ATTRACTOR Special Session

Toward Justifiable Trust in Autonomous Systems Incorporating Human Knowledge in Autonomous Systems through Machine Learning

Nikunj C. Oza, Ph.D.

Kamalika Das, Ph.D., David Iverson, Vijay Janakiraman, Ph.D.

nikunj.c.oza@nasa.gov

Trust in Autonomous Systems

- Humans must trust the decisions made by autonomous systems
- This trust can be increased through learning from domain experts
- Offline: From past mission operations' data
- Online: Obtain human feedback during operations

Machine Learning from Humans

- Offline: Inverse Reinforcement Learning
  - Past mission operations -> Mission planner/controller
  - Automatic Discovery of Precursors in Time series data (ADOPT)
- Offline/Online: Anomaly Detection
  - Offline: Past mission operations -> Statistically anomalous operations
  - Online: Current operations data -> anomalous/normal
- Offline/Online: Active Learning
  - Offline: Statistical Anomalies -> Operationally significant/Not
  - Online: Tie-breaker actions

Inverse Reinforcement Learning

- Reinforcement Learning: Given rewards of all states, learn a controller that maximizes expected cumulative reward (discounted sum of rewards).
- Problem: Need rewards of *all* states, or heuristics. Relative rewards between states difficult to determine.
- Inverse Reinforcement Learning (a.k.a., Apprenticeship Learning): Given examples of runs of a controller, learn the rewards of the states---assume that most runs represent optimal operations
- Has been demonstrated on single-agent systems (e.g., RC helicopters [Ng, et. al., 2004] and others) and multi-agent systems (e.g., traffic routing [Natarajan, et. al., 2010])

IRL for ATTRACTOR

- Set rewards where known
- Learn from example trajectories, standard search and rescue configurations
- Single-agent and multi-agent

ADOPT

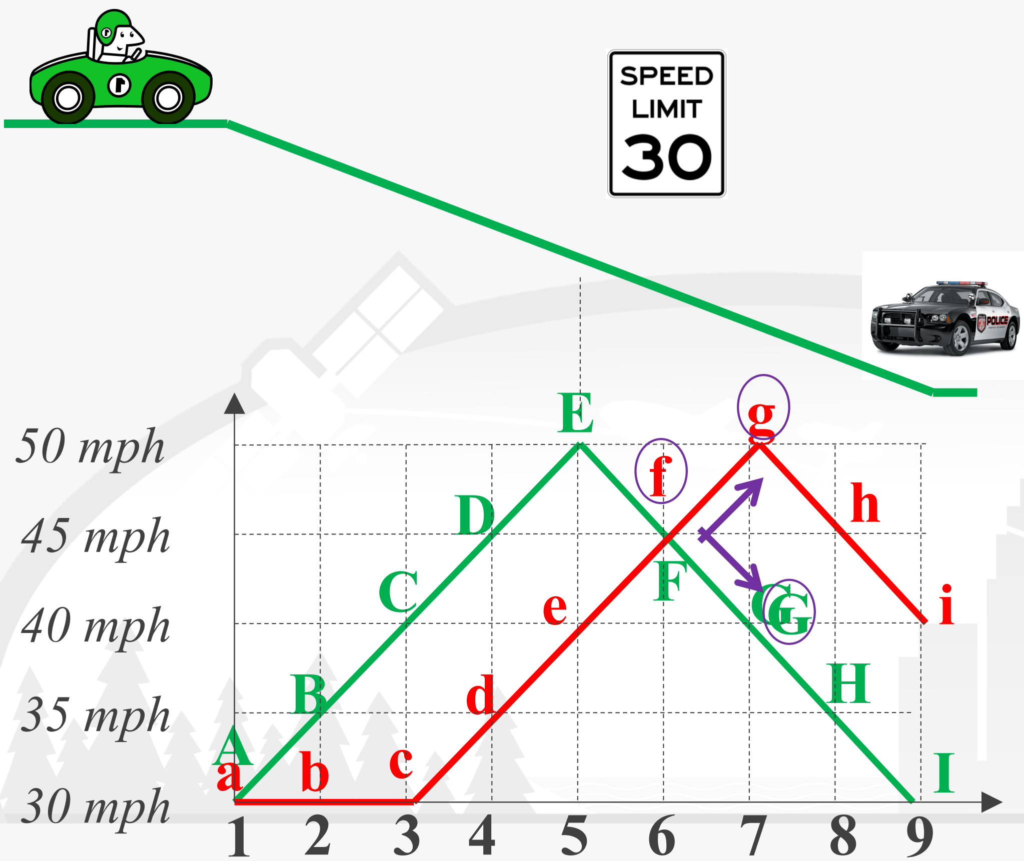

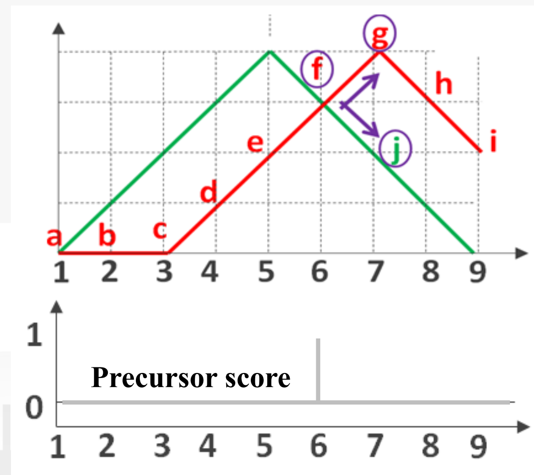

$$V(a_k^*) - V(a_k) > \delta.$$

ADOPT

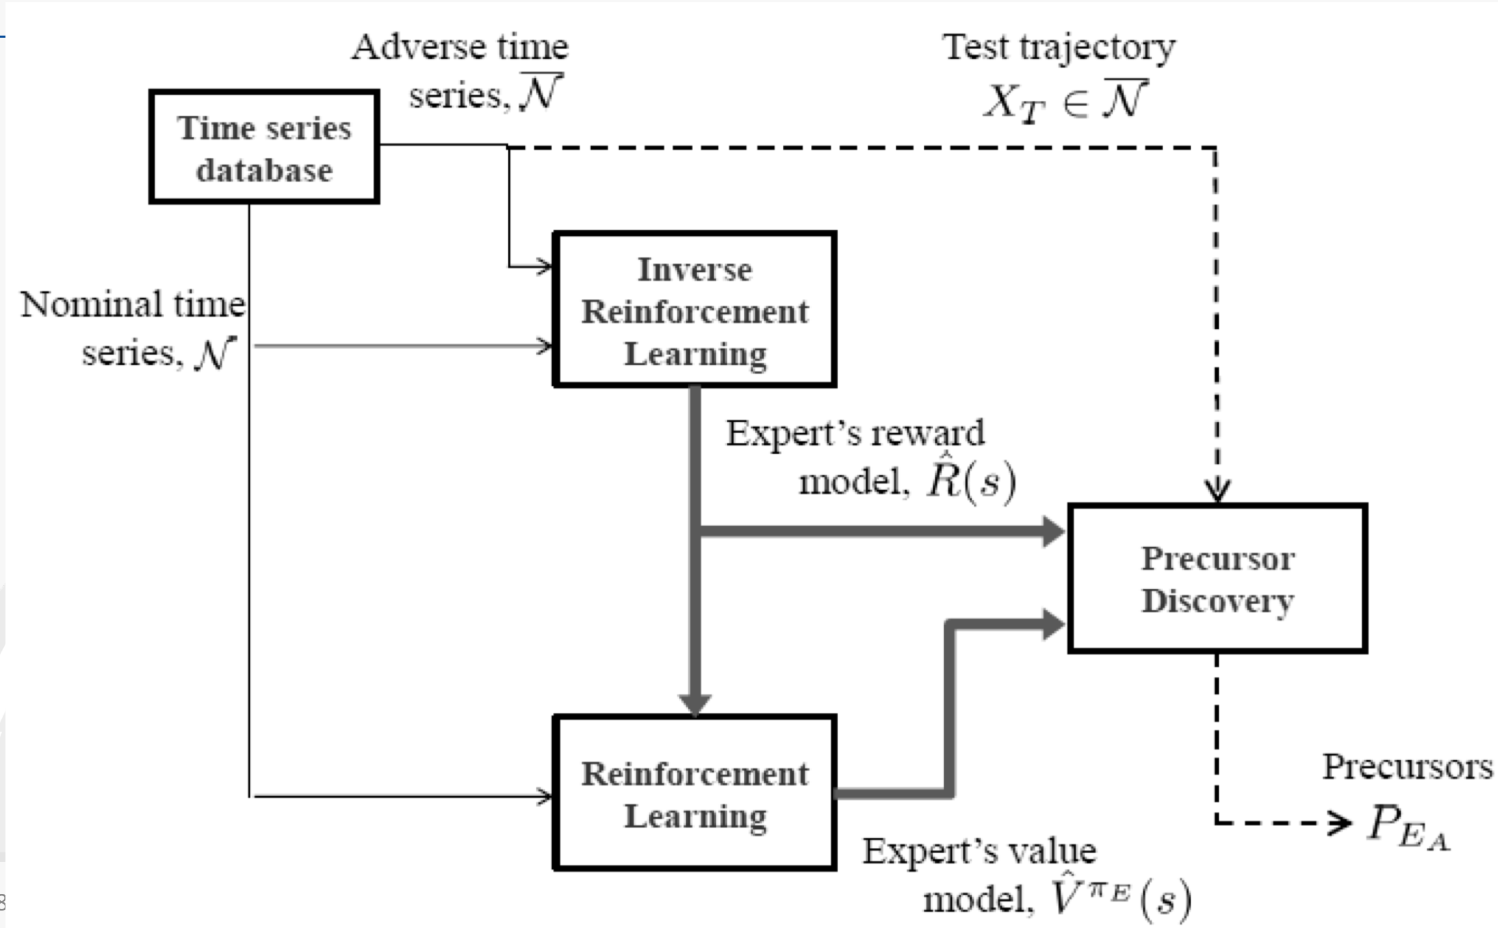

Adverse time series, $\overline{\mathcal{N}}$
Test trajectory $X_T \in \overline{\mathcal{N}}$
Time series database
Inverse Reinforcement Learning
Nominal time series, $\mathcal{N}$
Expert's reward model, $\hat{R}(s)$
Precursor Discovery
Reinforcement Learning
Precursors P_{E_A}
Expert's value model, $\hat{V}^{\pi_E}(s)$

Precursor to go-around events

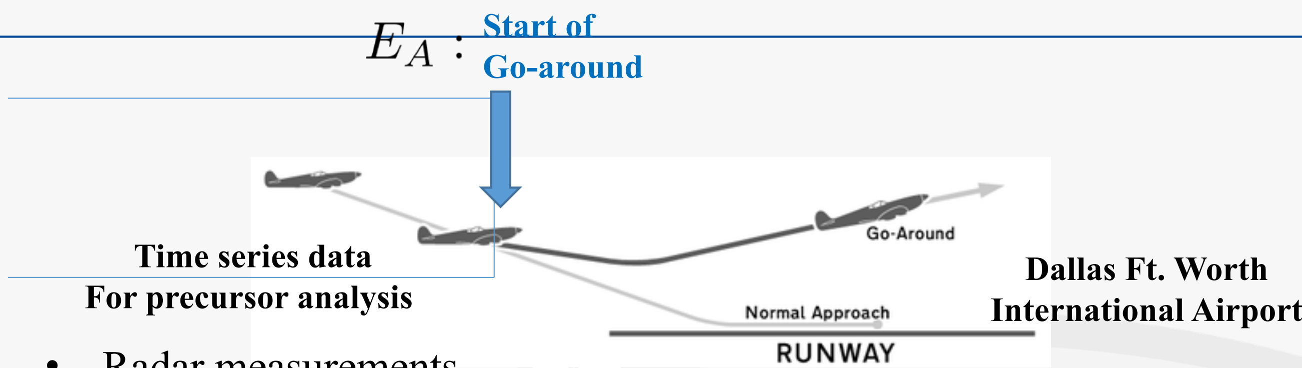

- Radar measurements
  - Latitude, longitude, altitude, ground speed.
- landing airport data
  - Runway configuration, counts and rates of departure and arrival, total air and taxi delays, meteorological data.
- Derived features
  - Headwind components, altitude in feet above ground level (AGL), horizontal/vertical distance between aircraft, and the corresponding closing rates to the nearest aircraft.
- Data dimensions
  - **41 time series variables, 250 time steps, 1000 flights.**

Precursor: Energy mismanagement

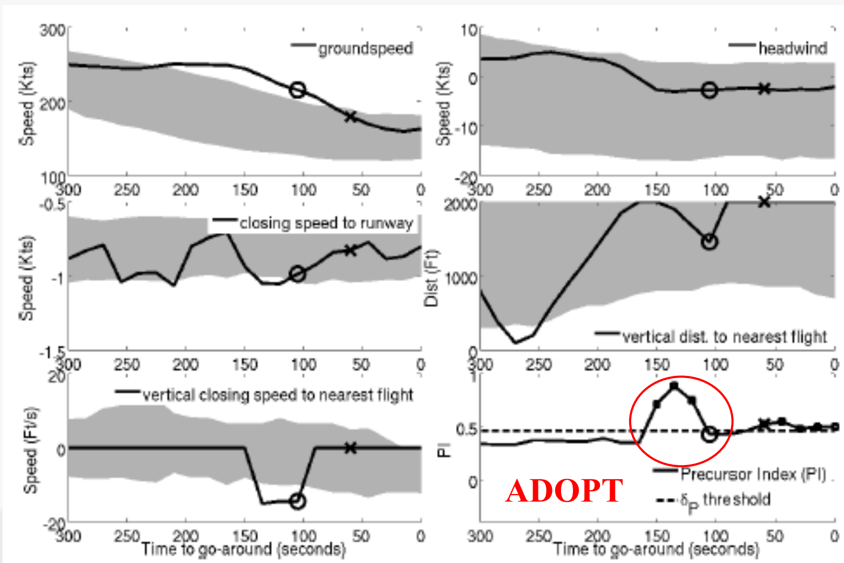

- **Some flight variables plotted against time to go-around.**
- **Shaded area – nominal distribution**
- **O is the outer-marker (5 miles to TD)**
- **X is when flight is at 1000 ft altitude**

Abstracted actions

- Expert actions (pilot switches, control commands) are not observed in data.
- Actions are abstracted based on state transitions.

3.1.1 Energy Mis-Management Scenario The flight in this example (see Figure 2) executed a go-around because of high speed prior to landing. It can be seen that between 160 and 130 seconds and around 50 seconds prior to the go-around, the PI value is increasing significantly and at the same time there is a significantly high ground speed above the shaded band (nominal data distribution) indicating a possible precursor to the go-around. The SME confirmed that the flight's high speed around the two markers (outer marker and 1000 ft altitude marker) could be a main reason behind the go-around. ADOPT was able to identify these points and even time instants prior to these as precursors.

ADOPT for ATTRACTOR

- Identify states that represent mission failure
- Identify precursors to them
- Need to abstract to mission-level and vehicle-level actions

Machine Learning from Humans

- Offline: Inverse Reinforcement Learning
  - Past mission operations -> Mission planner/controller
  - Automatic Discovery of Precursors in Time series data (ADOPT)
- Offline/Online: Anomaly Detection
  - Offline: Past mission operations -> Statistically anomalous operations
  - Online: Current operations data -> anomalous/normal
- Offline/Online: Active Learning
  - Offline: Statistical Anomalies -> Operationally significant/Not
  - Online: Tie-breaker actions

Anomaly Detection

- Training: Past mission operations -> Statistically anomalous operations
- Testing: Current operations data -> anomalous/normal
- Example algorithm: Multiple Kernel Anomaly Detection (MKAD) [Das, et. al. 2010]
  - Works with heterogeneous data
    - FOQA data: discrete and continuous data
    - Radar-track data: lat/lon, altitude, distance to nearest aircraft
  - Kernel functions (measures of similarity) for each modality
  - Same underlying optimization code

MKAD

Features from multiple data sources

Flight instances	f_1	f_2	…	…	f_5	…				…	f_n
1											
2											
3											
…											
…											
…											
m											

One class SVM objective function:

$$\min \quad Q = \frac{1}{2}\sum_{i,j} \alpha_i \alpha_j K(\mathbf{x}_i, \mathbf{x}_j)$$

$$\text{subject to} \quad 0 \le \alpha_i \le \frac{1}{\ell\nu}, \sum_i \alpha_i = 1, \rho \ge 0, \quad \nu \in [0,1]$$

MKAD Decision function:

$$f(\mathbf{z}) = sign\left(\sum_i \alpha_i \sum_p \eta_p K_p(\mathbf{x}_i, \mathbf{z}) - \rho\right)$$

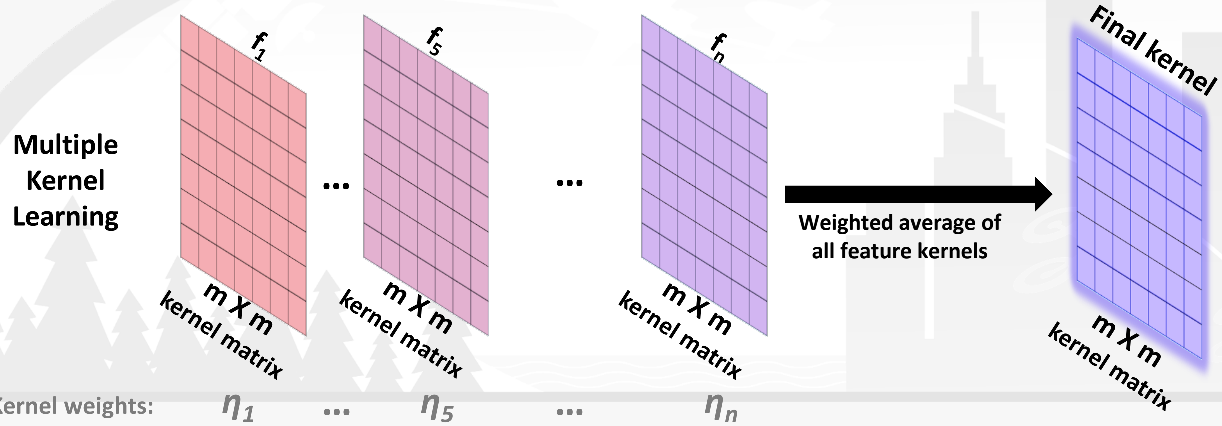

Anomaly Detection

- Other Algorithms
  - SequenceMiner: Finds anomalies in discrete sequences (e.g, aircraft mode sequences)
  - Multivariate Time Series (MTS) search: Quickly search for multivariate motifs in a large time series archive.
  - iOrca: Distance-based anomaly detection on continuous data.
  - Nu-Anomica and Bi-Criterion: More efficient versions of one-class SVM for anomaly detection.
- Online Monitoring Tools
  - Prototype daily and hourly anomaly reporting systems

Example trajectory anomaly

- Overshoots Extended Runway Centerline (ERC)
  by over 1 SM
- Over 250 Kts @2500 Ft.
- Angle of intercept > 40°
- Overshoots 2nd approach

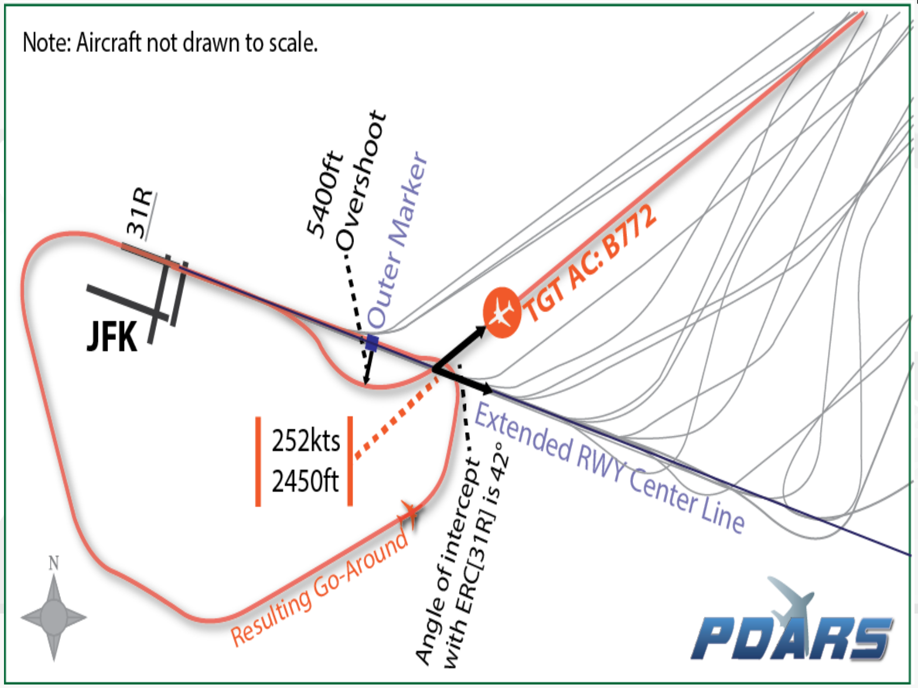

Anomaly Detection for ATTRACTOR

- Identify anomalous missions and parts of missions
- Identify anomalies at mission and vehicle level, interplay between them

Machine Learning from Humans

- Offline: Inverse Reinforcement Learning
  - Past mission operations -> Mission planner/controller
  - Automatic Discovery of Precursors in Time series data (ADOPT)
- Offline/Online: Anomaly Detection
  - Offline: Past mission operations -> Statistically anomalous operations
  - Online: Current operations data -> anomalous/normal
- Offline/Online: Active Learning
  - Offline: Statistical Anomalies -> Operationally significant/Not
  - Online: Tie-breaker actions

Active Learning

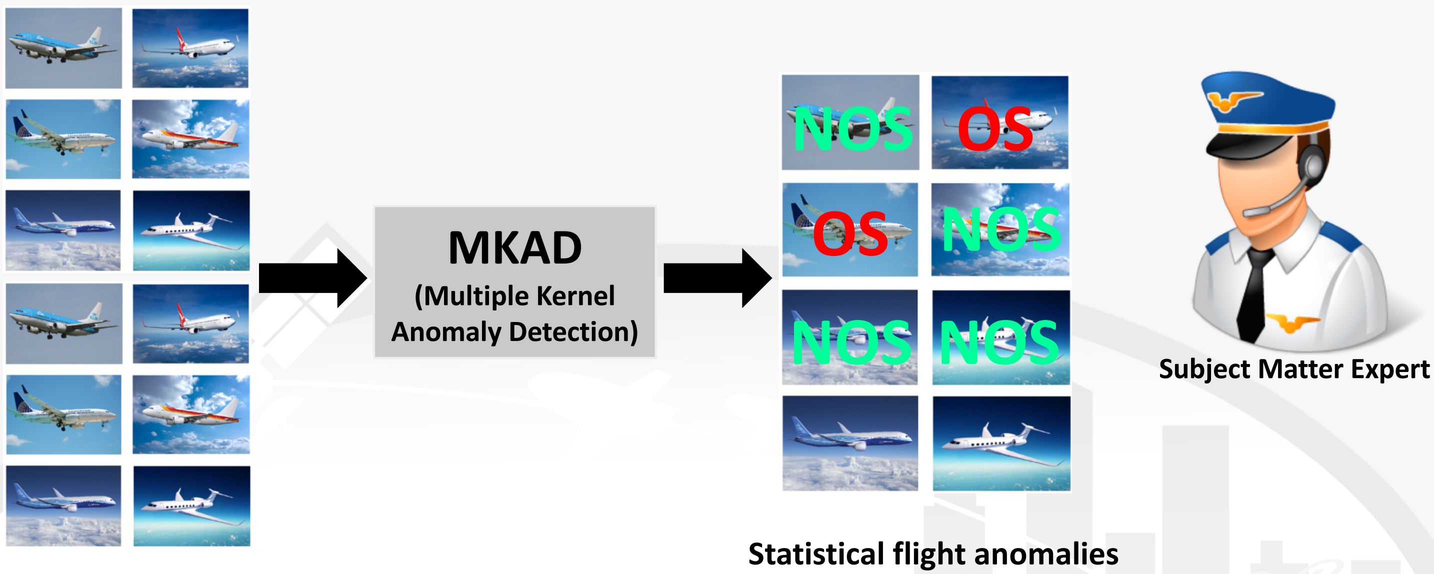

We want to **effectively** train a model to automatically identify operationally significant (OS) anomalies

Active Learning Framework

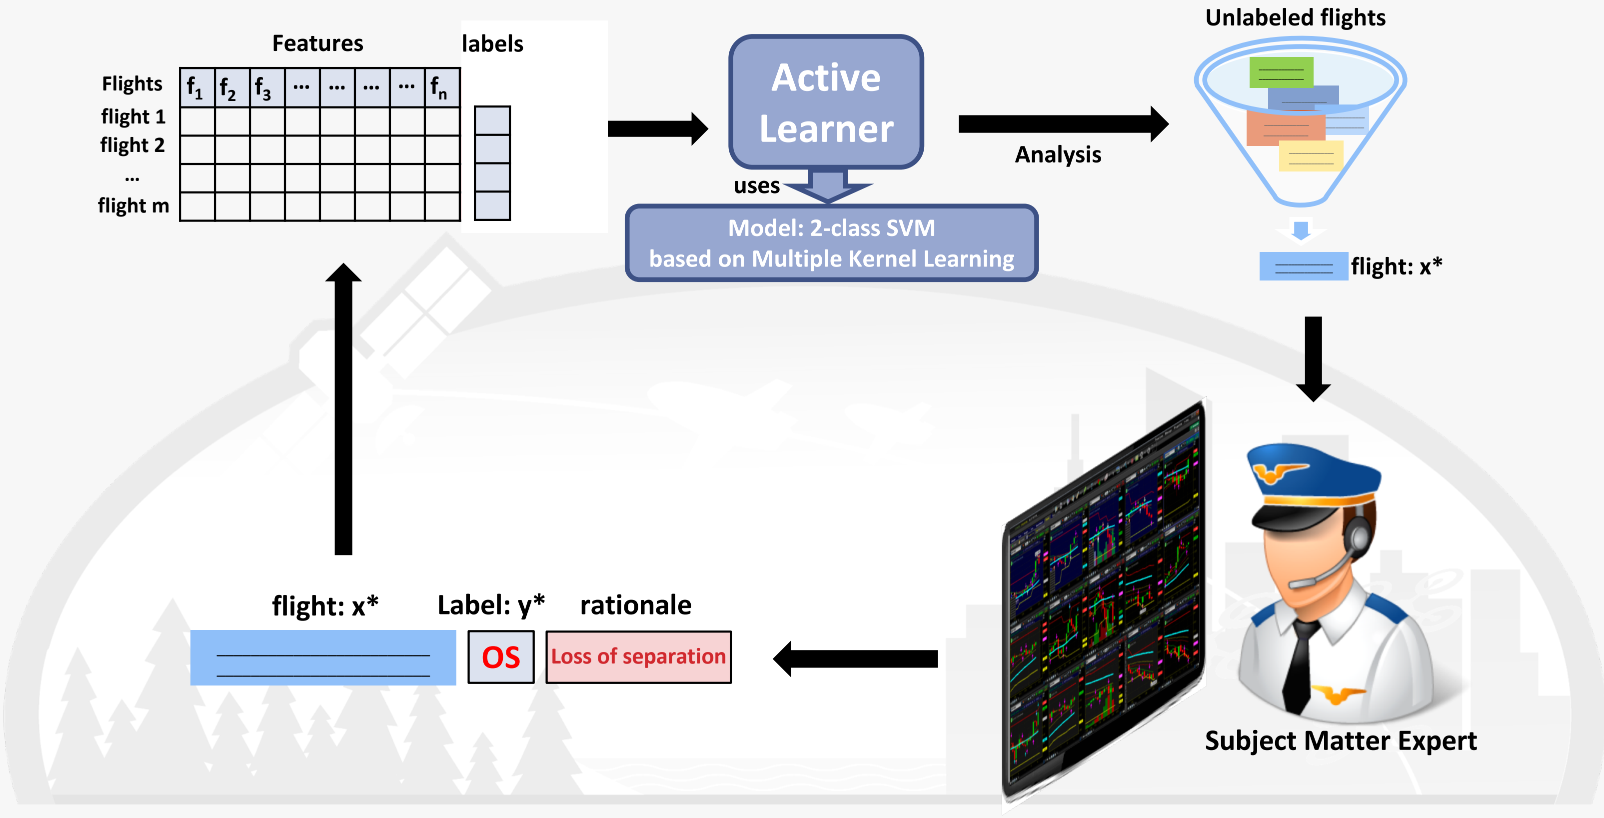

Rationale Features

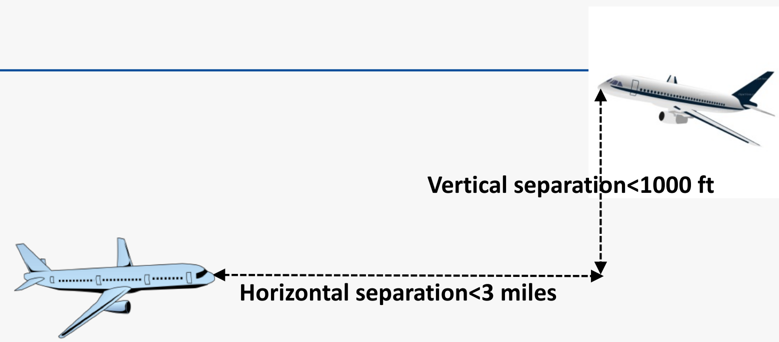

"Loss of separation"

- Horizontal separation < 3 miles AND Vertical separation < 1000 ft AND nearest neighboring flight is not on parallel runways and not part of the same flow

"Large overshoot"

- Maximum overshoot is greater than a threshold based on values of flights with positive labels

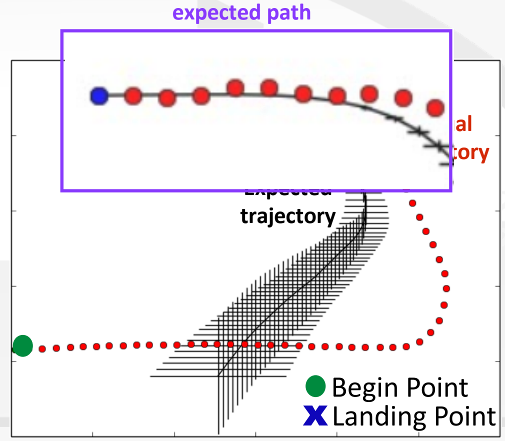

"Unusual flight path"

- Overall deviation from expected (average) trajectory of all landing flights on that runway

Active Learning for ATTRACTOR

- Label anomalies as operationally significant/not
- Choose among equal/near-equal valued actions

Future Work

- Humans must trust the decisions made by autonomous systems
- This trust can be increased through learning from domain experts
- Offline: From past mission operations’ data
  - **Simultaneously learning at mission (multi-vehicle) and vehicle level**
  - **Simultaneously learning actions and anomaly/precursor significances**
- Online: Obtain human feedback during operations
  - **Real-time/in-time versions of above**

Thank you!

Contact: **nikunj.c.oza@nasa.gov**

Papers

K. Das, I. Avrekh, B. Matthews, M. Sharma, and N. Oza, ASK-the-Expert: Active Learning Based Knowledge Discovery Using the Expert, *ECML-PKDD* 2017.

M. Sharma, K. Das, M. Bilgic, B. Matthews, D. Nielsen, and N. Oza, Active Learning with Rationales for Identifying Operationally Significant Anomalies in Aviation, *ECML-PKDD*, 2016.

V. Janakiraman, B. Matthews, and N. Oza, Discovery of Precursors to Adverse Events Using Time Series Data, *SDM* 2016.

B. Matthews, D. Nielsen, J. Schade, K. Chan, and M. Kiniry, Comparative Study of Metroplex Airspace and Procedures Using Machine Learning to Discover Flight Track Anomalies, 34^{th} *DASC*, 2015.

S. Das, B. Matthews, N. Oza, and A. Srivastava, Multiple Kernel Learning for Heterogeneous Anomaly Detection: Algorithm and Aviation Safety Case Study, *KDD* 2010.