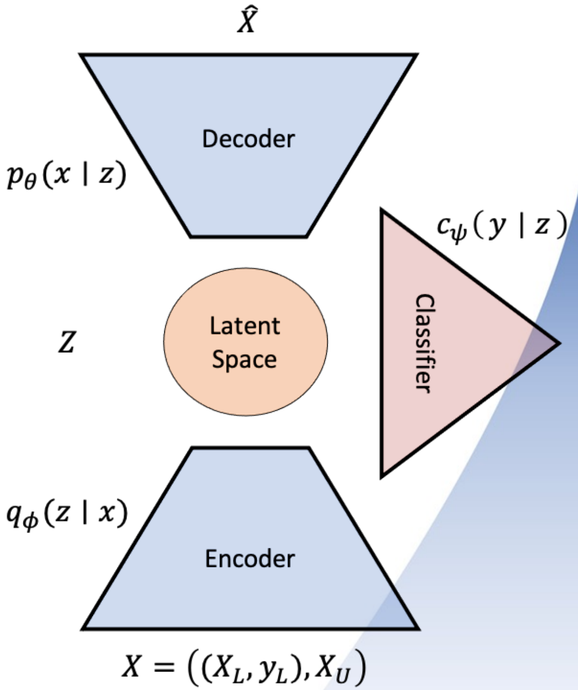

Memarzadeh et al. (2021), *AIAA Scitech & JAIS.*

RESAD: performance comparison

Objective = $w_c \mathbb{E}_{(X_L, y_L)}$ [classification performance] +

$w_e \mathbb{E}_{(x \in (X_L \cup X_U), y_L)}$ [latent space configuration/explainability] +

$w_r \mathbb{E}_{x \in (X_L \cup X_U)}$ [reconstruction fidelity]

Unsupervised learning ignores y_l, while supervised learning ignores X_U.

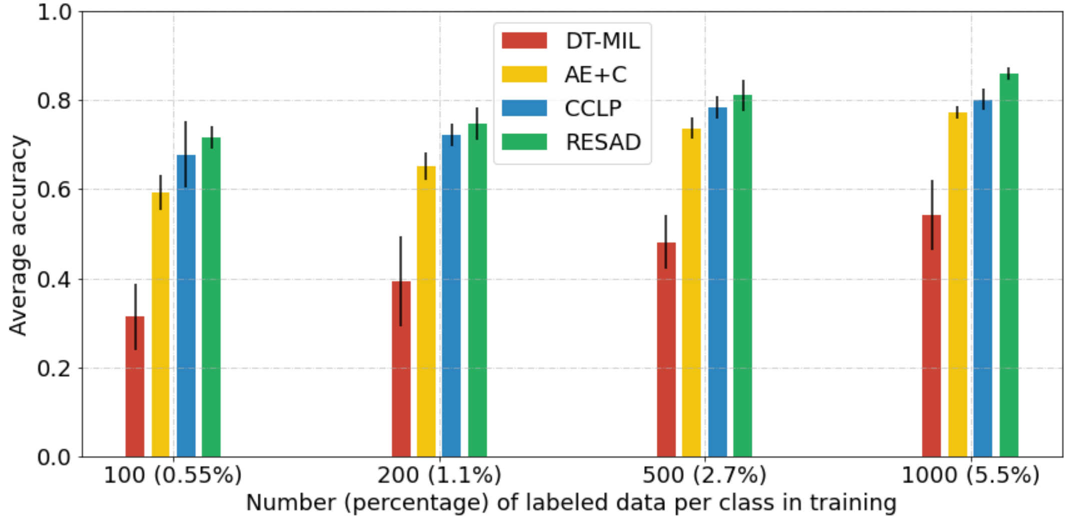

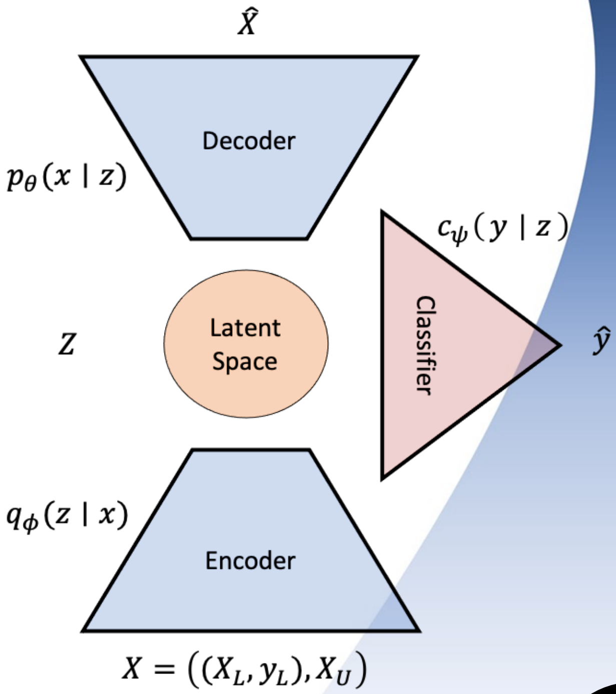

Memarzadeh et al. (2021), *AIAA Scitech & JAIS.*

Latent space configuration: the superiority of the CCLP approach

National Aeronautics and
Space Administration

Figures show 2D visualization of the 256D latent space of each model using t-Distributed Stochastic Neighbor Embedding (t-SNE), color-coded based on the actual class of the data.

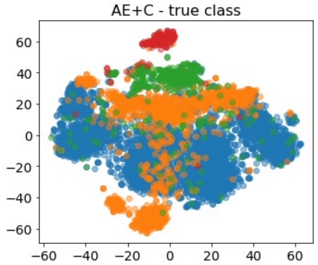

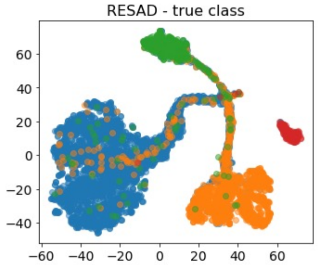

Latent space configuration: the superiority of the CCLP approach

National Aeronautics and Space Administration

Figures show 2D visualization of the 256D latent space of each model using t-Distributed Stochastic Neighbor Embedding (t-SNE), color-coded based on the actual class of the data.

Second column shows the results of K-Means clustering applied to the 256D latent space dividing the space into $n_c + 1$ clusters.

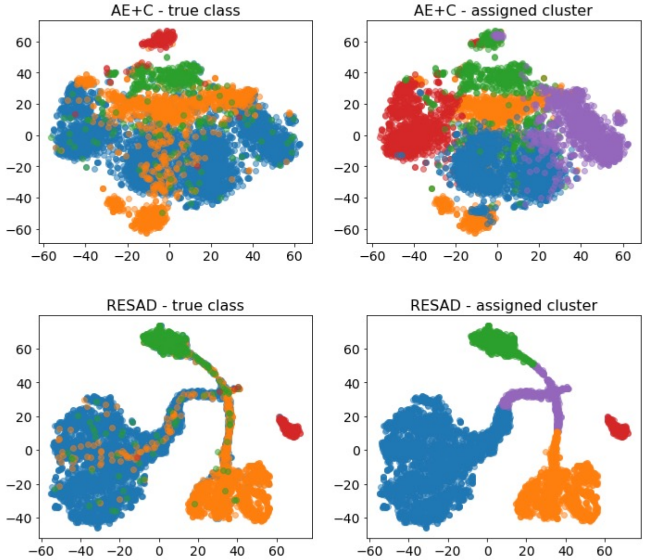

Latent space configuration: the superiority of the CCLP approach

National Aeronautics and Space Administration

We evaluate the relationship between clusters shaped in the latent space and the prediction uncertainty of the classifier. These results suggest a novel active learning strategy for selecting the most informative data to be labeled in future efforts.

CVAE – unsupervised encoding

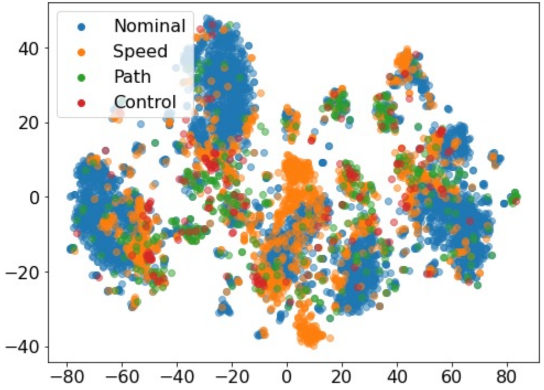

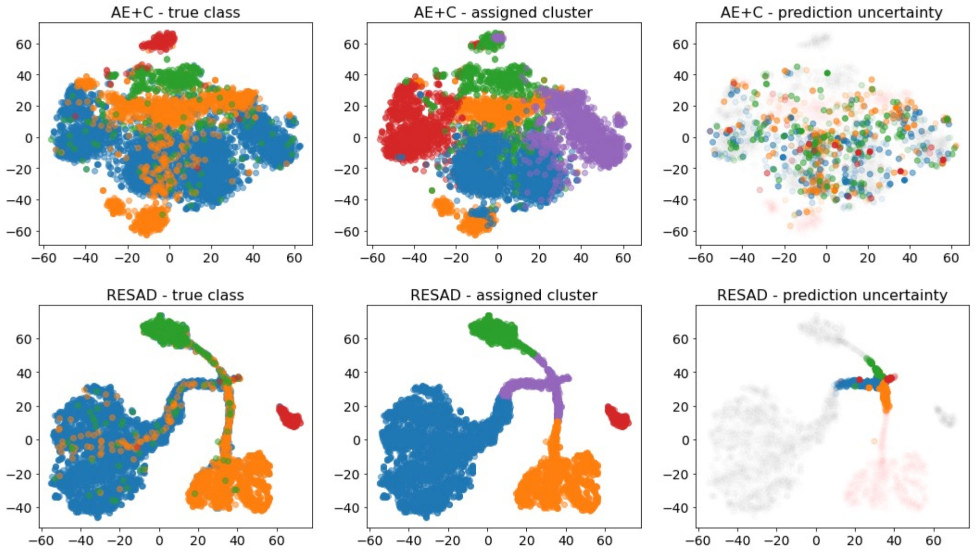

Acknowledgement and references

We acknowledge funding of this research from the NASA System-Wide Safety Project under contracts 80ARC020D0010 and NNA16BD14C.

References:

Memarzadeh, M., Matthews, B., and Templin, T., ”Multi-Class Anomaly Detection in Flight Data Using Semi-Supervised Explainable Deep Learning Model”, Journal of Aerospace Information Systems, 19(2), 2022.

Memarzadeh, M., Akbari Asanjan, A., and Matthews, B., “Robust and Explainable Semi-Supervised Deep Learning Model for Anomaly Detection in Aviation”, Aerospace, 9(8), 437, 2022.

Memarzadeh, M., Matthews, B., and Avrekh, I., “Unsupervised Anomaly Detection in Flight Data Using Convolutional Variational Auto-Encoder,” Aerospace, Vol. 7, 2020, p. 115.

Janakiraman, V. M., “Explaining Aviation Safety Incidents Using Deep Temporal Multiple Instance Learning,” KDD ’18: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2018, pp. 406–415.

Kingma, D., Rezende, D., Mohamed, S., and Welling, M., “Semi-supervised learning with deep generative models,” Advances in Neural Information Processing Systems (NeurIPS), 2014. URL https://arxiv.org/abs/1406.5298.

Kamnitsas, K., Castro, D., Le-Folgoc, L., Walker, I., Tanno, R., Rueckert, D., Glocker, B., Criminisi, A., and Nori, A., “Semi-supervised learning via compact latent space clustering,” Proceedings of the 35th International Conference on Machine Learning (ICML), 2018, pp. 2459–2468. URL https://arxiv.org/abs/1406.5298.

Houlsby, N. Huszar, F., Ghahramani, Z., and Lengyel, M., “Bayesian active learning for classification and preference learning.”, 2011. URL https://arxiv.org/abs/1112.5745

Gal, Y., and Ghahramani, Z., “Dropout as a Bayesian approximation: representing model uncertainty in deep learning.”, *ICML*, 48, 2016. pp 1050-1059.

Gal, Y., Islam, R., and Ghahramani, Z., “Deep Bayesian active learning with image data”, *ICML*, 70, 2017, pp. 1183-1192