Updates on Research and Collaborations

Matthew Hale
Department of Mechanical and Aerospace Engineering
University of Florida

Center of Excellence for Assured Autonomy in Contested Environments
Spring 2023 Review
April 26th, 2023

Collaborations with CoE PIs

- Ongoing collaboration with Dawn Hustig-Schultz and Ricardo Sanfelice (UCSC)
  - Developed a distributed asynchronous heavy ball algorithm
  - Paper to be presented at ACC 2023 in San Diego next month
  - A journal extension is in the works

- Ongoing collaboration with Parham Gohari, Mustafa Karabag, Cyrus Neary, and Ufuk Topcu (UT-Austin)
  - CDC 2023 paper on private RL (under review)
  - CDC 2023 paper on private stochastic matrices (under review)
  - Another paper under review at UAI for privacy in multi-agent planning
  - Various extensions are in the works

Collaborations with Air Force Colleagues

- Applied optimization to weapon-target assignment (WTA) problems with RW:
  Katherine Hendrickson, Prashant Ganesh, Kyle Volle, Paul Buzaud, Kevin Brink, and Matthew Hale, "Decentralized Weapon–Target Assignment Under Asynchronous Communications", *Journal of Guidance, Control, and Dynamics,* 2023, 46:2, 312-324.
- Work on anomaly detection with RY:
  M. Ubl, B. Robinson, and M.T. Hale, "Anomaly search over many sequences with switching costs," Under review at *Control Systems Letters*
- 2023 AIAA SciTech paper with Michael Anderson and cadets from USAFA
- Work on optimization in the loop with RV:
  G. Behrendt, A. Soderlund, M. Hale, and S. Phillips, "Autonomous satellite rendezvous and proximity operations with time constrained sub-optimal MPC," Accepted to 22nd IFAC World Congress, 2023.
- Engaging with AFRL every summer
  - William Warke was a Summer Scholar in 2018, 2019, 2022 at RW with Kevin Brink
  - I was a Summer Faculty Fellow in 2020 at RW with Kaitlin Fair/Kevin Brink
  - Matthew Ubl was a Summer Scholar in 2021 at RY with Ben Robinson
  - Gabriel Behrendt was a Summer Scholar in 2022 at RV with Sean Phillips
  - Alexander Benvenuti was a Summer Scholar in 2022 at RW with Scott Nivison
- **For summer 2023:**
  - William Warke will be at RW with Kevin Brink
  - Gabriel Behrendt will be at RW with Zach Bell
  - Alexander Benvenuti will be at RW with Brendan Bialy
  - Calvin Hawkins will be at RY with Ben Robinson

Herbert Wertheim College of Engineering
UNIVERSITY of FLORIDA

Mechanical and Aerospace Engineering

Anomaly Search Over Many Sequences With Switching Costs

Matthew Ubl, Ben Robinson, Matthew Hale

Center of Excellence for Assured Autonomy in Contested Environments Spring 2023 Review

26 April 2023

Anomaly Search Over Multiple Sequences

- **Anomaly Search:** Identifying data that deviates from the norm among many options
  - Cognitive Radio (open frequencies)
  - Clinical Trials (successful treatment plans)
  - Intrusion Detection (compromised system)
  - Area Surveillance (changes of interest)

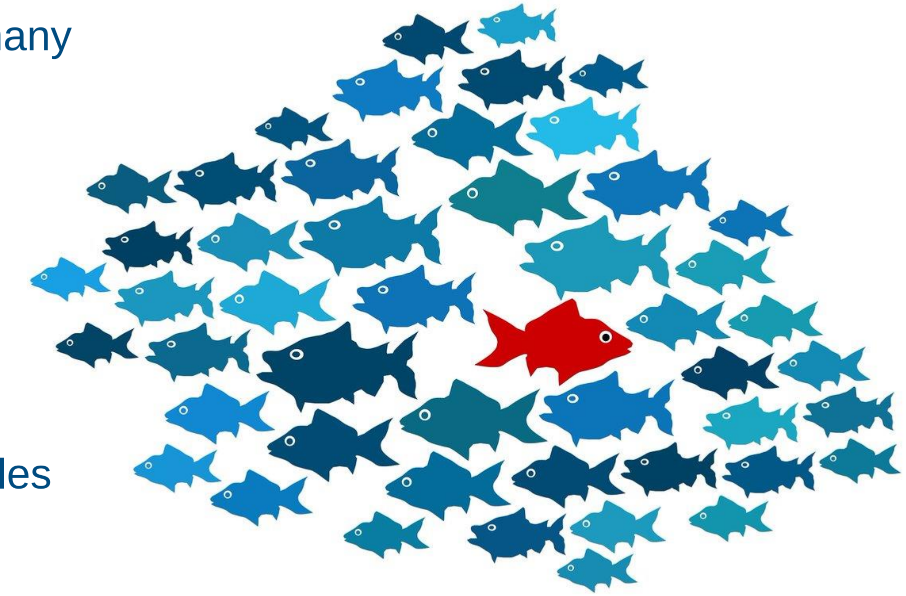

- Identifying anomalies accurately often requires many observations/samples
  - i.e., lots of information

Goal: Given multiple choices we can observe (data streams), how do we identify anomalous ones as quickly as possible, while still being accurate?

Want to Identify Anomalies Quickly and Accurately

Problem:

Given

- $k \leq \infty$ data streams,
- with anomalous streams occurring with prior probability $\hat{\pi}$,

Develop an algorithm to identify an anomalous data stream

- in as few observations (τ) as possible,
- with a false-identification error probability less than ϵ. That is,

$$\text{minimize } \mathbb{E}\{\tau\}$$
$$\text{such that } \mathbb{P}\left(H^{k_\tau} = H_0\right) \leq \epsilon$$

Existing State-of-the-Art Algorithm Uses One Threshold Parameter

- The optimal algorithm, based on the Sequential Probability Ratio Test, is known (Lai, 2011)
- Assume observations of nominal data streams follow distribution f_0, anomalous ones follow f_1
- The Log-Likelihood Ratio of an observation X_t^k is $\ell(X_t^k) = \log\left(\frac{f_1(X_t^k)}{f_0(X_t^k)}\right)$
- Add observations up to $\Lambda_T^k = \sum_{t=0}^{T} \ell(X_t^k)$
  - If Λ_T^k crosses a lower threshold $\gamma_L < 0$, declare it nominal, switch to next data stream
  - If Λ_T^k crosses an upper threshold $\gamma_U > 0$, **RED FLAG**, declare it anomalous

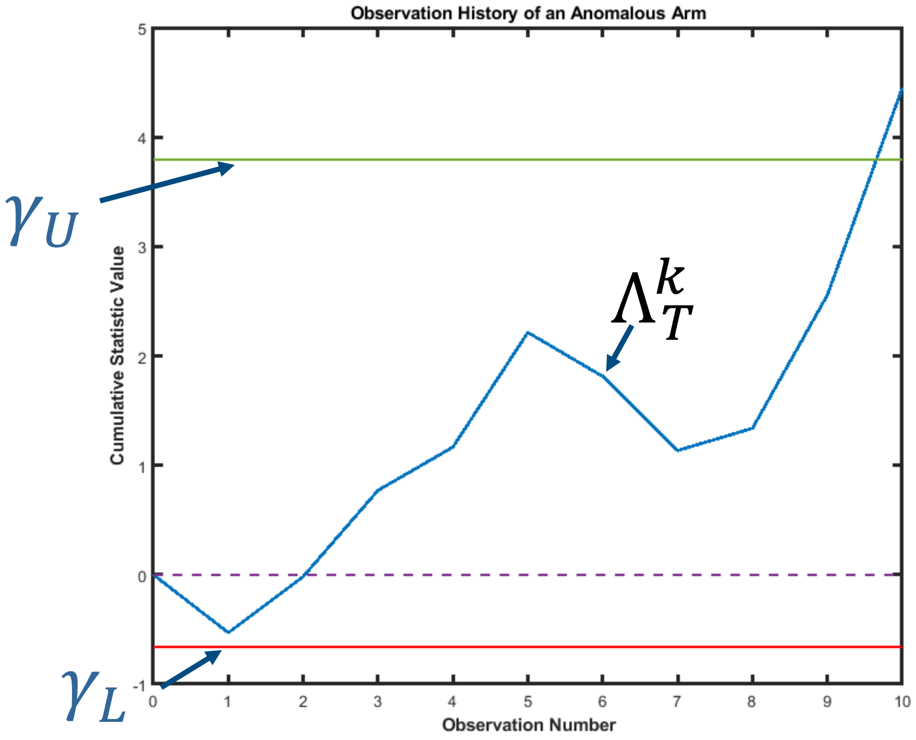

Optimal Thresholds (Lai, 2011)**:** The optimal thresholds of this algorithm are $\gamma_L = 0$ and $\gamma_U = \gamma_U^*$

Existing Algorithms Don't Handle Switching Costs Well

- **Issue #1:** We know that an optimal upper threshold γ_U^* exists, but it cannot be directly calculated
- **Issue #2:** This algorithm also does not consider the case of *switching costs*
- **Switching Costs:** When the observer switches from data stream k to data stream $k+1$, it incurs some switching cost λ
  - Dead-time when no useful observations can be taken
- Some algorithms exist that address switching costs, none perform very well
  - Separate exploration/exploitation steps
  - Pre-scheduling when we can switch (block scheduling)

Goal: Can we develop an algorithm to solve the anomaly detection problem efficiently while minimizing switching costs?

Want to Identify Anomalies Quickly (Including Switches) and Accurately

New Problem:

Given

- $k \leq \infty$ data streams,
- with anomalous streams occurring with prior probability $\hat{\pi}$,

Develop an algorithm to identify an anomalous data stream

- in as little time as possible,
- with a false-identification error probability less than ϵ. That is,

$$\text{minimize } \mathbb{E}\{\tau\} + \lambda \mathbb{E}\{S\}$$
$$\text{such that } \mathbb{P}\left(H^{k_\tau} = H_0\right) \leq \epsilon$$

We Can Use Approximately Optimal Thresholds

- What do we mean when we say we want to minimize switches?
  - We want to reduce the probability that we switch away from an anomalous data stream, i.e., our *false negative error*

- γ_U governs when we declare a data stream as anomalous, γ_L governs when we switch to a new one
  - We reduce the number of switches by tuning γ_L

Approximately Optimal Thresholds for Switching Costs (Ubl, Robinson, Hale, 2023): The approximately optimal thresholds for anomaly detection with switching costs are:

$$\gamma_U^* = \log\left(\frac{1-\hat{\pi}}{\hat{\pi}}\frac{1-\epsilon}{\epsilon}\right)$$

$$\gamma_L^* = \underset{\gamma_L \leq 0}{\arg\min}\, C(\cdot\,, \epsilon, \hat{\pi}, f_0, f_1) \longrightarrow C: 1-\text{Dimensional, strongly convex}$$

These thresholds approach optimality when $f_0 \approx f_1$, $\epsilon \ll 1$, and $\lambda \gg 0$

Simulation Results

- $\hat{\pi} = 0.1, \epsilon = 0.01, F_0 = \mathcal{N}(0,1), F_1 = \mathcal{N}(0,1.5)$

- Compared the algorithm from (Lai, 2011) to ours (Ubl, 2023) in terms of expected observation-switching cost

$$\mathbb{E}\{\tau\} + \lambda \mathbb{E}\{S\}$$

as switching cost λ grows

- Comparable performance for small λ, our algorithm quickly achieves better performance as λ grows large

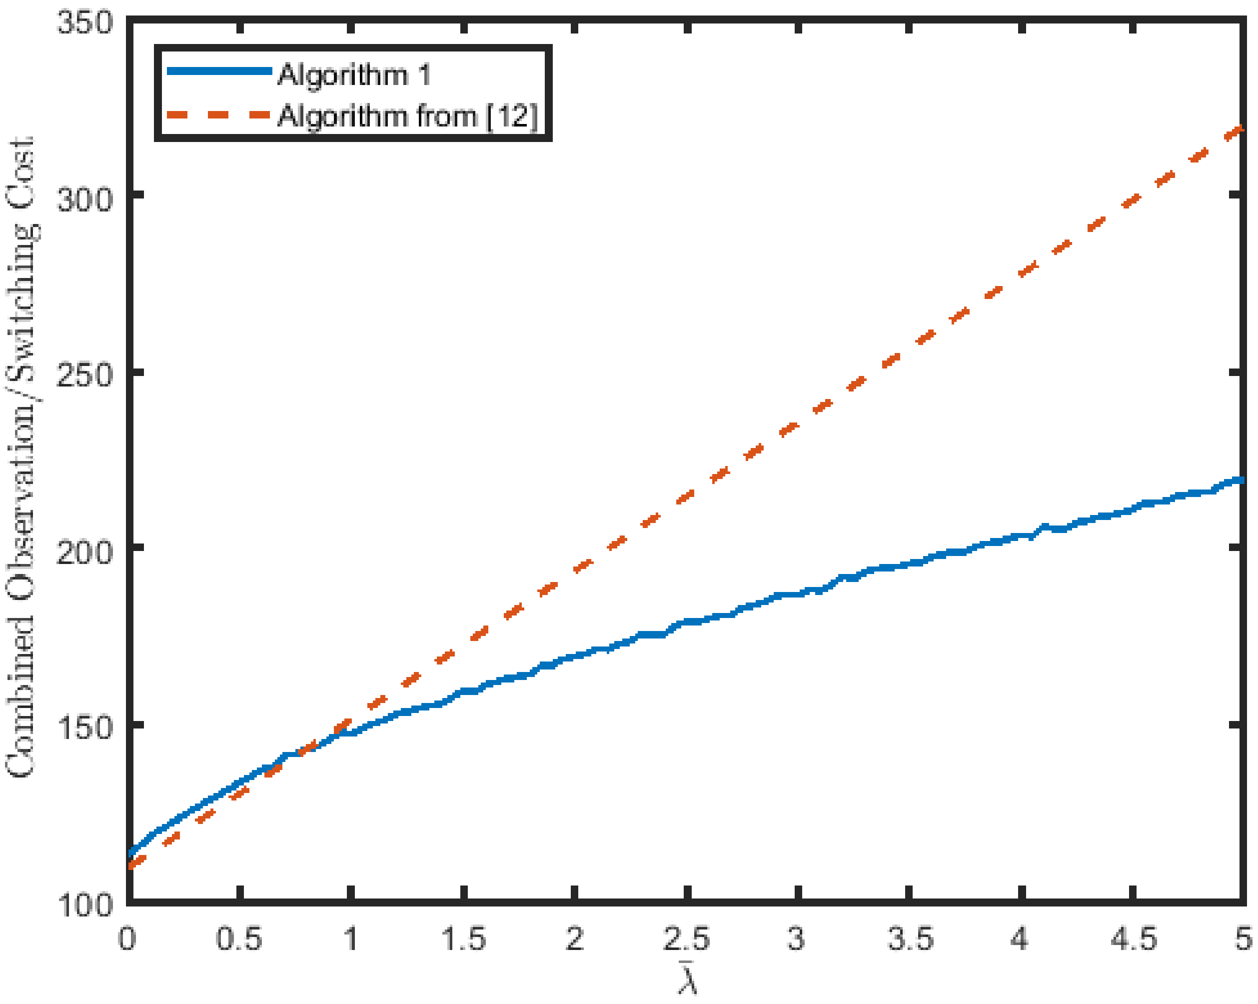

Thank you