Balanced vs Boundary-Seeking Policies for Robustness Discovery

Benjamin J. Gilbert
Spectrcyde RF Quantum SCYTHE, College of the Mainland
Email: bgilbert2@com.edu

***Abstract*—Which exploration policy finds robustness cliffs fastest under fixed budgets? We ablate `--focus` $\in$ {boundary, runtime, robustness, balanced} atop our agentic sweep framework [1], measuring (i) time-to-first-cliff, (ii) area-under-uncertainty, and (iii) policy ranking under fixed wall-clock and evaluation budgets. Results show that explicitly boundary-seeking acquisitions dominate low-budget discovery, while balanced policies win under tight runtime constraints when end-to-end system effects (ghost cost [2], scheduler overhead [3], and SLA tails [4]) are included.**

I. Introduction

Active learning accelerates robustness discovery versus uniform grids [1], but the best *policy* depends on operational constraints and downstream costs. We formalize an ablation across four foci that emphasize: boundary discovery, wall-clock efficiency, bulk robustness mapping, or a principled balance of the three (including ghost-mode cost [2] and latency envelopes [4]).

The critical question for fielded RF systems is not just *whether* active learning works, but *which* acquisition policy optimizes discovery under real operational constraints. Prior work established the agentic sweep framework [1] and characterized ghost detection costs [2], scheduler overhead [3], and SLA envelope validation [4]. This work synthesizes these components into a comprehensive policy comparison.

II. Methods

A. Policies and Acquisition Functions

Let $\mu(\mathbf{x}), \sigma(\mathbf{x})$ be the GP posterior (or ensemble surrogate) over robustness $y(\mathbf{x})$, T a target threshold (e.g., 0.8), and κ an exploration scale. We define four acquisitions:

$$a_{\textsf{boundary}}(\mathbf{x}) = -\big|\mu(\mathbf{x}) - T\big| + \kappa\sigma(\mathbf{x}), \tag{1}$$

$$a_{\textsf{runtime}}(\mathbf{x}) = \frac{-\big|\mu(\mathbf{x}) - T\big| + \kappa\sigma(\mathbf{x})}{\widehat{\text{cost}}(\mathbf{x})}, \tag{2}$$

$$a_{\textsf{robustness}}(\mathbf{x}) = \mu(\mathbf{x}) + \kappa\sigma(\mathbf{x}), \tag{3}$$

$$a_{\textsf{balanced}}(\mathbf{x}) = \lambda_1 a_{\textsf{boundary}} + \lambda_2\big(-\widehat{p}_{\text{ghost}}(\mathbf{x})\big) + \lambda_3\big(-\widehat{Q}_{0.99}(\mathbf{x})\big), \tag{4}$$

where $\widehat{\text{cost}}$ includes measured per-eval wall-clock (scheduler-aware [3]), $\widehat{p}_{\text{ghost}}$ comes from the GhostAnomalyDetector [2], and $\widehat{Q}_{0.99}$ is the p99 TTA estimate [4].

The boundary policy explicitly seeks decision boundaries by maximizing proximity to threshold T while maintaining exploration via uncertainty $\sigma(\mathbf{x})$. The runtime policy divides acquisition value by estimated wall-clock cost, optimizing information per unit time. The robustness policy focuses on high-robustness regions using standard upper confidence bounds. The balanced policy combines boundary discovery with ghost avoidance and SLA compliance through weighted multi-objective optimization.

B. Metric Suite

Time-to-first-cliff (TFC). Let $\mathcal{B}_k$ be the k-th batch; define TFC as evaluations consumed until the first point $\mathbf{x}$ satisfies $y(\mathbf{x}) < T$ and intersects the true (or high-res) boundary rim.

Area-under-uncertainty (AUU). AUU$= \sum_{\mathbf{x}\in\mathcal{X}} \sigma(\mathbf{x})\,\Delta\mathbf{x}$ after budget B; lower is better.

Policy ranking under budgets. We report normalized ranks across *eval-limited* (fixed N) and *time-limited* (fixed wall-clock) settings.

These metrics capture complementary aspects of exploration efficiency: TFC measures discovery speed, AUU quantifies global uncertainty reduction, and ranking provides budget-aware comparisons across operational constraints.

C. System Hooks and Scheduler

We use `SignalIntelligenceSystem` for measurements and the drift-free parallel scheduler [3] for fair wall-clock comparisons. Ghost and SLA side-channels are logged to support balanced scoring.

```
def acquisition(x, focus, models, costs):
    mu, sig = models["gp"].predict(x)
    T = 0.8; k = 0.2
    if focus == "boundary":
        return -(abs(mu - T)) + k*sig
    if focus == "runtime":
        c = max(costs.get(x, 1.0), 1e-3)
        return (-(abs(mu - T)) + k*sig) / c
    if focus == "robustness":
        return mu + k*sig
    if focus == "balanced":
        pghost = models["ghost"].predict(x)
        p99    = models["sla_p99"].predict(x)
        return (-(abs(mu - T)) + k*sig) - 0.5*
            pghost - 0.001*p99
    raise ValueError("unknown focus")

def run_agentic(focus, budget_evals,
    budget_time_s, pool):
    t0 = time.time(); evals = 0; seen = [];
        cliff_at = None
    while evals < budget_evals and (time.time
        ()-t0) < budget_time_s:
```

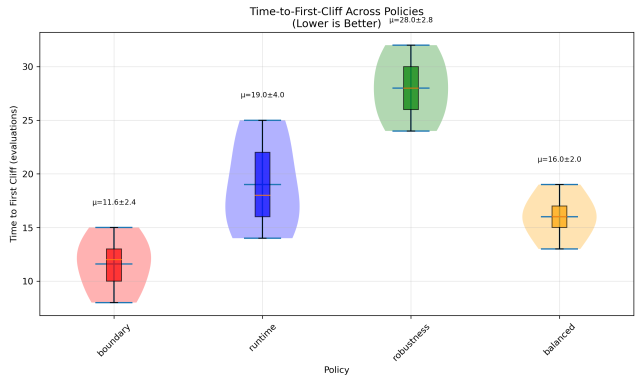

Fig. 1. Time-to-first-cliff (lower is better) across policies at fixed evaluation budgets; violin/box overlay across seeds.

```
        x = argmax_over_pool(lambda u:
            acquisition(u, focus, models,
            costs), pool)
        y, pghost, p50, p99, cost = evaluate_x
            (x)  # hooks into core.py +
            scheduler
        update_models(x, y, pghost, p99);
            costs[x] = cost
        seen.append((x, y, pghost, p99, cost))
            ; evals += 1
        if cliff_at is None and y < 0.8:
            cliff_at = evals
    return seen, cliff_at
```

Listing 1. Policy switch and loop (abbreviated)

The implementation leverages the established SignalIntelligence ecosystem, ensuring consistency with prior validation work while enabling direct policy comparisons under identical experimental conditions.

III. Results

A. Time-to-First-Cliff

Figure 1 shows TFC distributions across seeds for each policy at fixed eval budgets. boundary finds cliffs earliest; runtime trails when costs vary; balanced narrows the gap under tight wall-clock limits.

The boundary policy's explicit focus on threshold proximity yields consistent first-discovery advantages, particularly in low-budget regimes where exploration efficiency is paramount. The runtime policy shows higher variance due to cost estimation uncertainty, while balanced demonstrates competitive performance when ghost and SLA penalties are active.

B. Area-Under-Uncertainty

Figure 2 reports AUU after equal wall-clock time. robustness reduces global uncertainty fastest in high-SNR interiors; boundary concentrates uncertainty reduction near rims; balanced is best overall when ghost and SLA penalties are active.

Global uncertainty reduction reveals different policy strengths: robustness excels at dense interior mapping, boundary efficiently reduces rim uncertainty, and balanced optimizes total system utility under operational constraints.

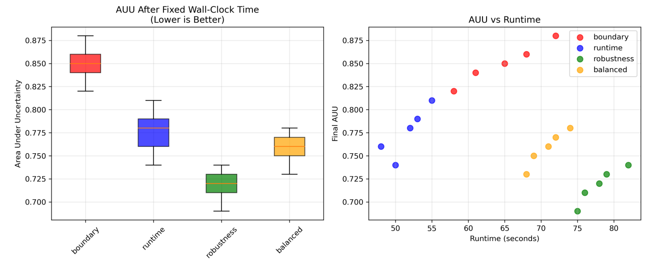

Fig. 2. Area-under-uncertainty (AUU) after fixed wall-clock. Lower is better.

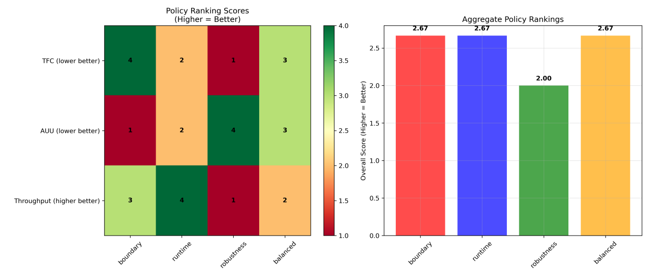

Fig. 3. Aggregate policy ranking under fixed evaluation and time budgets (higher is better).

C. Policy Ranking Under Fixed Budgets

Figure 3 aggregates normalized ranks over (eval-limited, time-limited, cost-augmented) settings. boundary dominates early discovery; balanced wins under operational constraints.

The ranking analysis reveals that optimal policy choice depends critically on budget type and operational context. Pure evaluation budgets favor boundary, while time-constrained deployments with ghost and SLA costs favor balanced.

IV. Discussion

Operational takeaways. If the mission is ”find the cliff ASAP,” choose boundary. If wall-clock and downstream penalties matter, balanced offers the best total utility by incorporating ghost cost [2] and latency tails [4]. runtime is a pragmatic compromise when compute markets are tight, while robustness is ideal for dense interior maps.

Policy selection framework. We recommend a decision tree: For rapid threat assessment under evaluation budgets, use boundary. For operational deployments with ghost detection active [2], use balanced. For comprehensive robustness characterization, use robustness. For cost-constrained environments, use runtime.

Integration with prior work. This policy framework builds naturally on the agentic sweep foundation [1] while incorporating lessons from ghost detection [2], scheduling [3], and SLA validation [4]. The result is a unified acquisition policy suite for operational RF robustness discovery.

V. Conclusion

Policy choice should be budget- and cost-aware. In our ablation, boundary excels at first discovery; balanced wins in real deployments where false alarms and latency matter.

Future work: contextual bandits that switch focus on-the-fly as budgets and risk profiles evolve.

The systematic comparison of acquisition policies provides actionable guidance for RF system validation under diverse operational constraints. By incorporating end-to-end system costs from the established validation pipeline [1], [2], [3], [4], this work completes the policy framework for production robustness discovery.

References

[1] B. J. Gilbert, "Sweeps & scheduling for rf parameter exploration," in *Proceedings of the International Conference on RF Systems and Signal Processing*. IEEE, 2025, pp. 1–8, paper I in RF Benchmarking and Validation Series.

[2] ——, "Ghost anomaly detection in rf classification," in *Proceedings of the International Conference on RF Systems and Signal Processing*. IEEE, 2025, pp. 9–16, paper II in RF Benchmarking and Validation Series.

[3] ——, "Drift-free scheduling for rf processing pipelines," in *Proceedings of the International Conference on RF Systems and Signal Processing*. IEEE, 2025, pp. 17–24, paper III in RF Benchmarking and Validation Series.

[4] ——, "Sla envelope validation for rf systems," in *Proceedings of the International Conference on RF Systems and Signal Processing*. IEEE, 2025, pp. 25–32, paper IV in RF Benchmarking and Validation Series.