Role-Specialized Agentic Orchestration

Adaptive Learning + Trustworthy Verification and Validation

Part I: ATLAS | Part II: AIVV

ATLAS
Self-evolution in drifting environments

AIVV
Neuro-symbolic V&V for autonomous systems

arXiv: 2602.02709

arXiv: 2604.02478

Prof. Guang Lin
Department of Mathematics and School of Mechanical Engineering
Purdue University

ICERM Workshop on Agentic Scientific Computing and Scientific Machine Learning
May 10, 2026

MAP

Talk Roadmap

From orchestration idea to two concrete research pipelines.

01 **Introduction**
Why agentic orchestration needs roles, gates, and statistical inner loops.

02 **ATLAS**
Adaptive self-evolution with task-distributed multi-LLM supporters.

03 **AIVV**
Trustworthy V&V through conformal gates and semantic councils.

04 **Conclusion**
Unified view plus future directions for scientific agent systems.

Keep the lens fixed: agents as engineered system components, not free-floating chatbots.

INTRO

The Paradigm of Agentic Orchestration

TAKEAWAY *Role-specialized LLM orchestration is a systems architecture for adaptive learning and trustworthy verification.*

Old view
A single assistant produces answers or code in isolation.

ORCHESTRATION →

New view
LLMs coordinate around data, policies, requirements, and safety gates.

Part I: ATLAS
Self-evolution and adaptation under non-stationary environments.

Part II: AIVV
Verification, validation, and safety bounds for dynamic maneuvering.

Core move: specialize the cognitive jobs, then couple them through measurable gates.

PART I

ATLAS

Adaptive Self-Evolutionary Research Agent with Task-Distributed Multi-LLM Supporters

Research frame
Long-horizon adaptation under drift

Training engine
EvoDPO with controlled reference promotion

Students
Ujin Jeon, Jiyong Kwon, Madison Ann Sullivan, Caleb Eunho Lee

ATLAS: https://arxiv.org/abs/2602.02709

ATLAS

The Challenge in Non-Stationary Environments

TAKEAWAY *Frozen agents and static fine-tuning pipelines are brittle when the world keeps moving.*

Frozen optimizers
Agents are treated as fixed solvers rather than developing systems.

Optimization only
Prior multi-agent work often improves search, not agent development.

Distribution drift
Changing tasks break fixed fine-tuned pipelines over long horizons.

Missing control
Robust self-evolution needs reference management and safety gates.

Research gap **How can an LLM agent improve itself progressively without drifting into stale references, mode collapse, or unstable preference updates?**

ATLAS

Introducing ATLAS

TAKEAWAY *ATLAS decomposes self-improvement into role-specialized support around one evolving research agent.*

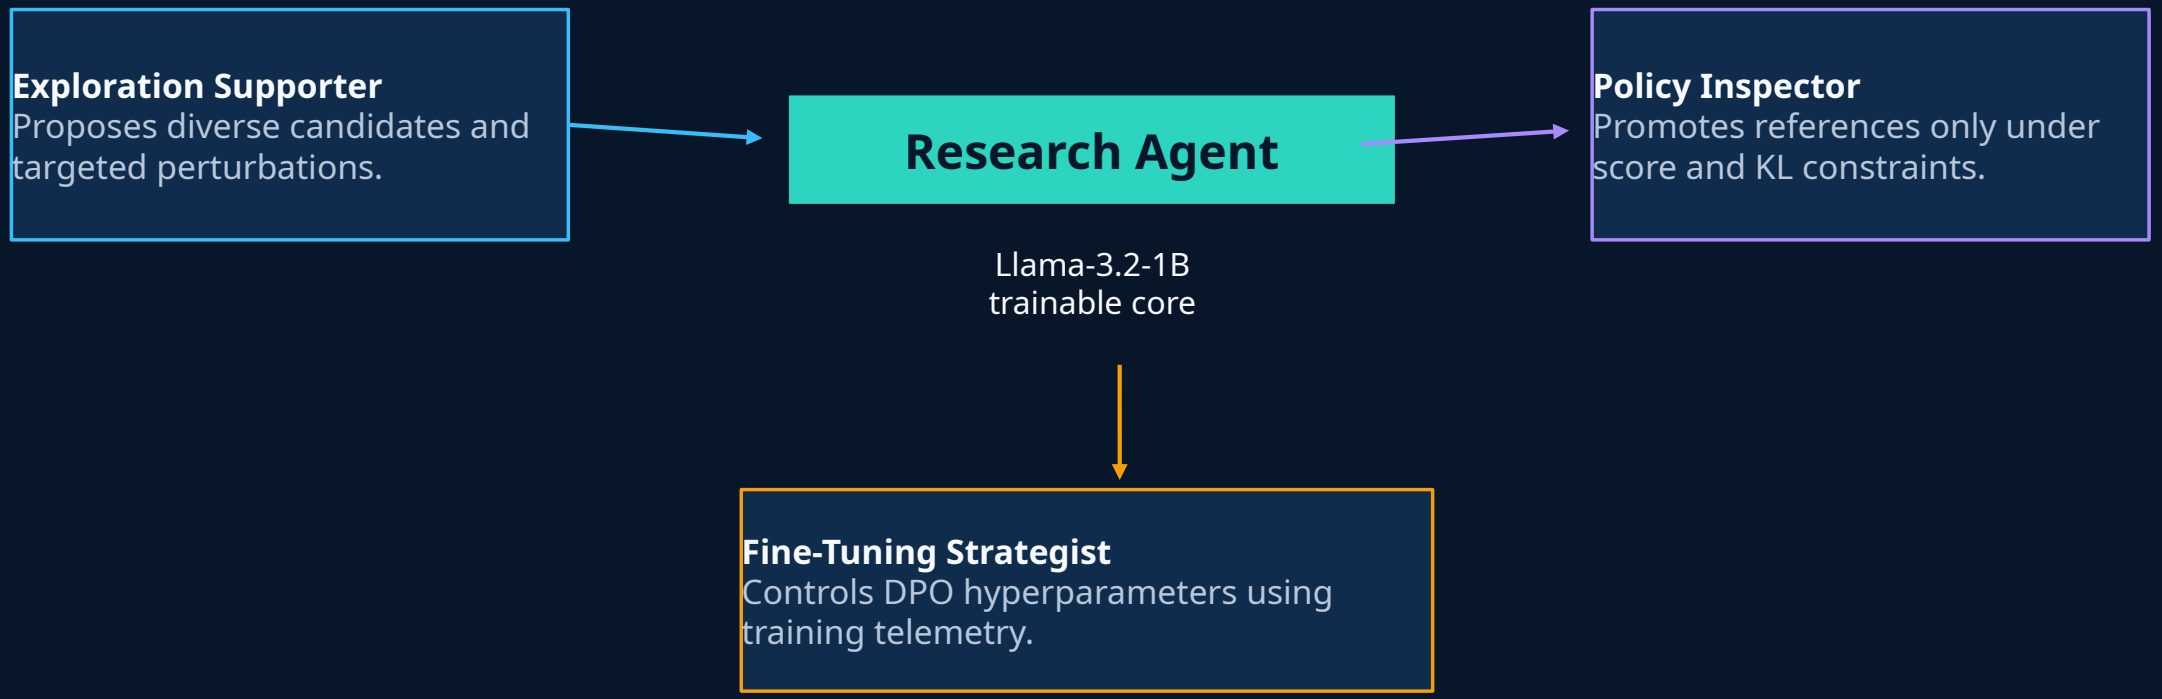

Core idea: distribute the jobs so preference learning stays adaptive, diverse, and inspectable.

ATLAS

ATLAS Workflow Overview

TAKEAWAY *The loop alternates candidate exploration with EvoDPO updates, while a policy-inspector gate controls reference promotion.*

Phase 1: Exploration
supporter-guided candidate generation

Phase 2: EvoDPO updates
strategist tuning + inspector promotion

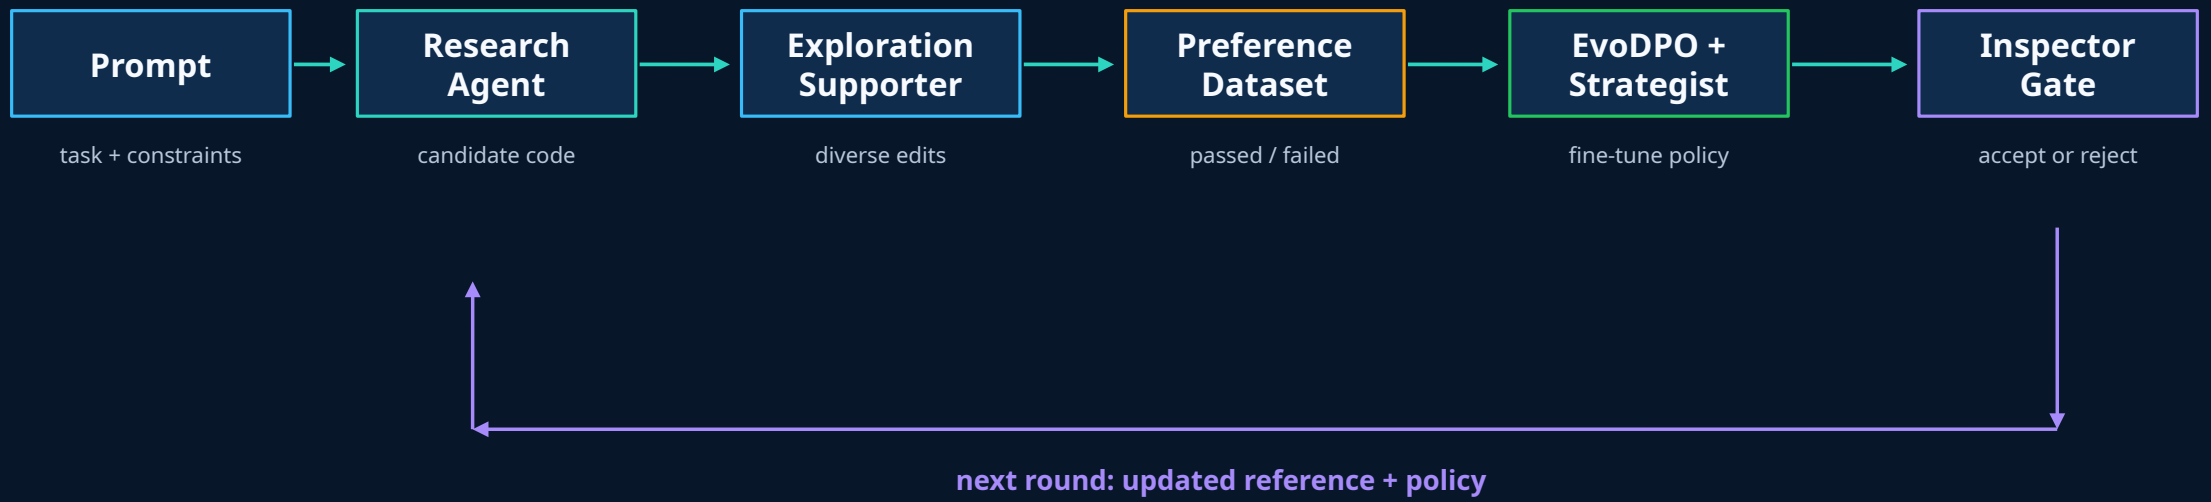

ATLAS

Task-Distributed Multi-LLM Supporters

TAKEAWAY *ATLAS gains control by assigning different cognitive jobs to different agents instead of asking one model to do everything.*

Model	Role	Task allocation
gpt-oss-120B	**Exploration Supporter**	candidate diversity + domain-aware perturbations
DeepSeek-R1-32B	**Fine-Tuning Strategist**	DPO telemetry control + hyperparameter steering
Llama4-latest	**Policy Inspector**	score/KL-gated reference promotion
Llama-3.2-1B	**Research Agent**	evolving executable solution writer

Design principle: separation of duties gives the training loop knobs, guards, and diversity.

ATLAS

The Self-Evolving Research Agent

TAKEAWAY *The research agent writes executable solutions, receives evaluator-induced preferences, and updates its policy.*

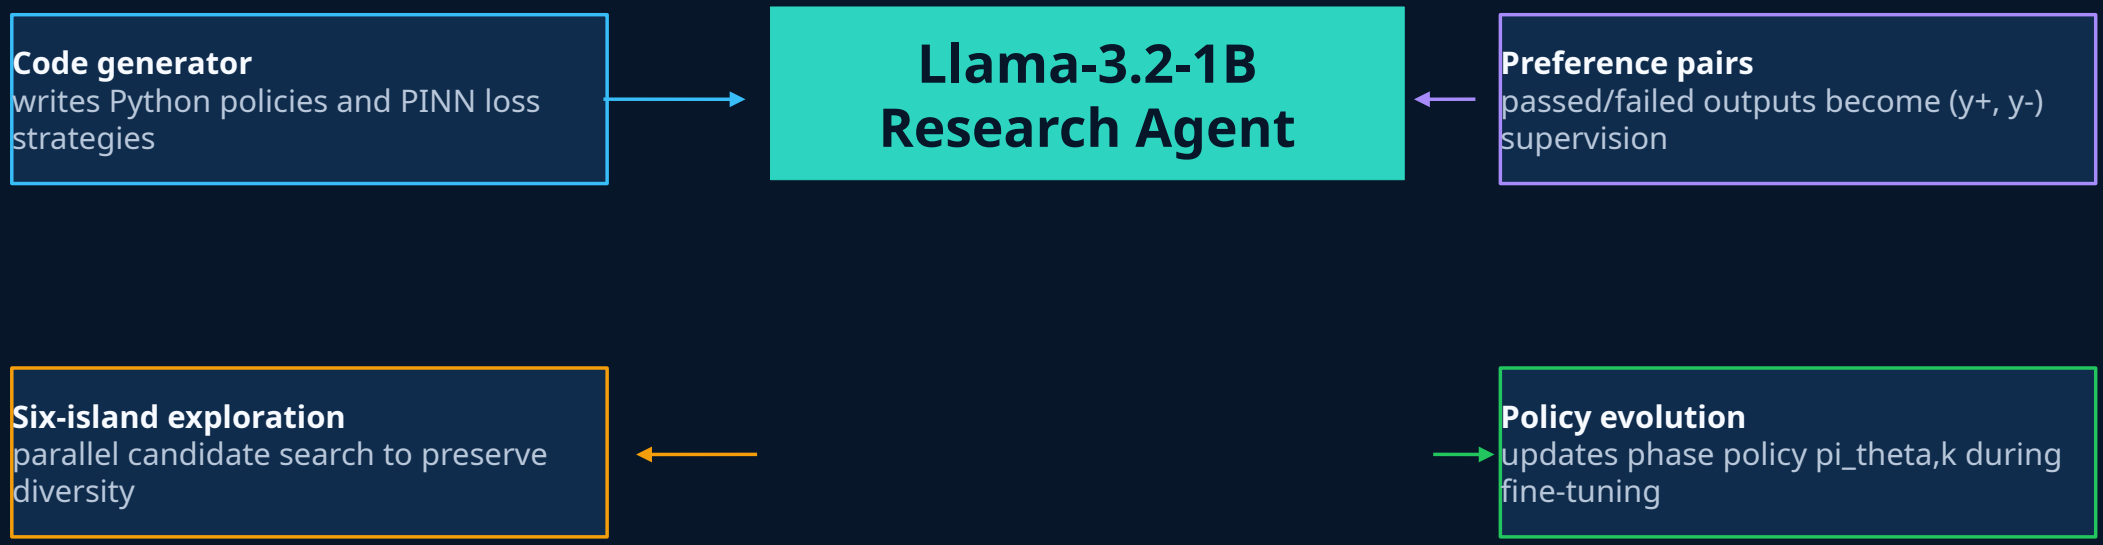

Mode collapse defense: track cluster coverage and keep both successful and failed candidates informative.

ATLAS

Role 1: Exploration Supporter

TAKEAWAY *The exploration supporter proposes targeted, domain-aware modifications to prevent mode collapse.*

Model
gpt-oss-120B

Function
Senior technical advisor providing static analysis on research-agent code outputs.

Responsibilities
Proposes diverse strategies, candidate completions, targeted perturbations, and bottleneck diagnoses.

Bandits
Adjust sliding window and regularization to respond to high drift.

PINNs
Recommend Huber loss with dynamic normalization to stabilize gradients.

Anti-collapse job: keep the candidate pool broad enough that preference learning has something real to learn from.

ATLAS

Role 2: Fine-Tuning Strategist

TAKEAWAY The strategist acts like a **training-control engineer**, adjusting *DPO hyperparameters* from telemetry instead of using a static recipe.

Model

DeepSeek-R1-32B

Function

Regulates DPO learning dynamics to stabilize preference-based updates

Responsibilities

Monitors training telemetry such as score distributions and accept/reject events; dynamically tunes hyperparameters

Adjustments and impact

Modulates inverse-temperature beta, learning-rate schedules, pair-selection thresholds, and training epochs; prevents overfitting to stale feedback and supports stable structural learning without mode collapse.

Design principle: training control should adapt as the policy and preference data evolve.

ATLAS

Role 3: Policy Inspector

TAKEAWAY The policy inspector is the safety valve: it promotes a new reference only when performance improves within a trusted region.

Model

Model: Llama4-latest

Function

Function: safety gate for adaptive reference management during Evolving DPO.

Responsibilities

Evaluates the proposed KL-regularized policy against the current reference and allows reference updates only when two conditions are satisfied:

1. **Sufficient score improvement:** $\Delta \hat{S}_k \geq \epsilon_s$
2. **Trust-region KL constraint:** $K\hat{L}_k \leq \delta_H$

Why it matters

Prevents unsafe or overly aggressive reference drift while still enabling long-horizon improvement.

ATLAS

The Limitation of Fixed References in DPO

TAKEAWAY **A fixed DPO reference** *becomes stale during long-horizon self-improvement, turning a stabilizer into a bottleneck.*

1 Standard DPO

Optimizes from pairwise preferences but relies on a fixed reference policy π_{ref}.

2 What changes over time

In phase-indexed self-improvement loops, the research agent, policy, and dataset D_k all evolve rapidly.

3 Failure mode

The mismatch between evolving data and a frozen π_{ref} causes slow adaptation, misaligned preferences, and eventual stagnation.

Key point: the reference must evolve safely if the policy and environment drift.

ALGORITHM

Core Algorithm: Evolving DPO (EvoDPO)

Solution: replace the static reference by a phase-indexed reference $\pi_{ref,k}$ that updates progressively.

STEP 1

Anchor to current reference

Optimize the research agent with the DPO objective relative to the current $\pi_{ref,k}$.

→

STEP 2

Propose candidate reference

Construct a candidate that seeks higher task score while penalizing divergence from the previous reference.

→

STEP 3

Estimate drift efficiently

Use a trajectory-conditioned next-token KL estimator to measure policy drift.

→

STEP 4

Safe promotion

Update the reference only if score and KL criteria are satisfied.

Reference evolution: $\underbrace{\pi_{ref,k} \rightarrow \tilde{\pi}_{ref,k+1}}_{\text{candidate proposal}}$ promote if $\underbrace{S(\tilde{\pi}_{ref,k+1}) \geq \tau_{score}}_{\text{score threshold}}$ and $\underbrace{\mathrm{KL}(\tilde{\pi}_{ref,k+1} \,||\pi_{ref,k}) \leq \tau_{kl}}_{\text{KL (drift) threshold}}$

Takeaway

EvoDPO replaces a static reference with a phase-indexed reference that can evolve safely under score and KL controls.

Theoretical Grounding: Dynamic Regret

We analyze the KL-regularized reference-update mechanism through a non-stationary, preference-based contextual bandit formulation.

Dynamic regret decomposition

$$R_T = R_T^{\text{error}} + R_T^{\text{bias}}$$

Learning error

Cumulative regret due to imperfect learning of the optimal policy in a drifting environment.

Reference-induced bias

Cumulative regret from using an outdated reference policy that drifts away from the environment's optimum.

Key theorem: under systematic reference refresh, dynamic regret is bounded by the environment drift budget V_T rather than suffering the linear bias of a fixed reference.

$$O\left(T^{\kappa} V_T + T^{1-\kappa/2} \sqrt{\log T} + T^{1-\kappa} + V_T \right)$$

TAKEAWAY

A fixed reference accumulates linear bias under drift; EvoDPO refreshes the reference so bias scales with the drift budget instead.

Experimental Setup: Two Distinct Domains

Domain 1: Non-Stationary Contextual Bandits

- Adaptive decision-making and policy evolution under concept drift
- Stress test for long-horizon adaptation

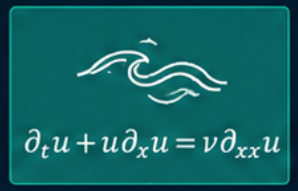

Domain 2: Scientific Machine Learning (SciML)

- Continuous optimization of loss design for the 1D viscous Burgers equation via PINNs
- Stress test for optimization under changing PDE stiffness

Both domains require sustained adaptation, not one-shot optimization.

TAKEAWAY The experiments test ATLAS in two complementary regimes: online decision-making under drift and scientific optimization under stiff PDE dynamics.

ATLAS

Domain 1: Non-Stationary Contextual Bandits

TAKEAWAY The bandit task is a clean stress test for adaptation under concept drift, where the agent must write policies without direct reward-oracle access.

01 Problem

A k-armed bandit with a reward parameter θ_t that drifts over time on the unit sphere, subject to a total variation budget V_T.

02 Agent task

The research agent writes a Python policy function that dynamically selects arms without direct access to expected rewards.

03 Adaptation challenge

Supporter agents propose and tune sliding-window LinUCB hyperparameters such as window size and regularization λ to track drift.

04 Evaluation metric

Negative Mean Regret (NMR); higher is better.

The task isolates adaptation quality from raw model size.

ATLAS

Domain 2: SciML (Burgers' Equation PINN)

TAKEAWAY The PINN task shows that agentic self-evolution can discover executable training objectives for stiff scientific machine learning problems.

01

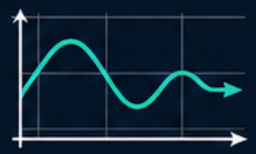

Problem

Neural-network training for PDEs is highly sensitive to physical parameters.

02

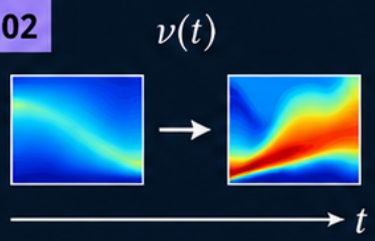

Drift dynamics

A time-varying viscosity schedule $\nu(t)$ shifts the system from a smooth diffusion-dominated regime to a stiff, shock-forming convection regime

03

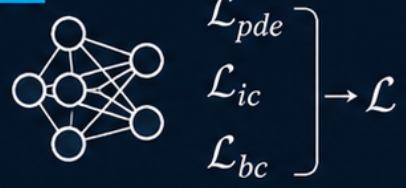

Agent task

Design a Python function to adaptively reweight the composite PINN loss $\mathcal{L}_{pde}$, $\mathcal{L}_{ic}$, and $\mathcal{L}_{bc}$ across epochs to avoid the trivial zero solution.

04

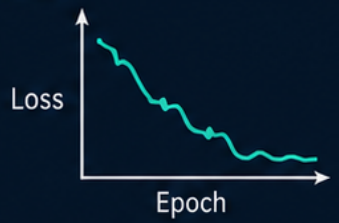

Evaluation metric

Validation loss on held-out collocation points; lower is better.

This domain tests whether ATLAS can improve scientific optimization objectives, not just decisions.

ATLAS

Baselines for Comparison

To isolate the impact of the proposed components, ATLAS is compared against two baselines.

METHOD	MAIN IDEA	WHAT IS MISSING
EvoTune (Static Baseline)	standard single-agent self-evolution; fixed reference policy and static hyperparameters; no supporter guidance.	No adaptive reference updates (EvoDPO); no supporter agents; no policy–inspector gate.
EvoDPO (Ablation)	single agent with an updating reference via EvoDPO; supporters and policy–inspector gating disabled.	No task-distributed supporters; no policy–inspector gate.
ATLAS (Ours)	EvoDPO plus task-distributed supporters and policy–inspector gate.	None — full system with adaptive reference, supporters, and gated evolution.

TAKEAWAY

The baselines isolate whether gains come from adaptive references, supporter agents, or their combined synergy.

This comparison separates algorithmic adaptivity from multi-agent support.

RESULTS

Performance Trajectories

ATLAS improves steadily over rounds, while the baselines plateau earlier under drift and stiffness.

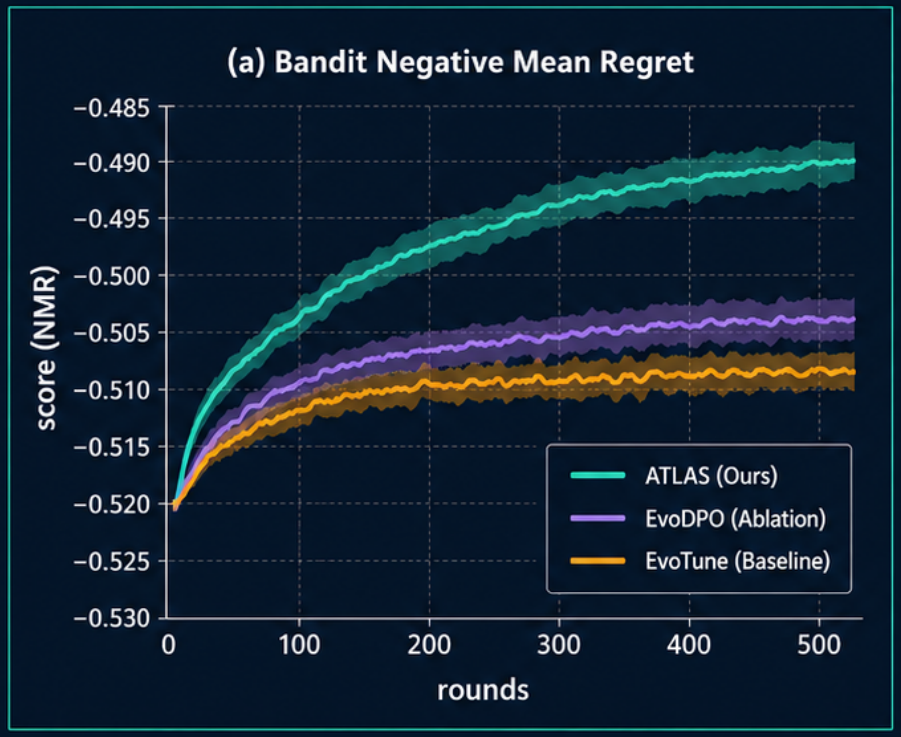

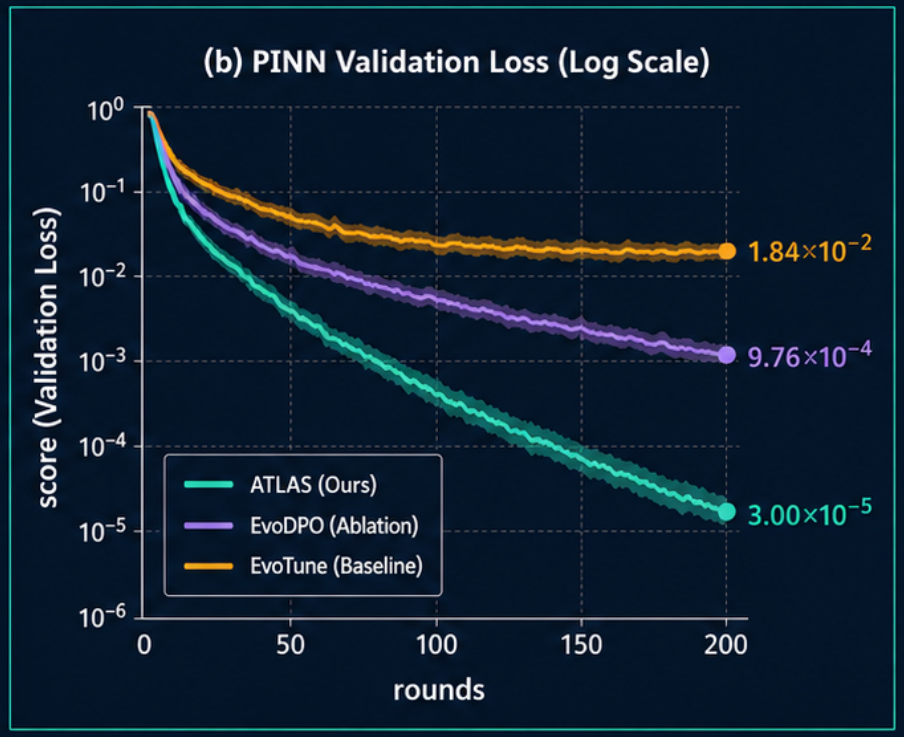

Analysis: EvoDPO slightly improves over the static baseline, but the **multi-supporter layer of ATLAS** is what continuously overcomes concept drift and physical stiffness.

RESULTS

Quantitative Summary

TAKEAWAY ATLAS achieves the strongest final bandit improvement and the largest PINN loss reduction.

Method	Bandit NMR Initial	Bandit NMR Final	Improvement	PINN Loss Initial	PINN Loss Final	Reduction
EvoTune	-0.621	-0.509	18.0%	0.880	1.84×10^{-2}	48x
EvoDPO	-0.621	-0.507	18.3%	0.880	9.76×10^{-4}	902x
ATLAS	**-0.621**	**-0.493**	**20.6%**	**0.880**	$\mathbf{3.00 \times 10^{-5}}$	**29,344x**

Key result: ATLAS couples adaptive references with role-specialized supporters to deliver the largest gain in both domains.

RESULTS

Agent Dynamics & Intervention Analysis

TAKEAWAY

The agents adapt their behavior by domain: **stricter gatekeeping in volatile bandits** and **more reference acceptance in structured PINN training**.

	Strategist Δβ	Strategist Threshold	Strategist Epoch	Inspector Accept Rate
Bandit	22.7%	31.1%	7.3%	16.4%
PINN	16.0%	24.0%	2.0%	38.0%

Bandit: strict conservatism gate prevents aggressive drift in a volatile task.

PINN: deterministic physics enables more consistent reference promotions.

ATLAS

Ablation Study: Isolating Contributions

01	Baseline (EvoTune)	Bandit final -0.509	PINN final 1.84×10^{-2}
02	+ Adaptivity (EvoDPO)	Bandit final -0.507	PINN final 9.76×10^{-4}
03	+ LLM Supporters (ATLAS)	Bandit final -0.493	PINN final 3.00×10^{-5}

Adaptivity: replacing a static reference with EvoDPO reduces PINN loss by nearly two orders of magnitude.

Synergy: adding supporters prevents stagnation, navigates bandit drift, and refines SciML loss by another order of magnitude.

TAKEAWAY The ablation supports the central mechanism: EvoDPO improves stability, while LLM supporters add search diversity and task-specific structure.

ATLAS

ATLAS Conclusion: Sustained Agentic Evolution

TAKEAWAY ATLAS demonstrates that long-horizon agent improvement becomes more stable when exploration, training strategy, and reference inspection are separated.

Architectural success

Task-specialized roles stabilize long-horizon preference learning.

Algorithmic impact

EvoDPO resolves reference mismatch through safely gated adaptive reference updates.

Robustness

Adaptivity plus multi-LLM oversight prevents stagnation in bandits and SciML optimization.

Next: AIVV uses the same orchestration principle for trustworthy verification and validation.

PART II

Part II: AIVV

Neuro-Symbolic LLM Agent-Integrated Verification and Validation for Trustworthy Autonomous Systems

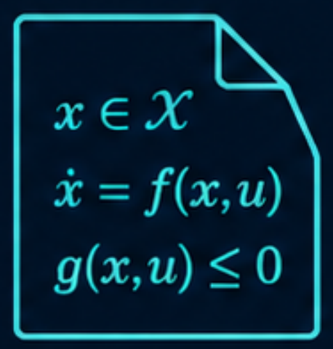

Mathematical Bounds
Dynamics, Constraints, Specifications

LLM Council
Multi-LLM Deliberation and Reasoning

Safety Gate
Verified, Validated, and Trusted

Prof. Guang Lin, Purdue University

My students: Jiyong Kwon, Ujin Jeon, Sooji Lee

AIVV: https://arxiv.org/abs/2604.02478

AIVV

The Scalability Crisis in Autonomous Systems

TAKEAWAY Manual HITL V&V cannot keep up with high-frequency, noisy telemetry from autonomous systems.

High-frequency telemetry

Mission-critical systems such as UUVs generate massive noisy streams with sparse true fault events.

HITL bottleneck

Algorithms flag faults in milliseconds, but human teams may need minutes or hours to triage and validate redesigns.

?

How do we scale trustworthy V&V without drowning engineers in false alarms?

Dynamic maneuvering

Large transient responses can trigger false positive alarms.

Need

Digitize human oversight using rapid semantic reasoning over requirements and telemetry.

AIVV

The Problem of Semantic Blindness

TAKEAWAY Purely mathematical detectors can detect deviations, but they cannot explain whether the deviation is a nuisance transient or a true failure.

What math sees

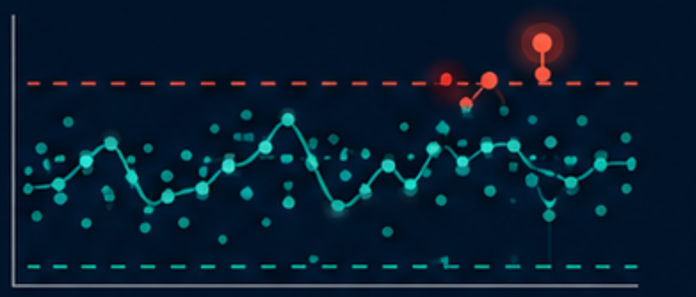

Data-driven LSTMs detect statistical anomalies, but do not know mission intent or maneuver context.

Tightening bounds catches gradual drift but can trigger a cascade of false alarms.

What semantics adds

Requirement-aware reasoning distinguishes nuisance transients from true mechanical or electrical failures.

Introducing AIVV:

Agent-Integrated Verification and Validation

Agents combine statistical detection with semantic understanding to reduce false alarms and focus on true mission-impacting failures.

AIVV

AIVV Architecture Overview

TAKEAWAY AIVV is a two-layer neuro-symbolic system: fast mathematical gating first, semantic multi-agent validation only when needed.

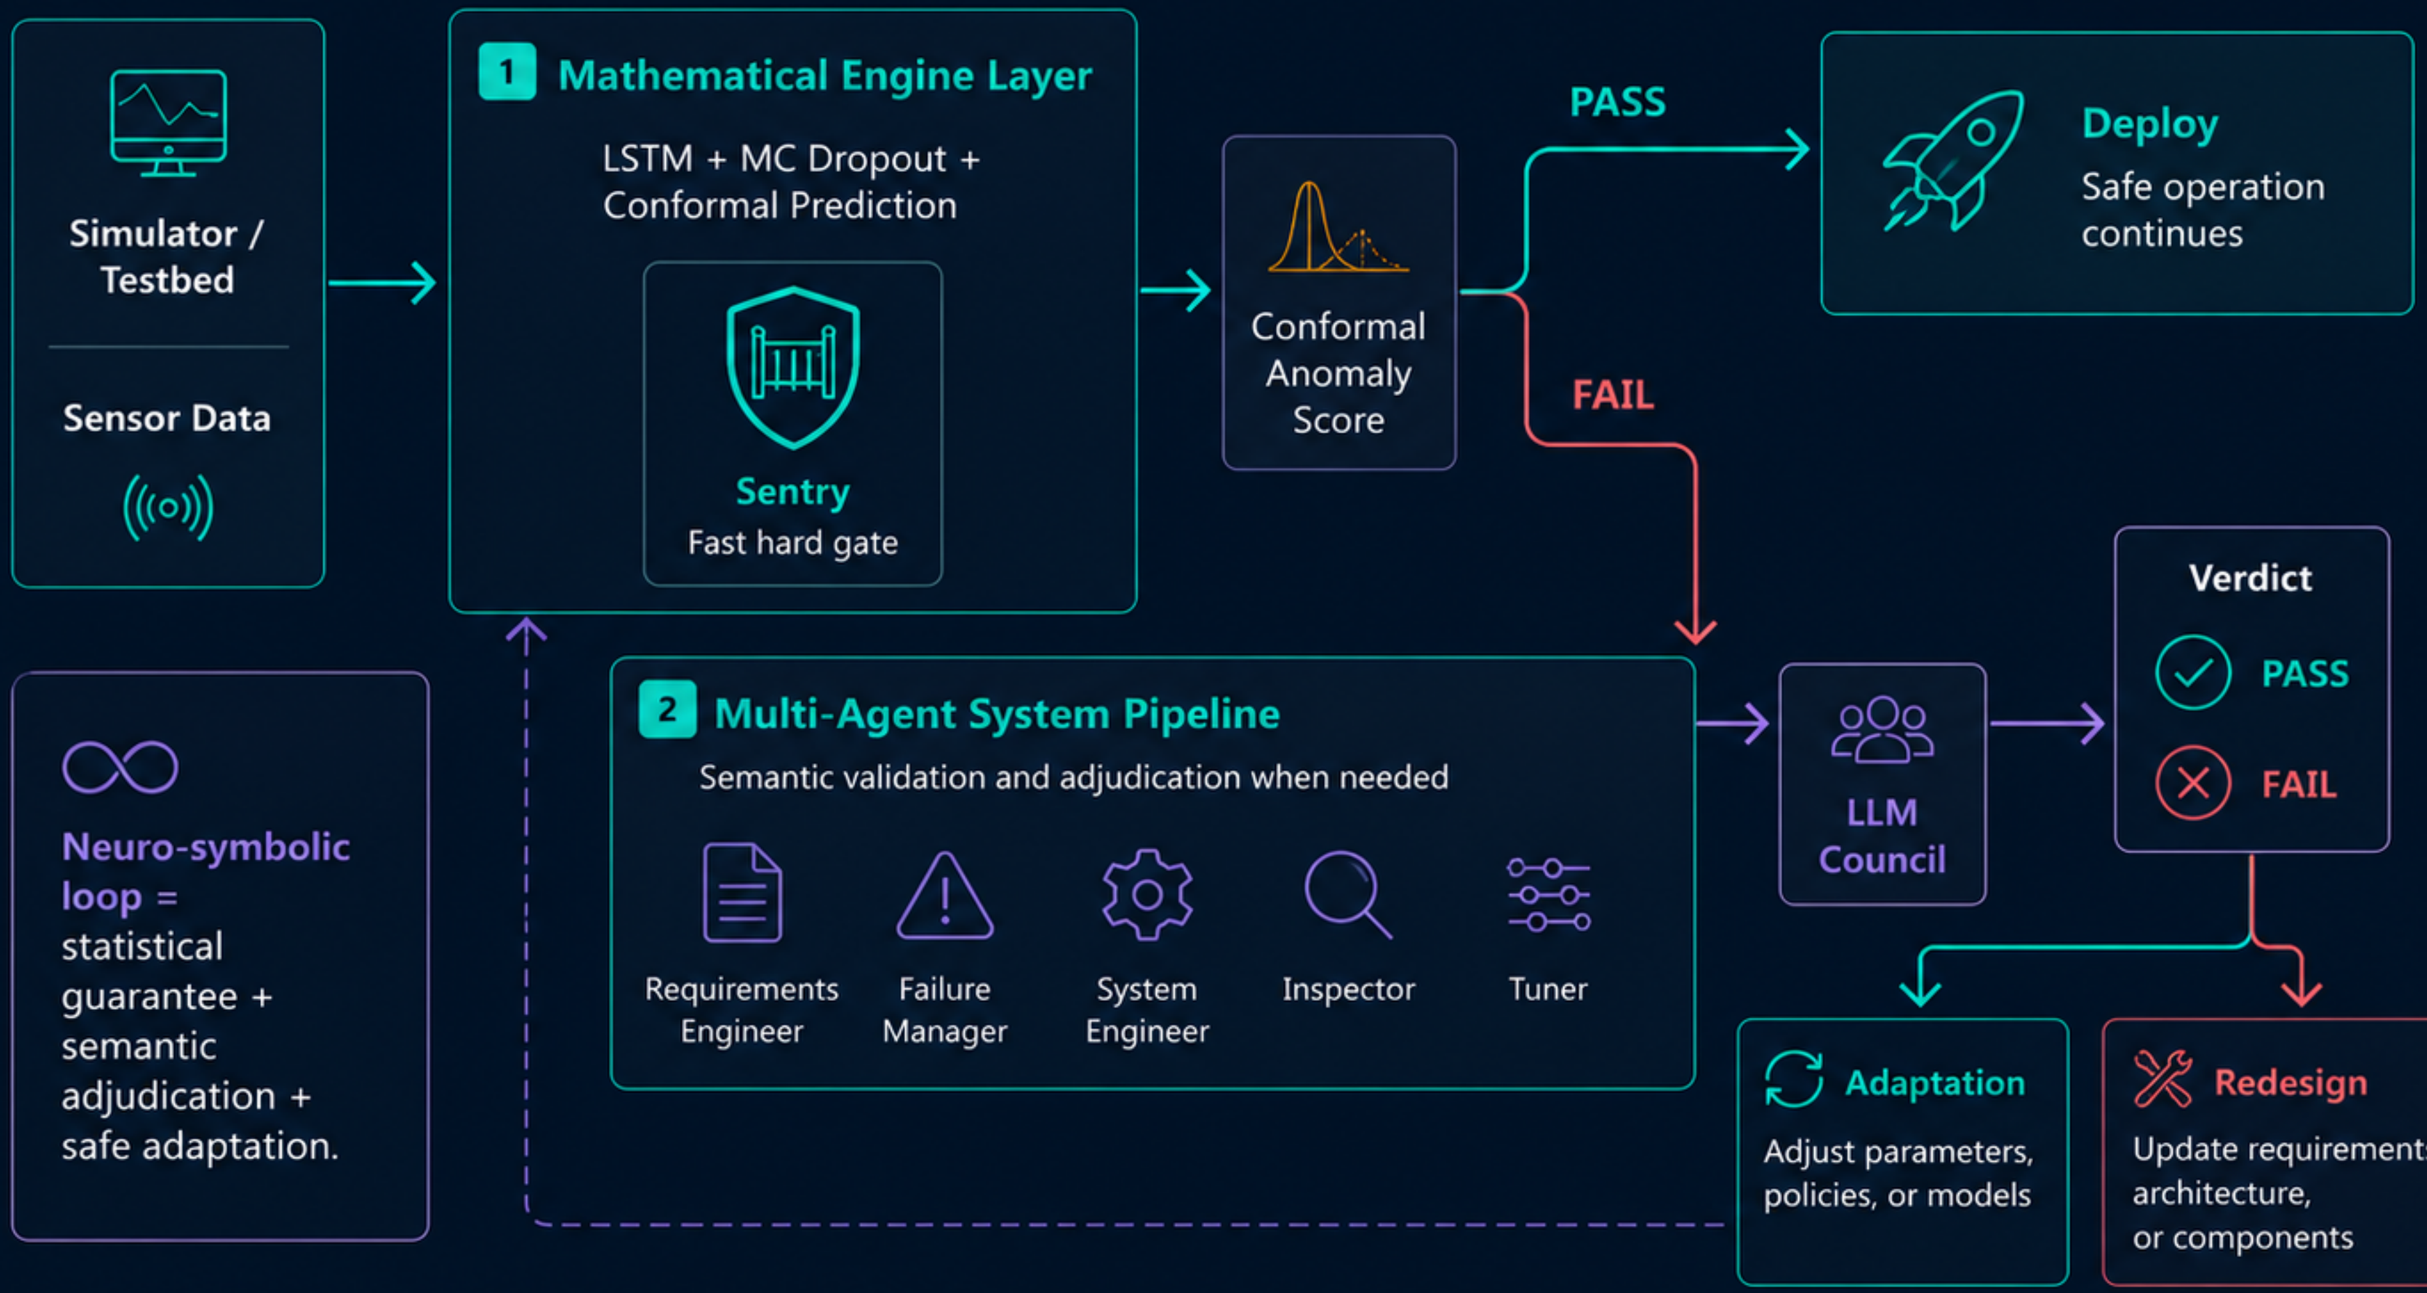

AIVV

Layer 1: Math Engine & Uncertainty Quantification (UQ)

TAKEAWAY

The math engine supplies a **statistically grounded front line** by combining sliding-window LSTM prediction with MC dropout uncertainty.

1 Sliding-window sequences

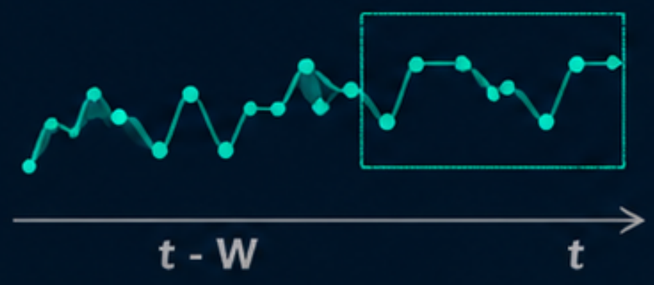

Recent multivariate observations with window length W.

2 LSTM predictor

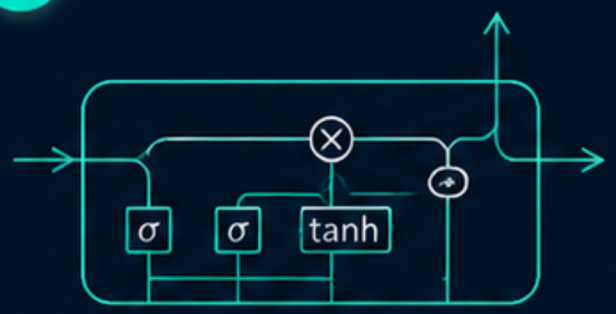

Captures temporal dynamics and predicts the next value.

3 MC Dropout uncertainty

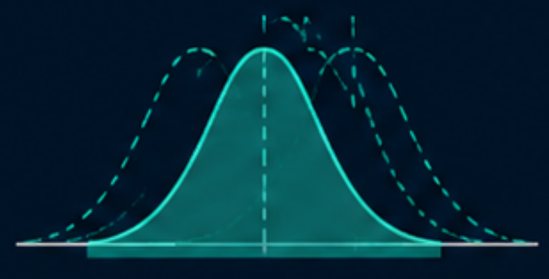

Stochastic forward passes with dropout enabled at inference.

μ **predictive mean $\hat{y}_t$**
(Expected next value)

σ **epistemic uncertainty σ_t**
(Model uncertainty)

Stochastic forward passes approximate Bayesian inference and capture temporal dependencies without requiring explicit physical models.

Role: fast front-line anomaly screening before LLM reasoning is invoked.

AIVV

Deterministic Conformal Gating (The Sentry)

Statistical guarantee: finite-sample marginal coverage from conformal prediction.

TAKEAWAY The **Sentry** provides the hard statistical gate: only residuals outside the calibrated conformal boundary are escalated.

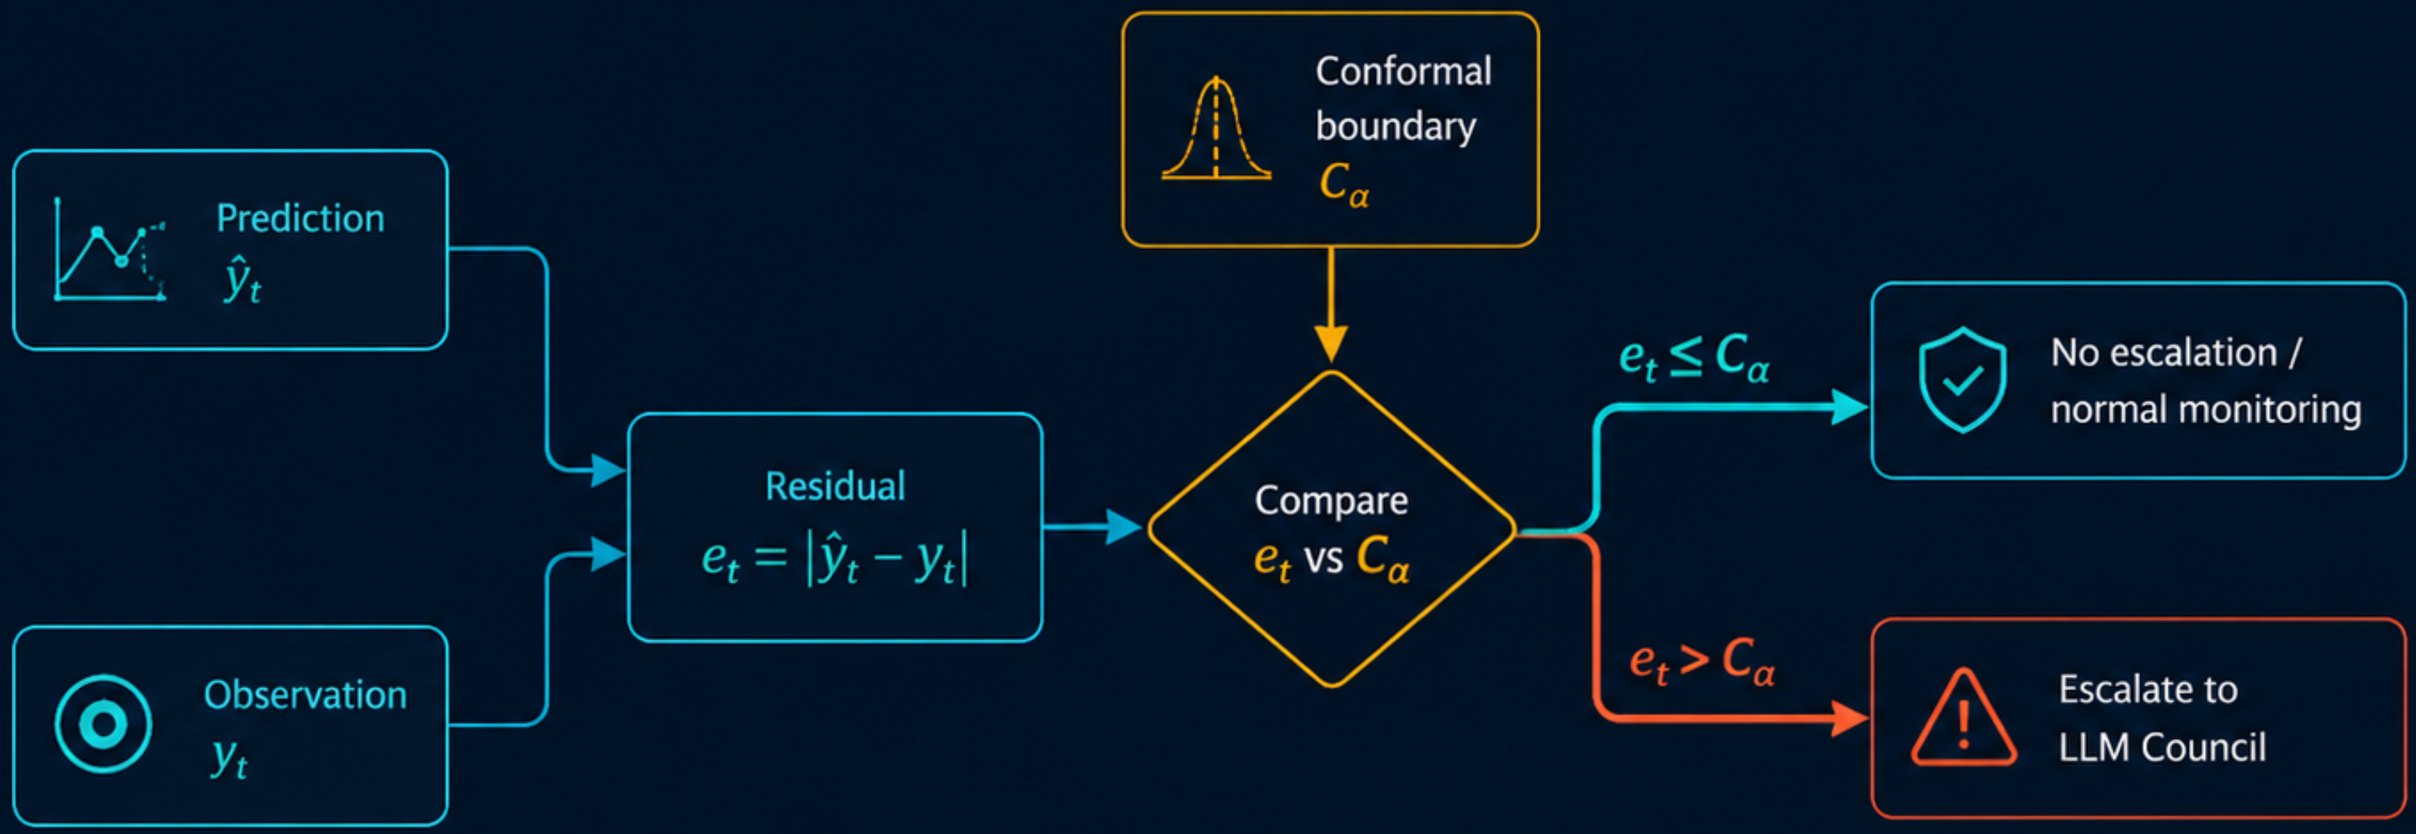

Purpose: preserve latency by invoking the LLM council only when calibrated bounds are violated.

AIVV

Phase 2: Deliberative Adjudication (The Council)

TAKEAWAY The LLM council adds semantic validation, using majority vote to reduce single-model hallucination risk.

1 When the **Sentry** flags an anomaly, it is escalated to a role-specialized LLM council.

2 **Objective:**

Provide deliberative, semantic oversight to differentiate nuisance faults from true failures.

3 **2-of-3 Majority Rule**

Three independent LLMs vote. A sample is confirmed as a true failure only if **at least two agents vote FAIL.**

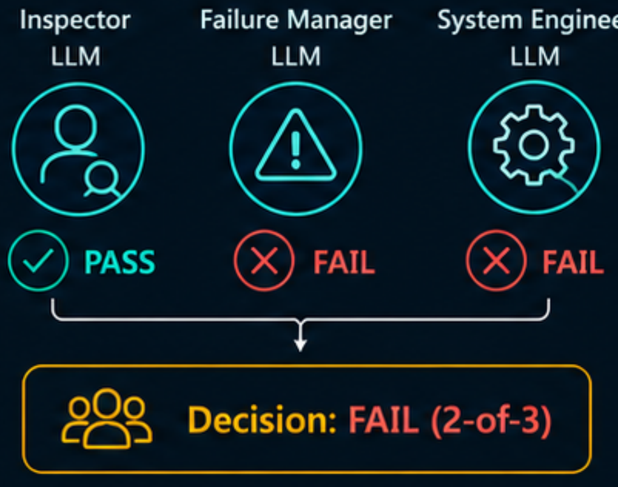

4 **Benefit:**

Digitizes the collaborative cross-validation process of a human engineering team, mitigating standalone LLM hallucinations.

Verdict

PASS

FAIL

AIVV

Council Roles: Requirements & Failure Managers

TAKEAWAY The Requirements Engineer and Failure Manager split semantic validation into normal-mode compliance and failure-mode severity analysis.

Requirements Engineer (LLaMA-4-17B)

Normal mode compliance

- Evaluates whether current behavior satisfies operational tolerances.
- **Returns:** compliance vote, cited section, analytical justification.

Failure Manager (GPT-OSS-120B)

Failure mode and effect analysis

- Evaluates severity of trajectory deviation.
- **Returns:** compliance vote, categorical failure-mode level, rationale.

Together, the two roles separate routine requirement checking from fault-severity reasoning.

AIVV

Council Role: System Engineer

TAKEAWAY The System Engineer connects diagnosis to mitigation by using **domain knowledge** to propose **physically meaningful corrective actions**.

System Engineer (LLaMA-3.3-70B)

Domain synthesis and mitigation strategy

- Uses explicit knowledge of UUV dynamics (Nomoto models, PID parameters).
- Understands math detection mechanisms (MC dropout, conformal prediction).

Actionable Artifacts

Aggregates findings from the Requirements Engineer and Failure Manager.

If a fault is confirmed, generates a **control-system gain-tuning proposal** to stabilize the system.

Key idea: diagnosis becomes **redesign guidance**, not just a pass/fail label.

AIVV

Role-Specific Model Alignment (Ablation)

TAKEAWAY The choice of which model fills which role matters; mismatched assignments can collapse validation performance.

Configuration	Requirements Engineer	Failure Manager / System Engineer	Inspector / Tuner	Validation FVR
Optimal	LLaMA-17B	GPT-120B / LLaMA-70B	Qwen-32B / GPT-20B	**100%**
Config 1	GPT-20B	LLaMA-17B / GPT-120B	LLaMA-70B / Qwen-32B	**56.0%**
Config 2	Qwen-32B	GPT-20B / LLaMA-17B	GPT-120B / LLaMA-70B	**93.33%**
Config 3	LLaMA-70B	Qwen-32B / GPT-20B	LLaMA-17B / GPT-120B	**44.0%**

Systematic cyclic shifts in role assignment reduce validation quality, showing that **model-role matching is crucial**.

AIVV

Phase 3: The Adaptation Pipeline

TAKEAWAY When a flagged event is judged to be nuisance behavior, AIVV uses it as a safe opportunity to **recalibrate** or **fine-tune** the math engine.

Council votes PASS
Anomaly is a nuisance/transient response.

Trigger dynamic retraining of the math engine.

Inspector Agent
(Qwen3-32B)

Synthesizes JSON logs and outputs one action: **RECALIBRATE**, **FINE TUNE**, or **TRY BOTH**.

Tuner Agent
(GPT-OSS-20B)

Applies the selected action to a temporary cloned engine.

The adaptation path is activated only for nuisance behavior, not true failures.

AIVV

Clone-and-Promote Safeguards

Challenge: Blindly updating the model in the loop can cause catastrophic forgetting.

TAKEAWAY AIVV avoids unsafe online model updates by **testing every adaptation on a clone** before deployment.

1 Tuner applies updates to a clone f^{θ}_{tmp}.

2 The **Sentry** is invoked a second time on the cloned model.

Sentry
(Second Pass)

3 **Promote:**
If the candidate resolves the conformal violation, it becomes the new deployed model.

PASS

Promote to Deployed Model

4 **Discard:**
If it fails, the deployed engine remains unchanged, protecting operational stability.

FAIL

Keep Current Deployed Model

Experimental Setup: REMUS 100 UUV

TAKEAWAY The AIVV experiments use a realistic UUV simulator with sensor noise, drift, and injected faults to stress-test validation.

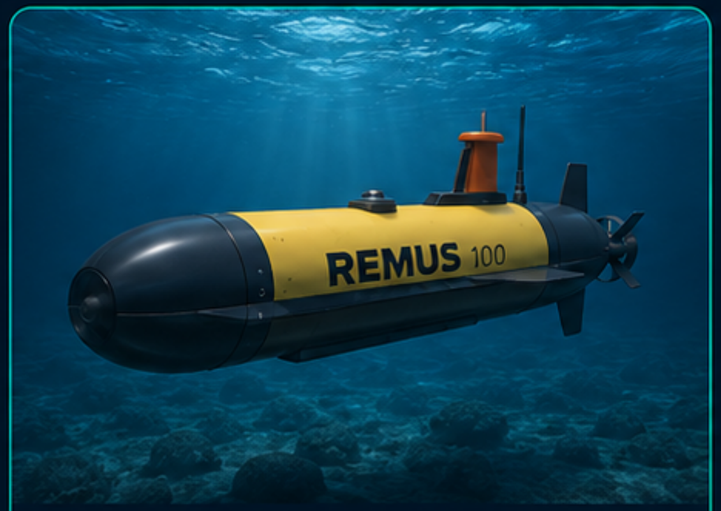

Environment:
Simulink model of the REMUS 100 UUV with IMU sensor fusion.

Conditions:
Added sensor-drift bias and environmental non-Gaussian noise to challenge the models.

Fault Injection:
Introduced electrical sensor failures and mechanical damper failures at specific timesteps to trigger hardover failure responses.

Hardover Response:
Faults that violate conformal bounds trigger escalation to the LLM Council.

Goal: evaluate trustworthy validation under realistic disturbance and fault conditions.

AIVV

Tested Mission Datasets

TAKEAWAY The three datasets increase realism and maneuver complexity, moving from hovering to structured mapping to a nonlinear complex mission.

1 Dataset 1: Hovering

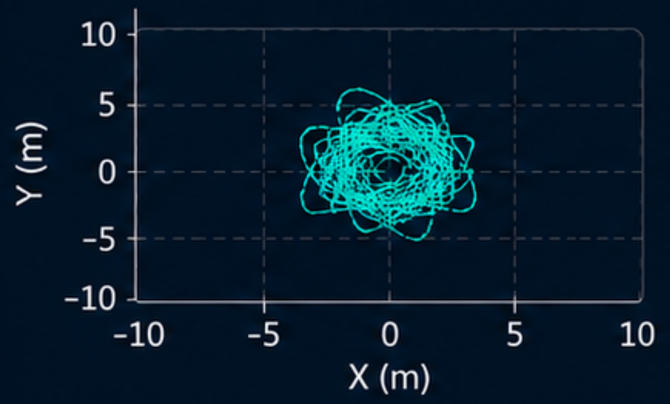

(a)

Station-keeping. Normal control oscillations vs. genuine faults.

2 Dataset 2: Lawnmower

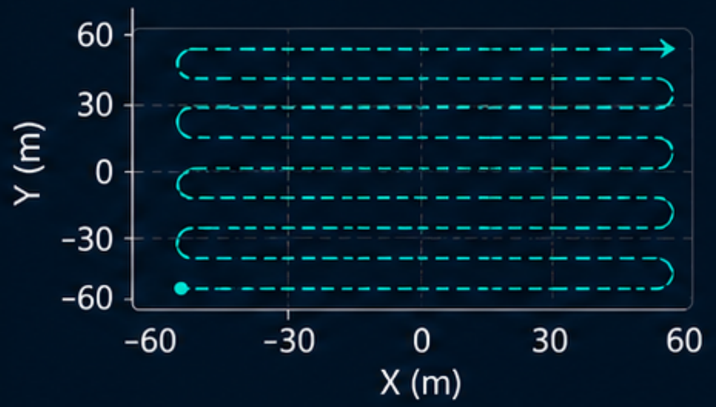

(b)

Structured periodic grid-mapping. Aggressive dynamic inputs vs. faults.

3 Dataset 3: Complex Mission

(c)

Aperiodic course corrections and varying speeds in a highly non-linear profile.

Qualitative Ablation: Mitigating False Alarms

TAKEAWAY The visual ablation shows the architecture progressively suppressing false alarms caused by transient maneuvers.

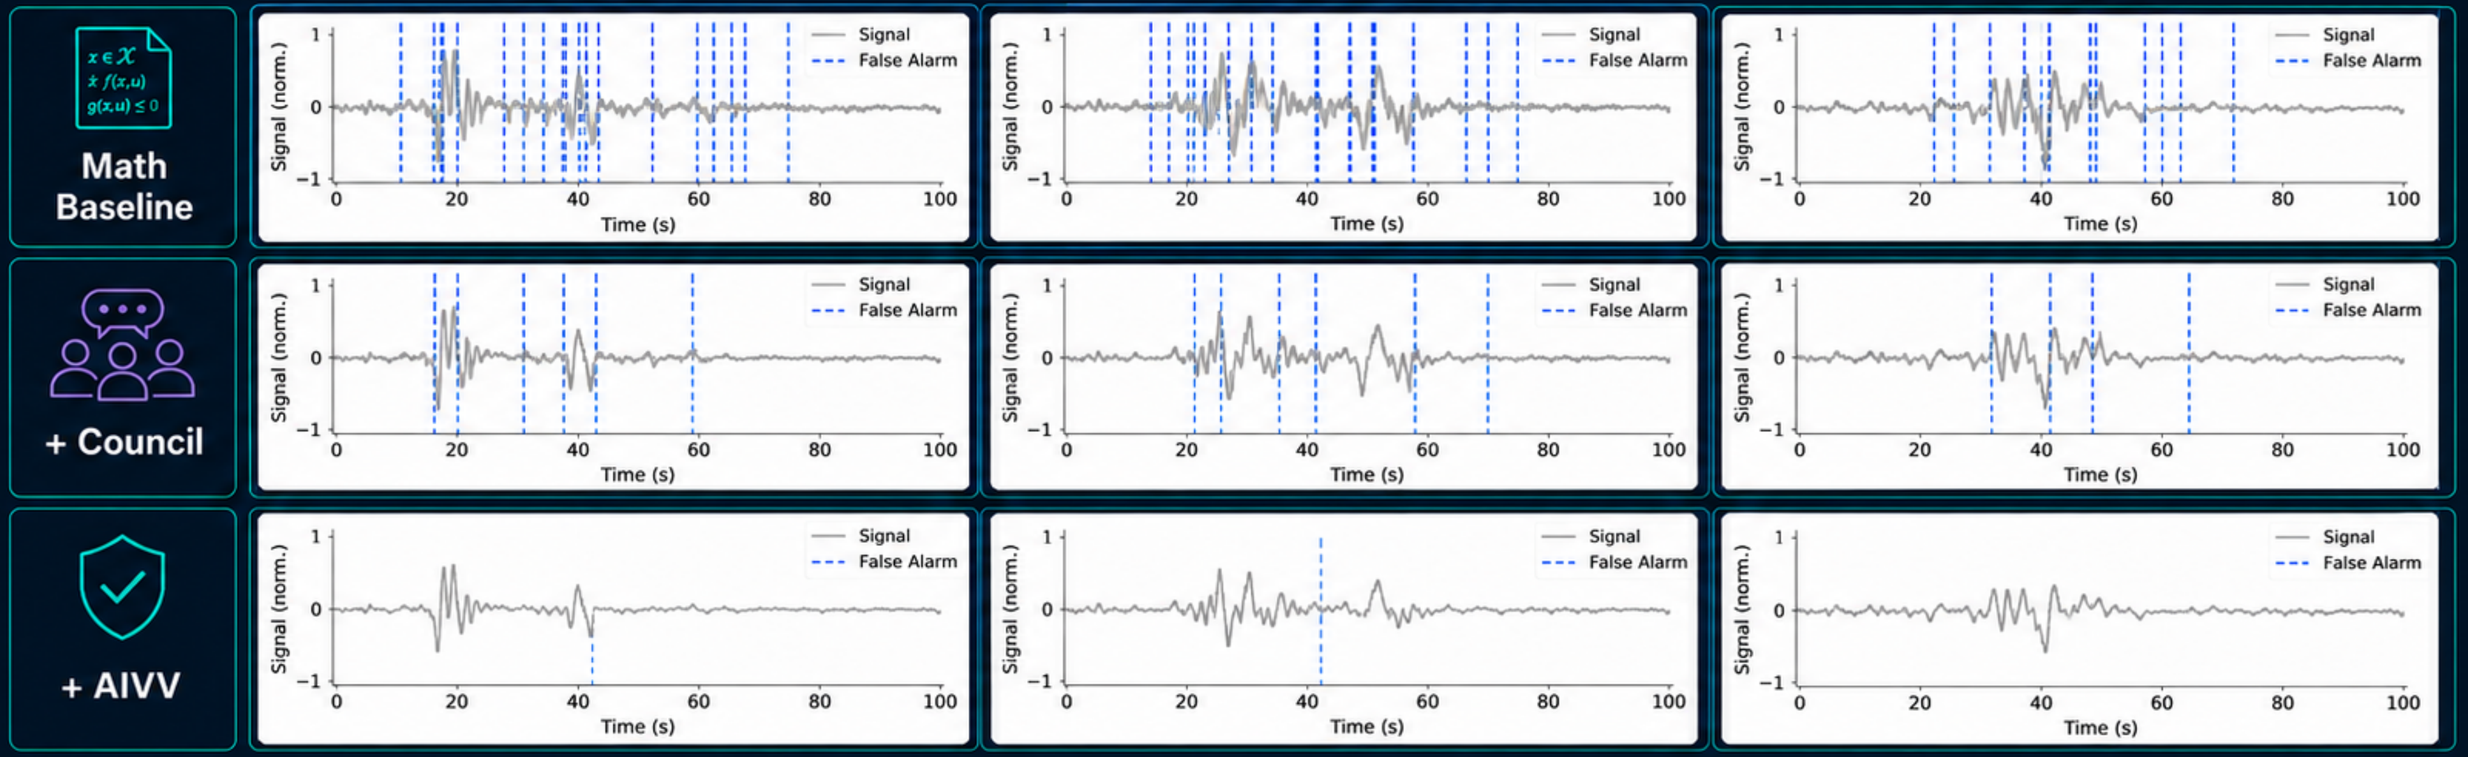

Progression (rows): Math Baseline → LLM Council → Full AIVV.

Observation: The mathematical baseline exhibits a severe false-positive rate during transients. The LLM Council visually suppresses these.

AIVV

Quantitative Validation: Failure Validation Rate (FVR)

TAKEAWAY Quantitatively, AIVV converts the visual false-alarm suppression into high FVR across all mission profiles.

Method	Dataset 1 (Hover)	Dataset 2 (Lawnmower)	Dataset 3 (Complex)
Math Engine + Sentry (Baseline)	**45.33%**	**0%**	**0%**
+ The Council	**98.67%**	**80%**	**73.33%**
+ Adaptation (AIVV)	**100%**	**89.33%**	**93.33%**

While math alone drops to 0% **FVR** in dynamic maneuvers (Datasets 2 and 3), the full AIVV pipeline safely suppresses false positives to maintain at least **89.33%** accuracy.

AIVV

Adaptation Performance Improvements

TAKEAWAY The clone-and-promote adaptation mechanism improves the math engine where the mission profile is complex.

The clone-and-promote dynamic tuning significantly improves true math engine accuracy, particularly in complex, realistic mission profiles.

Dataset	$\text{Accuracy}_{\text{initial}}$	$\text{Accuracy}_{\text{tuned}}$	Improvement (%)
Dataset 1 (Hovering)	**0.954**	**1.000**	**+4.82%**
Dataset 2 (Lawnmower)	**0.994**	**0.997**	**+0.30%**
Dataset 3 (Complex)	**0.688**	**0.847**	**+23.11%**

Council Loop Accuracy before and after the Adaptation Pipeline.

Largest improvement occurs in the most complex mission profile.

AIVV

Actionable Outcomes: Systemic Redesign

TAKEAWAY AIVV is not only an alarm filter; it translates validated failure analysis into structured redesign actions.

1 The Task

Translating failure-mode analysis into structured corrective actions.

Failure Analysis

2 System Engineer Output

Proposed adjusted PID gain-tuning parameters based on Failure Manager anomaly flags.

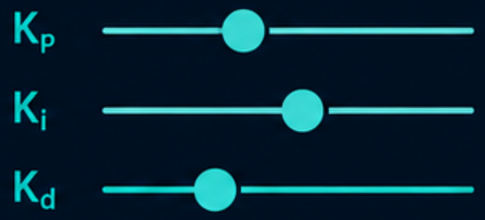

Gain-Tuning Proposal

3 Result

Bridging the V&V gap between fault identification and actionable system redesign.

Redesign Action

Validated diagnosis becomes a structured corrective action.

AIVV

System Engineer: Actionable V&V Output

TAKEAWAY The reasoning log demonstrates that the system can produce auditable, sample-specific gain-tuning recommendations.

The Task: Translating failure-mode analysis into structured corrective actions.

LLM Reasoning Log (Snippet)

```
[System Engineer -- Active Optimizer]
Gain-Tuning Proposals (23 unique samples, triggered by FM/RE FAIL):
- Sample 221 | Triggered by: RE | SE Vote=FAIL |
    Params: {Kp: 0.6, Ti: 19.0, Td: 1.0, Reference_Max_Velocity: 9.5}
    Reason: Since the requirements engineer voted FAIL due to an operational limit
            violation, the tuning proposal aims to reduce the error magnitude.
            The proportional gain (Kp) is increased to ...

- Sample 225 | Triggered by: FM+RE | SE Vote=FAIL |
    Params: {Kp: 0.7, Ti: 15.0, Td: 1.2, Reference_Max_Velocity: 9.0}
    Reason: Since both failure manager and requirements engineer voted FAIL,
            adjusted parameters are proposed. Increased Kp to 0.7 to reduce error,
            decreased Ti to 15.0 to reduce oscillations.
```

Result: The framework autonomously bridges the V&V gap between **fault identification** and **physically grounded** control updates.

AIVV

Gain-Tuning Verification on REMUS 100

TAKEAWAY Offline verification shows that the proposed gain changes can reduce post-fault error and oscillation, making autonomous redesign plausible.

Implemented Proposal: K_p : 0.5 → 0.7, T_i : 20 → 15, T_d : 1.0 → 1.2, Max Vel: 10.0 → 9.0

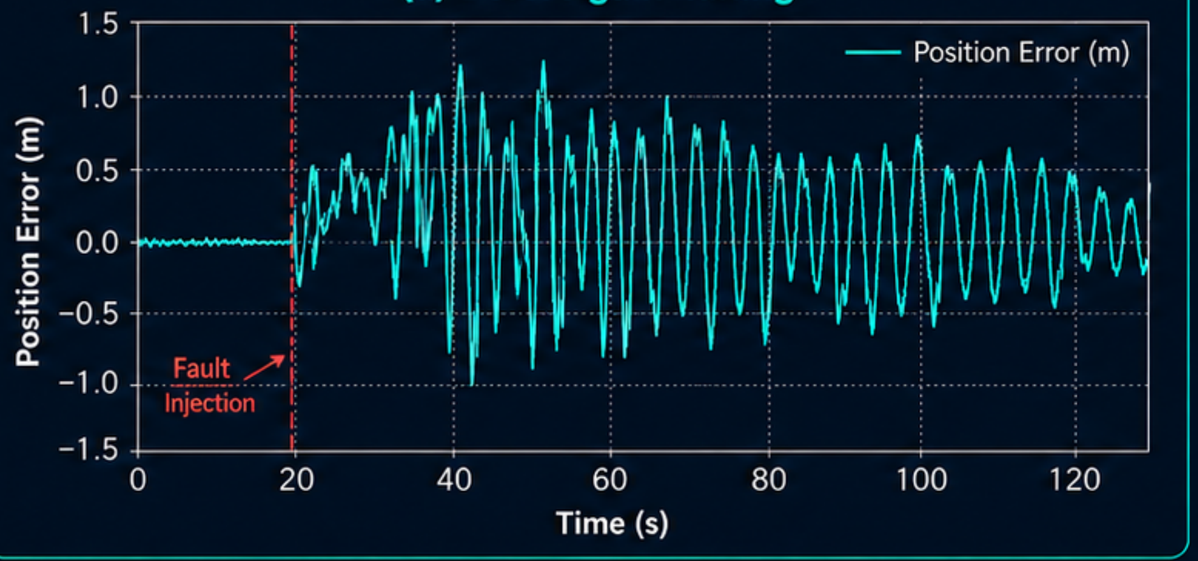

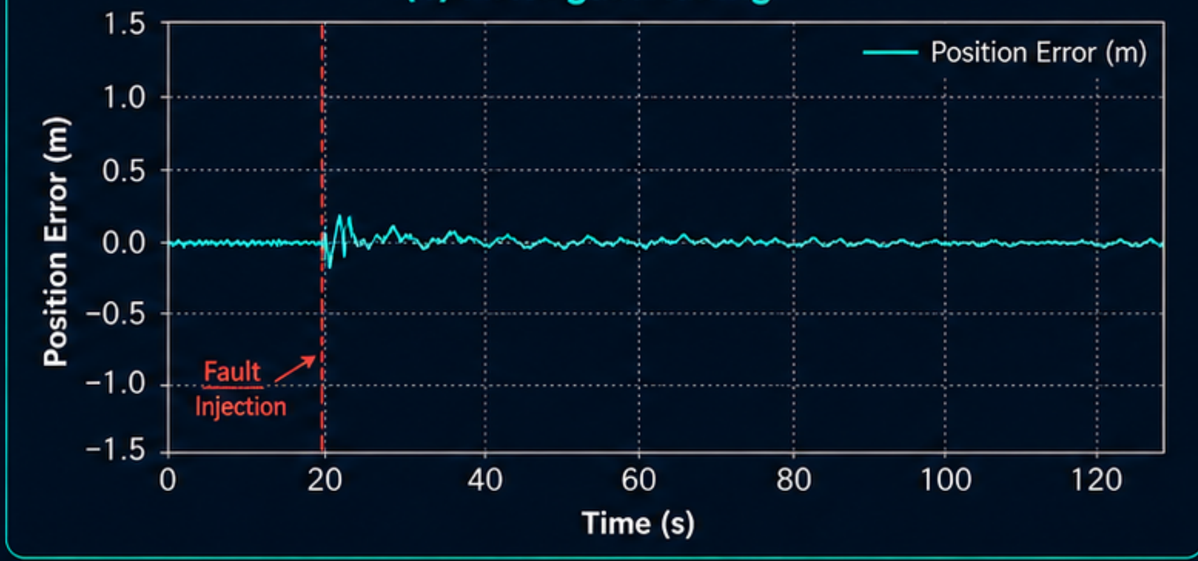

Proof of Concept: As an offline validation, these parameters reduce error magnitude and dampen post-fault oscillations.

AIVV Conclusion: Trustworthy System V&V

TAKEAWAY AIVV digitizes HITL-style V&V by anchoring LLM semantic reasoning to calibrated mathematical bounds and safe adaptation gates.

1

Neuro-Symbolic V&V

AIVV successfully digitizes the manual HITL process by anchoring LLM **semantic reasoning** to **calibrated conformal bounds.**

2

Reliability

Eliminates context-blind false positive alerts caused by dynamic maneuvering and noise, while guaranteeing **accurate true-fault validation.**

3

Safe Execution

Clone-and-promote methodology allows for safe, dynamic model **recalibration** directly in the loop.

AIVV combines calibrated math, semantic validation, and safe adaptation.

AIVV

Overall Conclusion & Future Directions

TAKEAWAY Agentic AI becomes most powerful when it is engineered as **role-specialized orchestration** around **statistical inner loops.**

1

A Unified Paradigm

By combining statistical inner loops with role-based LLM oversight, we achieve scalable adaptation (ATLAS) alongside mathematical reliability and safety (AIVV).

2

Beyond Standalone Assistants

Agentic AI is most effective when deployed as coordinating, specialized components within a structured orchestration framework.

3

Future Work

- **ATLAS:** Extending the evolutionary framework to broader scientific discovery tasks with noisy or sparse evaluation signals.
- **AIVV:** Investigating closed-loop execution of System Engineer gain-tuning proposals for fully autonomous, fault-tolerant system redesign.

Structured orchestration turns agentic AI into a reliable scientific and engineering system.

Thank You

Questions & Discussion

ATLAS: https://arxiv.org/abs/2602.02709

AIVV: https://arxiv.org/abs/2604.02478