

Ghost Modes and the Cost of Over-Recovery in RF Demodulation

Benjamin J. Gilbert
RF Quantum SCYTHE Research Group
College of the Mainland
Texas City, TX 77590
Email: bgilbert2@com.edu
ORCID: 0009-0006-2298-6538

September 23, 2025

Abstract

Ghost modes—false positives produced by RF demodulators—carry an operational cost. They waste computational resources, trigger spurious alarms and can mask true signals. In real-time systems, where performance criteria must be met every cycle[2], ghost hits compete with latency budgets and true hit rates. This paper quantifies the cost of ghost modes in a synthetic RF benchmark. We map iso-ghost contours across the SNR- Δf plane, construct cost curves as a function of a ghost weight λ under a latency cap, and draw a ROC-like utility curve under varying time budgets. The results reveal hotspots where ghost hits rise sharply, illustrate how economic preferences drive optimal operating points and provide guidance on policies to minimise false positives while respecting latency caps.

1 Introduction

RF demodulation pipelines seek to recover transmitted modes (true hits) while avoiding false positives (ghost hits). Ghost modes arise when noise or misconfiguration causes the demodulator to report spurious symbols. The cost of over-recovery is multifaceted: false alarms may trigger unnecessary protocol actions, increase downstream processing load and degrade user experience. In a hard real-time setting, missing performance deadlines constitutes failure[2], so ghost hits must be weighed against latency constraints and true hit performance. For instance, flight-control systems specify response times between 16 ms and 20 ms[2]; any algorithmic overhead introduced to suppress ghosts must not violate such deadlines.

Understanding where ghost hits occur and how to trade them against benefits is critical for designing robust RF systems. Dense sweeps of the parameter space are expensive[1], but simple surrogate models can provide insight. We use a synthetic benchmark with SNR and frequency offset Δf to simulate true and ghost hit rates and analyse cost trade-offs under latency caps. We define a utility function $U = \text{true_hits} - \lambda \text{ghost_hits}$ and study how λ influences the optimal operating region. Additionally, we examine how time budgets affect the trade-off between true positive and false positive rates.

2 Methods

2.1 Synthetic Model

We model true hits and ghost hits as logistic functions of the distance from an operating point: $\text{SNR} = 10$ dB and $\Delta f = 0$ kHz. Let

$$d(\text{SNR}, \Delta f) = \sqrt{(\text{SNR} - 10)^2 + (\Delta f)^2}, \quad (1)$$

then the true hit rate is

$$\text{true_hits} = \frac{1}{1 + \exp(0.6 [d(\text{SNR}, \Delta f) - 10])}, \quad (2)$$

and the ghost hit rate is

$$\text{ghost_hits} = \frac{1}{1 + \exp(-0.4 [d(\text{SNR}, \Delta f) - 10])}. \quad (3)$$

These choices ensure that true hits are high near the centre and decay with distance, while ghost hits are low near the centre and rise at the periphery. Latency follows

$$\text{runtime}_{\text{ms}} = 30 e^{\Delta f} e^{-0.1 \text{SNR}}, \quad (4)$$

so larger frequency offsets and lower SNRs incur higher runtime.

2.2 Iso-Ghost Contours

To visualise where ghost modes occur, we generate a 50×50 grid over $\text{SNR} \in [0, 20]$ dB and $\Delta f \in [0, 4]$ kHz and compute `ghost_hits` at each point. Figure 1 shows filled contours and selected iso-lines (0.2, 0.4, 0.6, 0.8) of `ghost_hits`. Regions near high Δf and low SNR exhibit high ghost rates; these are the "ghost rims" analogous to failure rims in robustness studies.

2.3 Cost Curves and Ghost Weight λ

We define a utility function

$$U(\text{SNR}, \Delta f; \lambda) = \text{true_hits} - \lambda \text{ghost_hits} \quad (5)$$

that balances true detections against the cost of ghost modes. The parameter $\lambda \geq 0$ reflects the relative penalty of false positives; a larger λ means ghost hits are more costly. We impose a latency cap of 100 ms and evaluate two summaries over the parameter grid: the maximum utility (best operating point) and the mean utility (average over admissible points) for $\lambda \in \{0, 0.2, 0.5, 1, 2, 5\}$. Figure 2 plots these curves.

2.4 ROC-Like Utility under Time Budgets

Inspired by receiver operating characteristic (ROC) analysis, we study how time budgets affect true positive and false positive rates. For latency thresholds $T \in \{10, 30, 50, 100, 200, 500\}$ ms we consider only points with $\text{runtime}_{\text{ms}} \leq T$ and compute the average true hit rate (TPR) and ghost hit rate (FPR). The resulting TPR–FPR pairs trace a curve analogous to an ROC (Figure 3). Labels denote the corresponding latency threshold.

3 Results

3.1 Iso-Ghost Contours

Figure 1 illustrates how ghost hit rates vary across parameter space. Ghosts are negligible near the operating point but rise rapidly for frequency offsets above 2 kHz and SNR below 5 dB. Iso-lines delineate ghost probability levels, enabling practitioners to avoid regions where ghost hits exceed tolerable thresholds (e.g., 0.4).




Figure 1: Iso-ghost contours over the SNR– Δf plane. Filled colours show the ghost hit rate; white contour lines mark levels 0.2, 0.4, 0.6 and 0.8. Ghosts are concentrated in bands of high frequency offset and low SNR.

3.2 Utility vs Ghost Weight λ

Figure 2 plots the maximum and mean utility under runtime ≤ 100 ms for different values of λ . When ghost hits are inexpensive ($\lambda = 0$), maximum utility occurs at the centre, where true hits are high; the average utility is also positive. As λ increases, the penalty for ghost hits reduces utility. For $\lambda = 2$ and above, the maximum and mean utilities drop significantly, reflecting a shift towards operating points with lower ghost rates (e.g., lower Δf or higher SNR). The curves provide a quantitative mapping from economic preferences (ghost cost) to optimal configurations.

3.3 ROC-Like Utility under Time Budgets

Figure 3 presents an ROC-like curve showing the trade-off between true positive rate (TPR) and false positive rate (FPR) as the latency budget increases. With a strict 10 ms budget, runtime constraints force the system to operate at high SNR and low Δf , leading to low ghost hits but also reduced true hits. As the budget relaxes to 50–100 ms, both TPR and FPR increase. At very




Figure 2: Utility as a function of ghost weight λ under a latency cap of 100 ms. The blue curve shows the maximum utility across admissible operating points, while the orange curve shows the mean utility. Higher λ penalises ghost hits and reduces the achievable utility.

lax budgets (200–500 ms), the system explores regions with high ghost rates and achieves only moderate gains in TPR. This curve helps select a latency cap that balances detection performance against false positives and computational cost.

4 Discussion

Our synthetic analysis highlights the dual role of latency and ghost penalties in configuring RF demodulators. Iso-ghost contours reveal parameter regimes that are inherently prone to false positives. Designers can avoid these by enforcing caps on frequency offset and maintaining adequate SNR. The utility curves demonstrate how the relative cost of ghost hits (λ) influences the optimal operating region: when ghosts are costly, one should favour regions with low ghost rates even if true hits are slightly reduced. The ROC-like analysis emphasises that relaxing latency budgets allows access to regions with higher detection rates but at the expense of increased false positives and computational cost. Such curves can inform service level agreements for real-time RF systems[2].

While the synthetic model is simplistic, the methodology generalises to higher-dimensional spaces and more realistic performance models. Combining ghost hit maps with latency maps and true hit maps facilitates multi-objective optimisation of RF pipelines. Future work will incorporate economic costs of missed detections, extend the analysis to adaptive and Bayesian decision thresholds, and validate policies on hardware.




Figure 3: ROC-like curve showing the average true positive rate (TPR) versus false positive rate (FPR) for different latency budgets. Points are annotated with their latency thresholds in milliseconds. Looser budgets allow higher TPR but also incur higher FPR.

5 Conclusion

Ghost modes impose tangible costs in RF demodulation pipelines. This paper quantified those costs using a synthetic benchmark and proposed a utility framework to balance true hits against ghost hits under latency caps. Iso-ghost contours identify danger zones in parameter space, cost curves translate economic preferences into operating choices, and ROC-like analysis reveals trade-offs between detection and false positives across time budgets. These tools can guide the design of demodulation policies that minimise false positives without compromising timeliness.

References

- [1] Robert B. Gramacy and Herbert K. H. Lee. Bayesian treed gaussian process models with application to computer modeling. *Journal of the American Statistical Association*, 99(467):1162–1178, 2004.
- [2] N. Hillary. Measuring performance for real-time systems. White Paper GRNTEEPFRMNCWP, Freescale Semiconductor, 2005. Discusses hard real-time system performance requirements, including deadlines that must be met every cycle and response time limits of 16-20 ms in flight control systems.