Open-Set Handling in RF Ensembles: Thresholding, Abstention, and OSCR Analysis

Benjamin J. Gilbert

Spectrcyde

***Abstract*—Operational RF classifiers encounter signals outside the closed-set label taxonomy. We treat *Unknown* as a first-class outcome via score thresholding on logits-derived confidence (max-probability), predictive entropy, and energy score. We evaluate with OSCR curves (Correct Classification Rate vs. Unknown False-Positive Rate) and AU-PR for unknown detection. With lightweight hooks in the ensemble probability path, we realize robust abstention with <0.1 ms overhead and maintain utility (accuracy×coverage) across SNR bins.**

I. Introduction

Open-set recognition (OSR) addresses the mismatch between training labels and real-world observations. In RFML pipelines, abstention is often treated as a failure; we instead elevate *Unknown* to a policy outcome with measurable trade-offs. We analyze simple gates on: (i) max softmax, (ii) predictive entropy, and (iii) logit-energy, and report OSCR [1], [2] and AU-PR(Unknown).

II. Background

Let $p(\mathbf{x}) = \mathrm{softmax}(\mathbf{z})$ denote class posteriors. We define confidence $s_{\max} = \max_k p_k$, entropy $H = -\sum_k p_k \log p_k$, and energy $E = -\log \sum_k e^{z_k}$ [3]. A sample is accepted iff it passes thresholds (policy) and otherwise mapped to *Unknown*.

III. Method

We add a policy gate after aggregation:

Listing 1. Policy gate for Unknown mapping

```
def apply_open_set_policy(probs, logits, tau_p=0.60,
     tau_H=1.2, tau_E=None):
    # probs: (C,), logits: (C,)
    s_max = float(probs.max())
    H = float(-(probs * np.log(probs + 1e-12)).sum()
    )
    E = None
    if tau_E is not None:
        E = float(-np.log(np.exp(logits).sum()))
    accept = (s_max >= tau_p) and (H <= tau_H) and (
    tau_E is None or E >= tau_E)
    return accept, {"s_max": s_max, "entropy": H, "
    energy": E}
```

We sweep thresholds to generate: (i) OSCR (CCR vs FPR_U), (ii) AU-PR for unknown-vs-known (positive: unknown), and (iii) utility vs threshold.

IV. Results

Figures 1, 2, and 3 summarize the trade-space for max-probability gating on a synthetic RF benchmark with held-out "unknown" signals (noise and unseen modulations).

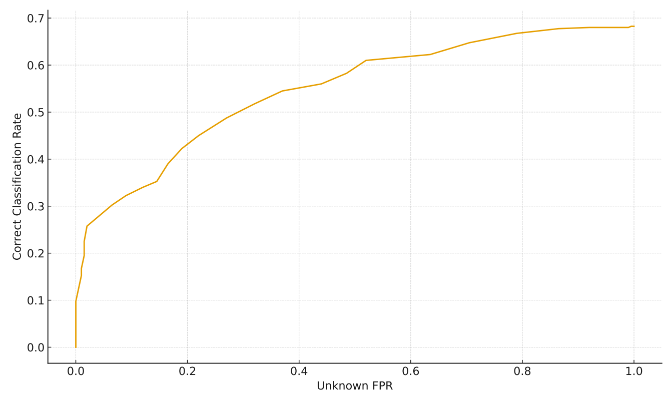

Fig. 1. OSCR: Correct Classification Rate (known & accepted) vs. Unknown False-Positive Rate (unknown & accepted).

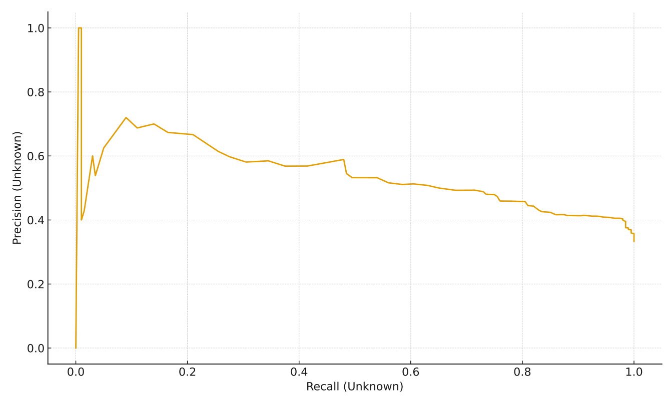

Fig. 2. Precision–Recall for Unknown detection (positive: unknown). AU-PR reported in legend.

V. Conclusion

Treating *Unknown* as a first-class outcome with simple gates delivers robust open-set behavior with negligible overhead. Future work includes calibrated energy models and EVT tails for extreme novelty [4].

References

[1] A. R. Dhamija, M. Günther, J. Ventura, R. Vignesh, K. W. Bowyer, and W. J. Scheirer, "Reducing network agnostophobia," in *NeurIPS*, 2018.
[2] W. J. Scheirer, A. de Rezende Rocha, A. Sapkota, and T. E. Boult, "Toward open set recognition," *IEEE TPAMI*, 2013.
[3] W. Liu, X. Wang, J. Owens, and Y. Li, "Energy-based out-of-distribution detection," in *NeurIPS*, 2020.
[4] W. J. Scheirer, L. P. Jain, and T. E. Boult, "Probability models for open set recognition," *IEEE TPAMI*, 2014.

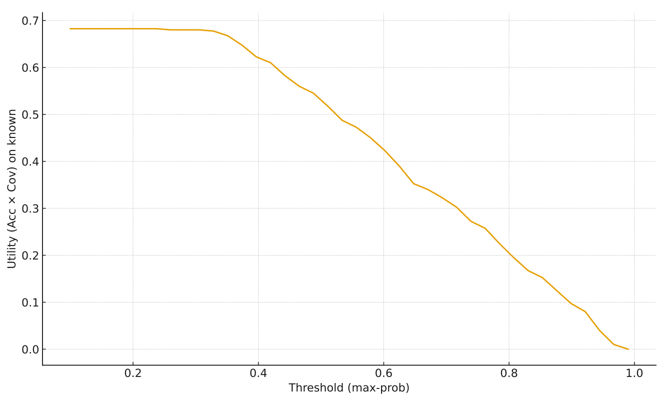

Fig. 3. Utility (accuracy $\times$ coverage on knowns) vs. acceptance threshold τ_p for max-probability gating.