Herbert Wertheim College of Engineering
Florida Institute for National Security (FINS)
UNIVERSITY *of* FLORIDA

LYAPUNOV-STABLE ADAPTIVE CONTROL FOR MULTIMODAL CONCEPT DRIFT

Tianyu Bell Pan, Mengdi Zhu, Alexa Cole, Ronald Wilson, Damon Woodard
Department of Electrical and Computer Engineering
Florida Institute of National Security
Applied Artificial Intelligence Group

12/3/2025
The 39th Annual Conference on
Neural Information Processing Systems (NeurIPS)

Florida Institute for National Security (FINS)

Introduction

This paper treats multimodal learning as a controlled dynamical system and designs a provably stable online controller to adapt safely under drift.

- **Goal**
  - We study how to keep multimodal models reliable over time when the data distribution drifts and different modalities degrade at different rates.

- **Motivation**
  - Modern multimodal systems are deployed in the wild.
  - Over time, content style, quality, and attacks change → ***Concept Drift***
  - Often, one modality becomes unreliable (e.g., noisy images, manipulated text) while others remain useful.
  - *Standard training or naïve fine-tuning has no guarantees:*
    - Performance may collapse or oscillate

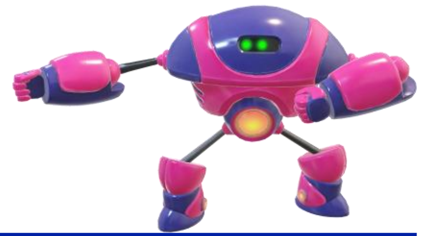

Core Idea: Learning as Control

- Wrap the existing model with a lightweight controller: **LS-OGD**.
- **Base model:**
  - Two modalities, late fusion: $z_t = \alpha_t z_t^{(1)} + (1 - \alpha_t) z_t^{(2)}$
- **At each time step or mini-batch, the controller:**
  - Observes recent performance
  - Detects evidence of drift
  - *Adjust:* The learning rate η_t and the fusion weight α_t
- **Aim to make the closed-loop system (model + modality) Lyapunov-stable:**
  - Error do not blow up
  - The system tracks changes
  - It recovers when drift stops
- Instead of blindly training, we steer the learner using **control-theoretic** principles.

How LS-OGD Works

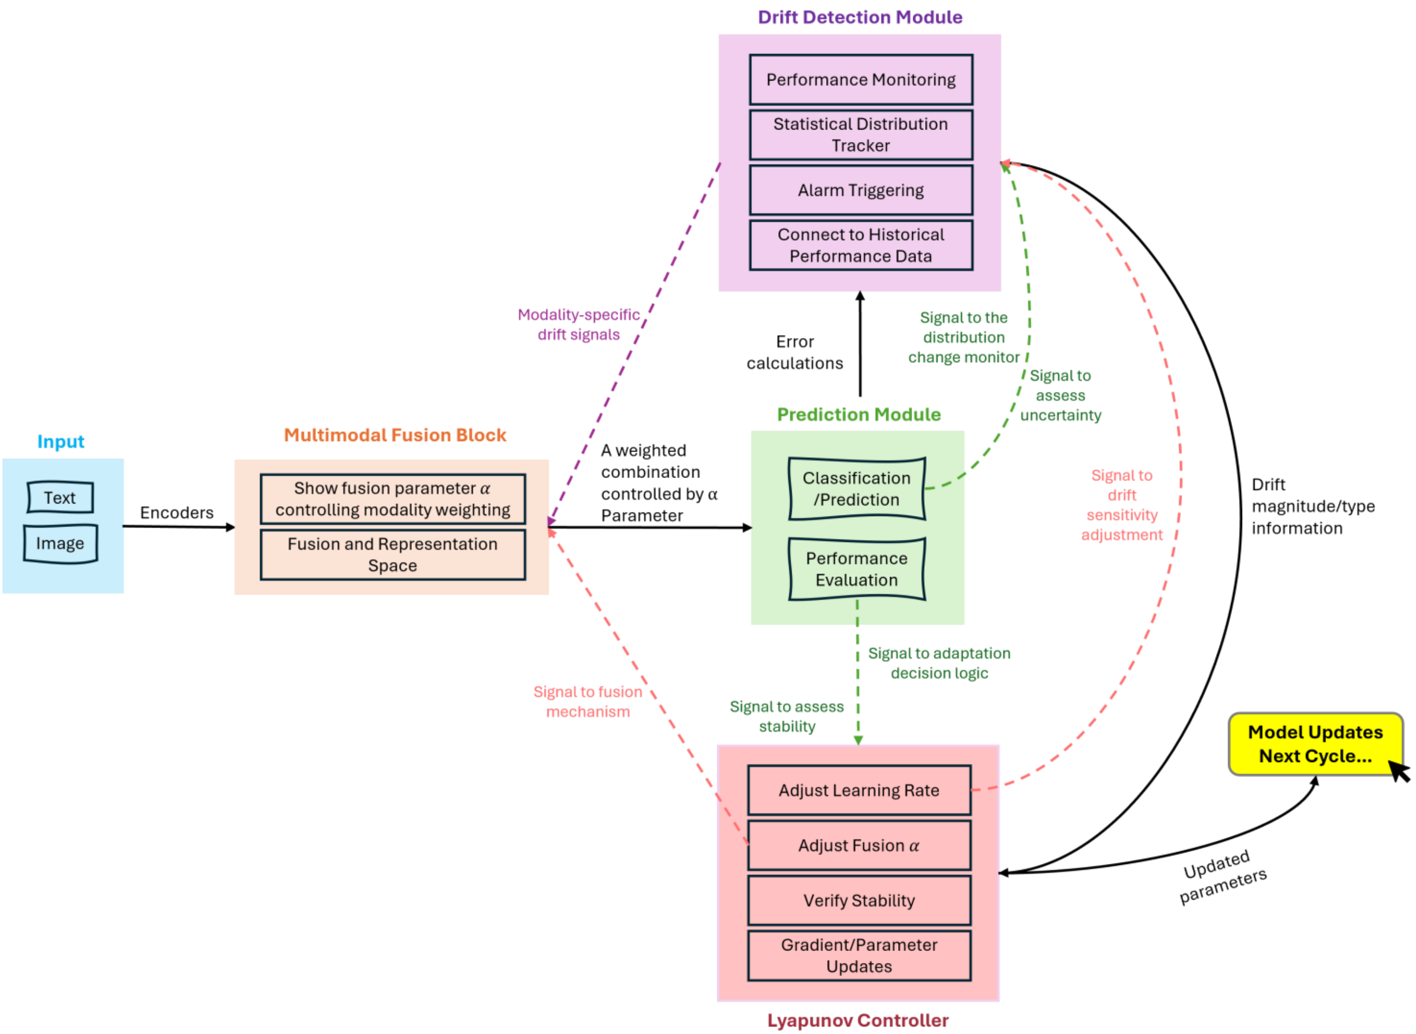

Jointly adapts the learning rate and multimodal-fusion weight

Maintain bounded prediction error under bounded concept drift

To converge to zero error once drift stops

Theory: What Do We Guarantee?

- **Lyapunov function**
  - Use $V(t) = \frac{1}{2} e_t^2$ as an energy of the prediction error.
- **Main result – Uniform Ultimate Boundedness (UUB)**
  - Under bounded drift and mild regularity assumptions.
  - LS-OGD ensures the prediction error is UUB: *It enters a small region and stays there.*
  - If drift eventually stops → Error converges to zero as in a stable system.
- **Fault-tolerant modality adaptation**
  - If one modality becomes persistently unreliable.
  - The controller automatically down-weights it ($\alpha \to 0$ or 1).
  - The system converges to the performance of the best remaining modality without losing stability
- **We don't just heuristically adapt, but we prove the adaptation is safe and effective**

Takeaways

- **Conceptual**
  - Bridges control theory and multimodal machine learning for streaming, non-stationary environments.
- **Practical**
  - *A plug-in controller:*
    - Minimal overhead.
    - Compatible with existing models.
    - Focuses on reliability and fault tolerance.
- **Impact**
  - *Relevant for any deployed multimodal system:*
    - Content moderation, misinformation detection.
    - Biometrics and surveillance.
    - Human-AI interfaces, robotics, etc.

If we want trustworthy multimodal models in the wild, we can't just train once and hope but controllers with guarantees. LS-OGD is a step in that direction~

Herbert Wertheim College of Engineering
Florida Institute for National Security (FINS)
UNIVERSITY *of* FLORIDA

Thank You!

Email: tpan1@ufl.edu

Florida Institute for National Security (FINS)

Appendix: Experiments & Evidence

- **Evaluated on M3A multimodal misinformation benchmark with controlled drifts:**
  - Image corruptions (blur, noise, compression).
  - Text shifts (lexical changes, appended phrases, etc.).
- **Compare:**
  - Standard multimodal model (fixed fusion, fixed learning rate).
  - Same model + LS-OGD controller.
- **Observations:**
  - Baseline degrades significantly under drift.
  - LS-OGD → *Detects performance drops*, *reweights toward the healthier modality*, *and maintains higher accuracy and more stable behavior over time.*
- The controller behaves exactly as the theory predicts: *bounded error, graceful adaptation, and robust handling of broken modalities.*