SAME IDENTITY, RELAXED EVIDENCE: MARGINCELL FOR RF DOMAIN GENERALIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Radio-frequency fingerprint identification (RFFI) provides lightweight physical-layer authentication solution in zero-trust wireless networks. In practical deployments, however, its recognition performance degrades sharply under unseen environments. Existing robust RFFI schemes follow a suppress-and-stabilize paradigm that seeks to eliminate environment-induced variation, leaving residual-shift tolerance unmodeled. To tackle this, we reformulate source-only RF domain generalization via decision-evidence geometry. We prove that true-vs-rival margins form decision-complete evidence coordinates, under which each identity occupies a relaxed region of decision-equivalent states rather than a single point. This geometry yields a maximal reserve-preserving evidence cell and an exact bottleneck robustness radius, while revealing cross-entropy's intrinsic contraction toward point evidence. Building upon this insight, we propose MarginCell, a source-only regularizer that combines a bounded dispersion reward for releasing decision-safe variation with a distance-to-cell barrier against reserve-eroding deviations. Experiments show consistent improvements in cross-environment recognition and complementary gains with diverse methods.

1 INTRODUCTION

Next-generation wireless networks face growing challenges in establishing lightweight and reliable trust among heterogeneous devices (Chen et al., 2026). Traditional cryptographic authentication relies on complex credential management, incurring additional overhead and potential security vulnerabilities. Radio-frequency fingerprint identification (RFFI) has emerged as a compelling non-cryptographic physical-layer solution (Li et al., 2026). It identifies transmitters from subtle device-specific signal distortions induced by inherent hardware imperfections (known as RF fingerprint (RFF)) (Cekic et al., 2021). This hardware-rooted and lightweight paradigm offers crucial support for secure and efficient zero-trust access control in evolving wireless networks.

Deep learning has been widely adopted to improve RFFI accuracy (Xie et al., 2025). However, in practical deployment, time-varying channels and heterogeneous acquisition conditions severely shift RF observations, causing models trained in source domains to degrade markedly under unseen conditions (Ma et al., 2025). This undermines the sustained reliability of RFFI, making robustness to unseen environments a central challenge for practical deployment. We therefore focus on *source-only RF domain generalization*, where training has access only to labeled observations from limited source environments, with no target-domain samples or statistics available.

Existing robust RFFI methods mainly follow a *suppress-and-stabilize* paradigm, i.e., they primarily attempt to mitigate environmental variation via physical compensation (Zha et al., 2025) or invariant representation learning (Wang et al., 2026). However, suppression learned from observed variations fails to fully eliminate residual shifts in unseen environments, leaving decision-level tolerance an unavoidable yet unexplored challenge. Generic domain generalization (DG) methods do not close this gap either. They rely on proxy objectives (e.g., representation invariance or risk minimization), without explicitly modeling the class competition which ultimately determines whether an identity is correctly recognized. For fine-grained RFFs characterized by subtle hardware differences, identity preservation hinges not on how far the representation moves overall, but on the margin between the true identity and its closest rivals. As shown in Figure 1, the same identity can exhibit markedly different rival-wise evidence states across environments while preserving the same decision, revealing

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

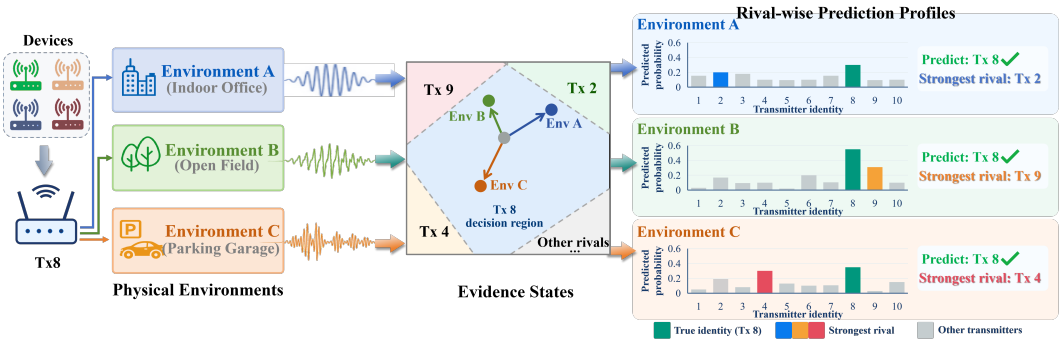


Figure 1: Environmental shifts reshape rival competition. The strongest rival changes while the predicted identity remains stable, revealing a range of decision-consistent evidence states.

a key insight: *decision consistency does not require evidence equality*. Instead, each identity occupies a decision-equivalent evidence region capable of absorbing environmental variations, with risk emerging only when the weakest-rival reserve is exhausted. What is therefore needed is an explicit *decision-evidence geometry* to delineate this admissible freedom and its reserve boundary. Yet, at fixed class-mean margins, cross-entropy (CE) favors point evidence and contracts the freedom that should be preserved for cross-domain tolerance. This leads to the central dilemma: *how can we release evidence freedom without eroding rival-wise decision reserve?*

We resolve this dilemma by shifting RF domain generalization from proxy-level objectives to explicit decision-evidence geometry. We prove that class-relative true-vs-rival pairwise margins form decision-complete evidence coordinates, converting implicit class competition into a measurable geometry directly aligned with rival decision boundaries. Based on this formulation, we derive the maximal reserve-preserving evidence cell, identify the threshold required to overcome CE-induced contraction, and prove that the bottleneck relative margin exactly determines the robustness radius to bounded evidence shifts. Guided by these insights, we propose **MarginCell**, a source-only regularizer operating in decision-evidence space under a release-and-bound principle: a bounded dispersion reward releases decision-safe evidence variation, while a distance-to-cell barrier constrains reserve-eroding deviations.

In summary, our main contributions are as follows:

- We reframe source-only RF domain generalization through the lens of decision-evidence geometry. By establishing true-vs-rival margins as decision-complete coordinates, we characterize the maximal admissible evidence cell and derive its exact robustness radius.
- We propose MarginCell, a release-and-bound regularizer: (a) a bounded dispersion reward that counteracts CE-induced contraction to release decision-safe evidence freedom; and (b) a distance-to-cell barrier that restricts reserve-eroding deviations.
- We evaluate the proposed scheme on multiple public benchmarks. The results demonstrate excellent cross-environment robustness and superior gains across diverse methods.

2 Related Work

Robust RFFI. RFFI has attracted growing attention for lightweight, intrinsic physical-layer device authentication (Li et al., 2026; Han et al., 2025), yet RF observations are sensitive to dynamic channels and acquisition conditions, posing a key bottleneck to reliable real-world deployment. Existing robust RFFI methods mitigate acquisition shifts through *signal-side physical suppression*, (e.g., channel calibration (Zha et al., 2025), spectral or hardware cues (Xie et al., 2025)) or *representation-side invariance*, including adversarial or federated feature alignment (Li et al., 2025b; Wang et al., 2026), and domain-invariant generative modeling (DIVA; Shi et al., 2026). Both of them follow a *suppress-and-stabilize* paradigm by modeling or aligning observed variations, which leaves residual shifts under unseen conditions. This raises a more fundamental question: *how much residual evidence shift can a decision tolerate when shifts cannot be fully removed?* MarginCell therefore reframes robust RFFI from *shift suppression* to *shift tolerance*, using decision reserve to quantify this tolerance and release-and-bound regularization to control evidence freedom.

Domain Generalization. DG learns predictors from source domains that transfer to unseen domains. Existing methods mainly pursue *domain invariance* through distribution alignment with MMD (Muandet et al., 2013) and domain-reliance unlearning with IU (Zhao et al., 2026); *risk robustness* through GroupDRO (Sagawa et al., 2019) and V-REx (Krueger et al., 2021); and *optimization and specialization* through flat-minima search with SFT (Li et al., 2025a), gradient guidance with GGA (Ballas & Diou, 2025), and pattern specialization with TFPS (Sun et al., 2025). However, these methods operate on proxy quantities (representations, risks, or gradients), rather than modeling the class-competitive evidence that determines prediction, leaving no explicit geometry to distinguish decision-preserving variation from reserve-eroding deviations. MarginCell instead moves beyond *proxy-level regularization* to *decision-evidence geometry*, explicitly constructing decision-complete pairwise margins to represent competitive evidence and its decision-preserving reserve.

Margin-Based Regularization. Margin-based methods reshape decision geometry through fixed angular margins (SphereFace; Liu et al., 2017; ArcFace; Deng et al., 2019) or adaptive/class-sensitive margins (CLOC; Pitawela et al., 2025; MR^2; Zhu et al., 2026). Related output regularizers, such as Label Smoothing (Szegedy et al., 2016), LogitNorm (Wei et al., 2022), and MbLS (Liu et al., 2022), instead regulate confidence, logit scale, or margin slack. However, existing methods treat margins or logits as quantities to shape, rather than as *decision-complete evidence coordinates* admitting decision-equivalent freedom. Crucially, this freedom is not self-realizing: under fixed class-mean margins, cross-entropy contracts evidence toward a point state. MarginCell therefore shifts from *margin-as-constraint* to *margin-as-evidence*, explicitly releasing decision-safe variation.

3 Theoretical Motivation

In this section, we isolate a residual question at the classifier output: *once the identity is correctly preserved, how much variation should remain in the evidence supporting that identity?* To answer this, we formulate cross-environment recognition in a class-relative margin space, where the same identity can be supported by multiple decision-consistent evidence states. Through a progressive theoretical analysis, we establish the maximal geometric allowance for these states, the optimization force required to release them, and the risk constraints they must satisfy.

3.1 Class-Relative Decision Evidence

Let a K-class classifier output logits $z_\theta(x) \in \mathbb{R}^K$, and let y be the true identity. For each competitor $k \neq y$, define the pairwise margin $m_{yk}(x) = z_y(x) - z_k(x) = \log \frac{p_y(x)}{p_k(x)}$, where $p = \mathrm{softmax}(z)$. Thus, m_{yk} measures the model's pairwise log-evidence for y against competitor k.

Proposition 3.1 (Decision-Complete Margins)**.** *For a fixed candidate identity y, the margin vector $m_y = (m_{yk})_{k\neq y} \in \mathbb{R}^{K-1}$ uniquely determines the softmax probabilities and the centered logits. Moreover, if ties are counted as errors, $\hat{y} = y$ if and only if $m_{yk} > 0$ for every $k \neq y$.*

The full proof is in Appendix A.1. Intuitively, softmax probabilities and predictions depend only on relative logit differences, fully determined by the $K-1$ pairwise margins to class y. Thus, pairwise margins are not arbitrary hidden-space distances, but decision-complete coordinates of the classifier's competitive geometry after quotienting out the irrelevant common logit shift.

Raw margins, however, can have very different scales across class-rival pairs. Let $M_y = (M_{yk})_{k\neq y}$ denote the positive class-rival mean margins, i.e., $M_{yk} = \frac{1}{n_y}\sum_{i:y_i=y} m_{ik}$. We define the relative evidence coordinate as

$$r_{yk}(x) = \frac{m_{yk}(x) - M_{yk}}{M_{yk}}. \tag{1}$$

This maps the reference state to 0 and every pairwise decision boundary $m_{yk} = 0$ to the common location $r_{yk} = -1$. Hence the correct-decision region is

$$\mathcal{R}_y = \{r \in \mathbb{R}^{K-1} : r_k > -1,\ \forall k \neq y\}. \tag{2}$$

Thus, $r = 0$ is merely the class reference: every $r \in \mathcal{R}_y$ yields the same identity prediction y. In this sense, *identity consistency does not require evidence equality*, a principle we term as *"Same Identity, Relaxed Evidence"*. Yet $\mathcal{R}_y$ is a region with finite negative tolerance but unbounded positive freedom; the remaining question is how much variation should be retained.

3.2 MAXIMAL ADMISSIBLE CELL

We seek the largest evidence region that preserves a prescribed decision reserve with unknown variation in new environments. Specifically, let the relative evidence undergo a perturbation δ with $\|\delta\|_\infty \le \rho$, while requiring every perturbed pairwise margin to retain at least a fraction $\kappa \in (0,1)$ of its class-rival mean margin. We assume $\kappa + \rho < 1$ so that a non-degenerate admissible region exists.

Theorem 3.2 (Maximal Admissible Cell). *Let $d = K - 1$. If a centrally symmetric set $\mathcal{A} = -\mathcal{A}$ satisfies $M_{yk}(1 + r_k + \delta_k) \ge \kappa M_{yk}$ for every $r \in \mathcal{A}$, every $\|\delta\|_\infty \le \rho$, and every rival k, then*

$$\boxed{\mathcal{A} \subseteq \mathcal{C}_y(\tau) := [-\tau, \tau]^d, \qquad \tau = 1 - \kappa - \rho.} \tag{3}$$

Conversely, $\mathcal{C}_y(\tau)$ itself satisfies the reserve condition. It is therefore the largest centrally symmetric admissible region under this ℓ_∞ evidence-shift model. (Proof in Appendix A.2)

Thus, this cell is not an arbitrary heuristic box, but the maximal symmetric region preserving the prescribed decision reserve under directional uncertainty. This gives the boundary τ a direct decision-reserve interpretation, formalizing the trade-off between evidence flexibility and robustness.

3.3 COLLAPSE-DISPERSION TRANSITION

Despite the legitimate non-zero degree of freedom established for same-identity evidence, geometric allowance alone does not guarantee that standard optimization will utilize it. We therefore examine the intrinsic bias of cross-entropy (CE) at fixed class-mean margins, and characterize when the point evidence loses stability.

Fix class y and abbreviate its class-rival mean-margin vector as $M := M_y \in \mathbb{R}^d_{>0}$. Consider mean-preserving relative-evidence deviations r_i, such that $m_i = M + D_M r_i$, where $D_M = \mathrm{Diag}(M)$ and $\frac{1}{n}\sum_i r_i = 0$. Since $\phi(m) = \log(1 + \sum_k e^{-m_k})$ is strictly convex, $\frac{1}{n}\sum_i \phi(M + D_M r_i) \ge \phi(M)$, with equality if and only if $r_i = 0$ for all i. Thus, *at a fixed class-mean decision level, cross-entropy uniquely favors point evidence*.

To quantify this contraction, let $a_k = e^{-M_k}/(1+\sum_j e^{-M_j})$ denote the softmax probability assigned to rival class k at the reference state. Using the Hessian of CE with respect to the raw margins together with the mapping $m = M + D_M r$, its local curvature in the relative-evidence space is

$$C_y = D_M \left[\mathrm{Diag}(a) - aa^\top\right] D_M, \tag{4}$$

where $v^\top C_y v$ measures the local CE cost for deviating the reference along evidence direction v. We next quantify how much local anti-contraction is needed to overcome this CE-induced point stability.

Theorem 3.3 (Evidence-Release Threshold). *Under the mean-preserving constraint, consider any twice-differentiable regularizer $\Omega(R) = \frac{1}{n}\sum_i \omega(r_i)$ that is stationary at $r = 0$ and has isotropic local anti-contraction curvature $\nabla^2\omega(0) = -2\mu I$ with $\mu > 0$. Let J_α denote CE plus this regularizer with strength α. Then the critical strength separating local point stability from instability is*

$$\boxed{\alpha_c = \frac{\lambda_{\min}(C_y)}{2\mu}.} \tag{5}$$

If $\alpha < \alpha_c$, the point-evidence state $R = \mathbf{0}$ remains strictly locally stable, and otherwise if $\alpha > \alpha_c$. More generally, a unit evidence direction v is released once $2\alpha\mu > v^\top C_y v$, meaning that lower-curvature evidence modes escape point contraction first. (Proof in Appendix A.3)

The mean-preserving constraint eliminates first-order CE variations, leaving point-evidence stability governed at second order by the curvature competition $C_y - 2\alpha\mu I$. For an eigendirection v_j of C_y with eigenvalue λ_j, the release condition reduces to $\alpha > \alpha_{c,j} := \lambda_j/(2\mu)$, yielding a direction-specific quantitative prediction. Figures in Appendix C.1 further provide direct numerical support for this prediction of Theorem 3.3.

Thus, admissible evidence freedom must be actively released from CE contraction. However, Theorem 3.3 determines only *when* a direction is released, not whether the resulting deviation is safe. We next identify variations that threaten identification through the rival closest to the decision boundary.

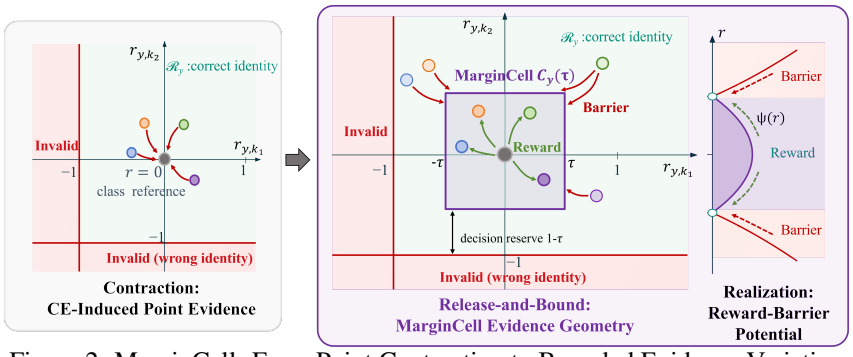


Figure 2: MarginCell: From Point Contraction to Bounded Evidence Variation.

3.4 Bottleneck-Risk Transfer

Since classification fails as soon as any true-vs-rival margin crosses zero, we define the bottleneck relative margin to capture the weakest competitor-wise reserve to the decision boundary:

$$s(x, y) = \min_{k \neq y} \frac{m_{yk}(x)}{M_{yk}} = 1 + \min_k r_{yk}(x). \tag{6}$$

Let R_T^{bal} and $\Pr_S^{\text{bal}}$ denote the class-balanced target 0-1 risk and source probability measure.

Theorem 3.4 (Bottleneck-Risk Transfer)**.** *For any $\rho \geq 0$, a sample remains strictly correct under every relative-evidence shift satisfying $\|\delta\|_\infty \leq \rho$ if and only if $s(x, y) > \rho$. Hence, $s(x, y)$ is exactly the ℓ_∞ robustness radius in relative-evidence space. Moreover, if there exists a label-preserving source-target coupling such that the class-balanced probability of an evidence displacement larger than ρ is at most η, then*

$$\boxed{R_T^{\text{bal}} \leq \eta + \Pr_S^{\text{bal}}[s(X, Y) \leq \rho]\,.} \tag{7}$$

Since $s(r) = 1 + \min_k r_k$ is 1-Lipschitz under ℓ_∞, target failure after a shift of at most ρ requires the coupled source sample to satisfy $s \leq \rho$. The complete proof is given in Appendix A.4.

This result turns evidence geometry into a risk-relevant criterion: nonzero variation is harmless, what matters is the negative tail depleting the bottleneck reserve toward a decision boundary. Accordingly, if $r \in \mathcal{C}_y(\tau)$, then $s(x, y) \geq 1 - \tau$, so the cell permits evidence variation while preserving a $1 - \tau$ bottleneck reserve. Evidence freedom should thus be released, but not at the expense of this reserve.

3.5 From Theory to MarginCell

Together, Theorems 3.2-3.4 yield a constrained principle for evidence freedom: delineate the maximal reserve-preserving cell, overcome CE-induced contraction to realize admissible variation, and prevent excursions that erode the bottleneck reserve. This naturally motivates unified reward-barrier construction (MarginCell in Section 4): release admissible evidence variation while controlling reserve-eroding deviations. *Margins define the evidence; the cell bounds its freedom.*

4 Methodology

Problem Formulation. Let $\mathcal{X} \subseteq \mathbb{C}^L$ denote the input space of length-L complex baseband RF signal sequences, and $\mathcal{Y} = \{1, \ldots, K\}$ be the fixed transmitter identity space. We consider a cross-environment DG setting, where a classifier $f_\theta : \mathcal{X} \to \mathbb{R}^K$ maps signal x to output logits $z_\theta(x)$.

During training, data is drawn from multiple labeled source environments, $\mathcal{D}_S = \bigcup_{e=1}^{E_S} \mathcal{D}_e^S$, where each $\mathcal{D}_e^S = \{(x_i^e, y_i^e)\}_{i=1}^{N_e}$ is governed by an environment distribution P_e^S. MarginCell leverages observations and identity labels, entirely independent of environment indices. Upon deployment, the model is evaluated on unseen target environments $\mathcal{D}_T = \bigcup_{j=1}^{E_T} \mathcal{D}_j^T$, whose distributions P_j^T may differ from the source domains due to variations in signal acquisition time or propagation conditions.

All environments share the same identity set $\mathcal{Y}$, while target observations and their distribution statistics are inaccessible for both training and model selection. The learning objective is to minimize the expected target risk $\mathcal{R}_T(\theta) = \mathbb{E}_{(x,y)\sim P_T}\left[\mathbb{I}(\hat{y}_\theta(x) \neq y)\right]$ with $\hat{y}_\theta(x) = \arg\max_k z_{\theta,k}(x)$.

MarginCell Overview. As established in Section 3, identity consistency permits nonzero relative-evidence variation as long as the bottleneck decision reserve is preserved. We therefore reframe cross-environment generalization as a decision-evidence boundary control problem: preserving admissible evidence variation to accommodate inherent physical variability, while preventing reserve-eroding excursions. MarginCell introduces two core mechanisms: a shared class-rival reference that places intra-class evidence variation on a relative scale (Section 4.1), and a unified reward-barrier potential that releases and bounds this variation (Section 4.2). The illustration is shown in Figure 2.

4.1 Relative Evidence Coordinates

We use the mean margin of intra-class samples as a shared reference to convert individual sample evidence into normalized relative offsets. Specifically, for a subset $\mathcal{B}_y$ of class y within a source minibatch $\mathcal{B}$ (where $n_y = |\mathcal{B}_y|$), the dynamic reference and relative coordinate for sample $i \in \mathcal{B}_y$ against competitor k are defined as:

$$M_{yk} = \frac{1}{n_y}\sum_{i\in\mathcal{B}_y} m_{ik}, \quad r_{ik} = \frac{m_{ik} - M_{yk}}{M_{yk} + \epsilon} \tag{8}$$

where $\epsilon > 0$ stabilizes normalization, and $r_{ik} < 0$ (> 0) indicates weaker (stronger) evidence than the current class-rival mean. Notably, the minibatch reference M_{yk} is computed dynamically and remains differentiable, coupling intra-class evidence through a shared scale so that the class-wise evidence geometry evolves jointly without an external prototype.

4.2 Unified Reward-Barrier Potential Function

With relative evidence coordinates established, our objective is twofold: releasing admissible evidence variations from CE-induced point contraction while controlling excessive deviations that may deplete the bottleneck reserve. To this end, we formulate a unified reward-barrier potential function:

$$\psi_{\tau,\beta}(r) = \underbrace{-\min\left(\frac{r^2}{\tau^2}, 1\right)}_{\text{Dispersion Reward}} + \underbrace{\beta\frac{[|r| - \tau]_+^2}{\tau^2}}_{\text{Boundary Barrier}} \tag{9}$$

where $[u]_+ = \max\{u, 0\}$, $\tau \in (0, 1)$ denotes the half-width of the cell $\mathcal{C}_y(\tau) = [-\tau, \tau]^{K-1}$, and $\beta > 0$ dictates the out-of-bound penalty strength. Sharing a unified boundary τ, the two terms operate separately to govern *release* and *control*:

Dispersion Reward: For $|r| < \tau$, the reward is $-r^2/\tau^2$. Featuring zero gradient and isotropic curvature $-2\tau^{-2}\mathbf{I}$ at $r = 0$, it injects the anti-contraction force required by Theorem 3.3. Superimposed with CE's anisotropic curvature, it delegates release directions to the classifier's intrinsic decision geometry: only low-curvature modes ($\lambda_j < 2\alpha\mu$) are released, while high-curvature ones remain contracted. For RFFI task, this allows a transmitter to retain cross-environment variations in specific rival-wise evidence without disrupting the global identity-discriminative structure. Beyond τ, the reward saturates and precludes unbounded margin expansion.

Boundary Barrier: Reward saturation only halts the incentive for expansion, it does not retract out-of-bounds evidence. Accordingly, the barrier imposes a quadratic penalty on the excursion distance outside the cell. This preserves in-cell freedom while controlling overshoot. In particular, it prevents excessive negative deviations from further eroding the weakest rival's bottleneck reserve, whose low-reserve tail directly enters the cross-environment risk bound given by Theorem 3.4. Consistent with the symmetric geometric budget established in Theorem 3.2, we adopt a symmetric distance-to-cell penalty rather than treating all out-of-cell states as classification errors.

4.3 Overall Objective and Inference Properties

Overall Objective. To avoid normalizing by near-zero class-rival references, we restrict the regularizer to valid rivals $\mathcal{V}_y = \{k \neq y : M_{yk} > \gamma\}$ and eligible classes $\mathcal{Y}_{\mathcal{B}}^+ = \{y : n_y \geq 2,\ |\mathcal{V}_y| > 0\}$.The validity mask is detached from backpropagation, while the retained M_{yk} remain differentiable.

Table 1: MarginCell performance comparison with other methods across datasets. (**Bold**: Best; Underline: Second-best).

Type	Method	WiSig (2→2 d.)			UCLA-CORES (3→2 d.)			LoRa (6→2 env.)			*Overall Avg.*		
		Acc	WEnv	$\text{TIA}_{@20}$	Acc	WEnv	$\text{TIA}_{@20}$	Acc	WEnv	$\text{TIA}_{@20}$	Acc	WEnv	$\text{TIA}_{@20}$
Reference	ERM	83.50	83.00	51.00	65.37	60.94	31.61	87.93	82.60	63.00	78.93	75.51	48.54
DG Methods	TFPS	91.75	86.30	76.35	70.82	67.36	32.84	79.65	66.58	54.63	80.74	73.41	54.61
	GroupDRO	88.15	86.73	64.80	69.24	65.99	44.77	89.26	83.10	62.19	82.22	78.61	57.25
	MMD-DG	89.02	86.47	67.40	71.84	68.11	46.65	92.75	88.85	76.75	84.54	81.14	63.60
	V-REx	89.92	88.13	70.10	69.95	68.91	41.87	90.11	84.88	65.44	83.33	80.64	59.14
Avg. (DG)		89.71	86.91	69.66	70.46	67.59	41.53	87.94	80.85	64.75	82.70	78.45	58.65
Robust RFFI	CDSI-RFF	89.92	83.03	70.05	70.23	68.20	21.41	39.80	31.15	14.00	66.65	60.79	35.15
	CLPS	94.33	89.47	83.00	73.33	70.91	42.64	88.18	83.05	58.44	85.28	81.14	61.36
	CIS	93.22	**91.43**	79.95	68.20	67.32	29.32	89.05	82.75	63.44	83.49	80.50	57.57
	CKR-RFF	88.90	85.83	67.00	70.75	68.38	36.68	90.94	84.13	71.00	83.53	79.45	58.23
Avg. (Robust RFFI)		91.59	87.44	75.00	70.63	68.70	32.51	76.99	70.27	51.72	79.74	75.47	53.08
Regularization	LS	90.90	87.93	73.00	73.55	71.85	47.83	91.88	87.18	69.88	85.44	82.32	63.57
	MbLS	89.75	87.10	69.45	72.58	70.95	47.78	91.46	87.70	72.13	84.60	81.92	63.12
	LogitNorm	91.38	88.77	74.55	70.78	67.25	41.22	91.63	84.93	72.00	84.60	80.32	62.59
Avg. (Reg.)		90.68	87.93	72.33	72.30	70.02	45.61	91.66	86.60	71.34	84.88	81.52	63.09
-	**MarginCell**	**95.15**	91.33	**85.70**	**76.74**	**75.17**	**53.44**	**93.34**	**89.23**	**79.19**	**88.41**	**85.24**	**72.78**
Improv.	*vs. Best*	+0.82	-0.10	**+2.70**	+3.19	+3.32	**+5.61**	+0.59	+0.38	**+2.44**	+2.97	+2.92	**+9.18**
	vs. Mean	+4.49	+3.95	**+13.37**	+5.72	+6.51	**+14.08**	+8.37	+10.66	**+17.38**	+6.19	+7.04	**+14.94**

For a minibatch $\mathcal{B}$, MarginCell aggregates all valid sample-rival coordinates as:

$$\mathcal{R}_{\text{MC}}(\mathcal{B}) = \frac{1}{|\mathcal{Y}_{\mathcal{B}}^{+}|} \sum_{y \in \mathcal{Y}_{\mathcal{B}}^{+}} \frac{1}{n_y |\mathcal{V}_y|} \sum_{i \in \mathcal{B}_y} \sum_{k \in \mathcal{V}_y} \psi_{\tau,\beta}(r_{ik}) \tag{10}$$

If no class is eligible, i.e., $|\mathcal{Y}_{\mathcal{B}}^{+}| = 0$, we set $\mathcal{R}_{\text{MC}}(\mathcal{B}) = 0$, leaving the update driven solely by cross-entropy. The complete training objective is:

$$\mathcal{L}_{\text{total}} = \frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \ell_{\text{CE}}(z_\theta(x_i), y_i) + \lambda w_t \, \mathcal{R}_{\text{MC}}(\mathcal{B}) \tag{11}$$

where λ controls the overall regularization strength and $w_t \in [0, 1]$ progressively activates MarginCell after the initial CE-dominated stage.

Inference and Plug-in Compatibility. At inference, neither M_{yk}, r_{ik} nor $\mathcal{R}_{\text{MC}}$ is required. The prediction follows the standard rule:

$$\hat{y}_\theta(x) = \arg\max_k z_{\theta,k}(x) \tag{12}$$

MarginCell therefore leaves the deployment pipeline unchanged and introduces no additional inference branch. During training, it requires only source labels and classifier logits. Therefore, it also can be added as a plug-in module to existing source-only DG frameworks. While conventional DG methods focus on macro-level cross-environment alignment, MarginCell operates orthogonally on intra-identity decision-evidence geometry, establishing a natural synergy with existing DG targets.

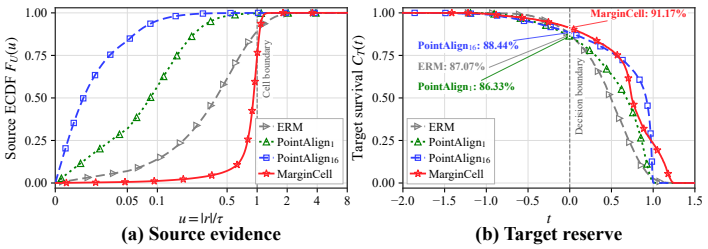
5 Experiments

5.1 Experimental Setup

Benchmarks and Metrics. We evaluate MarginCell on three representative public RFFI benchmarks with distinct acquisition protocols: WiSig (Hanna et al., 2022), UCLA-CORES (Hanna et al., 2020), and LoRa (Shen et al., 2022), with environments partitioned into seen source and unseen target sets under a unified source-only DG protocol. To evaluate cross-environment performance, we utilize target accuracy (Acc) for overall performance, worst-environment accuracy (WEnv) for environment-level robustness, and tail-identity accuracy ($\text{TIA}_{@20}$) for hard-identity robustness. Detailed dataset statistics with preprocessing configurations and source/target environment splits, and evaluation metrics are provided in Appendix B.1 and Appendix B.2, respectively. Our code is available at `https://anonymous.4open.science/r/cm_RFFI`.

Table 2: Plug-in performance of MarginCell across baselines and datasets. "+MC" refers to results with MarginCell. $\Delta > 0$ indicates absolute gains. (**Bold**: Best; Underline: Second-best).

Base Objective	WiSig				UCLA-CORES				LoRa			
	+MC Acc	ΔAcc	+MC WEnv	ΔWEnv	+MC Acc	ΔAcc	+MC WEnv	ΔWEnv	+MC Acc	ΔAcc	+MC WEnv	ΔWEnv
GroupDRO	**97.92**	**+9.77**	**96.17**	**+9.44**	77.38	**+8.14**	**75.85**	**+9.86**	93.12	+3.86	88.80	+5.70
V-REx	91.17	+1.25	89.33	+1.20	70.29	+0.34	70.28	+1.37	93.27	+3.16	89.05	+4.17
MMD-DG	97.17	+8.15	94.50	+8.03	74.25	+2.41	73.06	+4.95	**95.60**	+2.85	**92.55**	+3.70
CLPS	96.17	+1.84	92.33	+2.86	**79.68**	+6.35	74.19	+3.28	93.88	**+5.70**	89.10	**+6.05**
CIS	95.22	+2.00	92.17	+0.74	71.17	+2.97	70.90	+3.58	89.30	+0.25	87.85	+5.10



(a) Source evidence (b) Target reserve

Figure 3: Source Evidence Dispersion and Target Bottleneck Reserve.

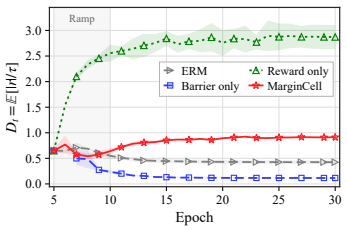


Figure 4: Training Dynamics of Relative Evidence.

Implementation and Baselines. We employ a uniform ResNet-18 backbone and compare MarginCell against ERM and representative baselines from three groups: (i) DG methods, including GroupDRO (Sagawa et al., 2019), MMD-DG (Muandet et al., 2013), V-REx (Krueger et al., 2021), and TFPS (Sun et al., 2025); (ii) robust RFFI methods, including CDSI-RFF (Zhao et al., 2024), CLPS (Tang et al., 2024), CIS (Shen et al., 2022), and CKR-RFF (Zha et al., 2025); and (iii) output regularizers, including Label Smoothing (Szegedy et al., 2016), MbLS (Liu et al., 2022), and Logit-Norm (Wei et al., 2022). Detailed baseline configurations and implementation with hyperparameter settings are provided in Appendix B.3 and Appendix B.4, respectively.

5.2 Comparison with Other Methods

Overall performance. We first compare MarginCell with SOTA DG, robust RFFI, and output-regularization baselines. As shown in Table 1, MarginCell achieves SOTA cross-environment performance across all benchmarks, with >95% accuracy on WiSig. Notably, the average gain enlarges to +9.18% on $\text{TIA}_{@20}$, consistent with MarginCell's rival-wise design that regularizes true-vs-rival evidence and protects the weakest-rival reserve. Multi-seed mean ± std results are in Appendix C.2.

In contrast, existing methods remain vulnerable to unseen residual class competition, particularly for hard identities: CIS drops from 79.95% $\text{TIA}_{@20}$ on WiSig to 29.32% on UCLA-CORES, while CDSI-RFF falls to 14.00% on LoRa. This instability exposes the limitation of the prevailing suppress-and-stabilize paradigm: suppressing observed variation does not explicitly model and enhance decision tolerance to residual evidence shifts. Generic DG and output regularizers likewise leave this gap unresolved, acting on proxy quantities or global outputs rather than explicitly protecting the weakest-rival reserve (e.g., GroupDRO with 57.25% $\text{TIA}_{@20}$ and all output regularizers below 48% on UCLA-CORES). MarginCell instead directly regulates true-vs-rival evidence and protects the bottleneck reserve, yielding consistent superior cross-environment robustness.

Plug-in Compatibility. Integrating MarginCell into existing DG and RFFI methods highlights its strong compatibility. As shown in Table 2, MarginCell consistently improves Acc and WEnv across datasets, with gains up to +9.77% on WiSig and +6.05% on LoRa. The benefit also extends to hard identities, with +12.72% $\text{TIA}_{@20}$ on average (full results in Appendix C.3). These gains highlight MarginCell's potential to complement diverse objectives by extending optimization into decision-evidence space, releasing bounded within-identity freedom while preserving the bottleneck reserve.

5.3 Evidence Geometry and Dynamics Analysis

Geometry of Bounded Freedom. To investigate if stricter evidence consistency yields better generalization, we construct PointAlign controls with varying strengths to pull intra-identity relative

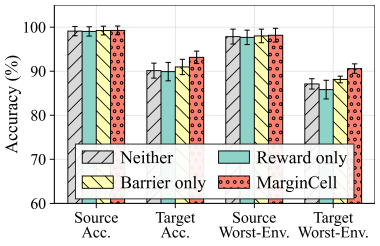


Figure 5: Ablation on Reward and Barrier Components.

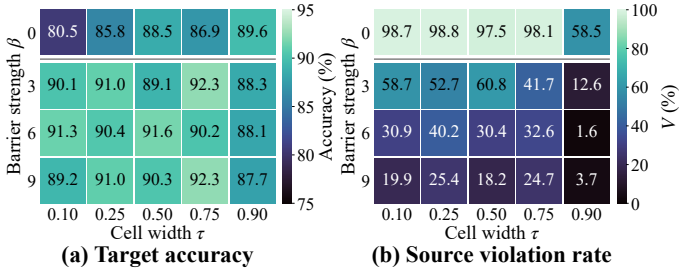


Figure 6: Joint Sensitivity to Cell Width and Barrier Strength.

evidence toward the class reference $r = 0$. Figure 3(a) shows the empirical cumulative distribution function (ECDF) of normalized source deviation $u = |r|/\tau$, where smaller u indicates stronger concentration and $u = 1$ marks the cell boundary. Stronger PointAlign shifts evidence toward $u = 0$, with PointAlign$_{16}$ producing the tightest concentration; MarginCell instead forms a controlled soft-boundary regime, occupying the available intra-cell freedom with limited overshoot.

Yet tighter point compression fails to benefit target decisions. Figure 3(b) evaluates target reserve survival $C_T(t) = \Pr[s > t]$, i.e., the fraction of samples with reserve above t, with $t = 0$ as the decision boundary. Under shifts, the highly collapsed PointAlign$_{16}$ is more vulnerable, plummeting to only 88.44% at $t = 0$, while MarginCell's bounded structural freedom yields stronger target robustness, with 91.17% survival at $t = 0$ and substantial positive reserve for $t > 0$. Appendix C.4 further clearly visualizes this source-target transition in the decision-evidence plane.

Release-and-Bound Dynamics. Figure 4 tracks the cell-normalized evidence dispersion $D_t = \mathbb{E}[|r|/\tau]$, illustrating how release-and-bound dynamics form bounded freedom during training. Once the regularization ramp begins, ERM exhibits persistent CE-induced contraction, further amplified by Barrier-only; conversely, Reward-only rapidly pushes evidence dispersion beyond the cell scale. By synergizing these opposing forces, MarginCell stabilizes at a finite, non-zero dispersion sweet spot ($D_t \approx 0.9$), avoiding both excessive contraction and uncontrolled expansion.

5.4 Component Ablation

Figure 5 reports averaged component ablations across benchmarks. Source performance remains stable, while removing either component degrades target Acc and WEnv, indicating improved transfer rather than source fitting. Reward-only releases CE-contracted evidence freedom but lacks boundary control (-5% WEnv); Barrier-only suppresses unsafe excursions but under-utilizes admissible variation (-3% Acc). Their combination yields peak target performance, supporting the complementary roles of evidence release and boundary control. Results across datasets are detailed in Appendix C.5.

5.5 Hyperparameter Sensitivity

In Figure 6, we investigate how the cell width τ and barrier strength β affect target accuracy and the violation rate $V = \Pr(|r| > \tau)$. MarginCell maintains strong accuracy over a broad (τ, β) range. Increasing β generally suppresses violations, whereas overly wide cells can yield low V (e.g., $\tau = 0.9$) without accuracy gain. This shows that MarginCell benefits not from simply minimizing violations, but from balancing τ-controlled evidence freedom with β-controlled boundary control.

6 Conclusion

In this paper, we propose a novel DG framework for RFFI task, offering a new perspective that shifts from suppress-and-stabilize to proactive decision-evidence geometric design. By explicitly modeling tolerance to residual evidence shifts and delineating its safety boundary via decision-equivalent geometry, we establish cross-environment robustness requires not evidence consistency, but bounded decision-equivalent freedom with sufficient weakest-rival reserve. Guided by this insight, we propose MarginCell, with a release-and-bound mechanism to release decision-safe evidence variation while constraining reserve-eroding deviations. Experiments validate its superior cross-environment robustness and broad compatibility with existing methods, especially for hard identities. We hope this work can inspire further exploration of decision-evidence geometry in DG research.

AI USE STATEMENT

In this work, we used generative AI tools to help polish the manuscript. We did not use AI for other purposes. All AI-assisted edits were reviewed by the authors. We take full responsibility for the final content of this work.

REPRODUCIBILITY STATEMENT

We describe our experimental setup, benchmarks, and evaluation metrics in Section 5.1. Detailed dataset splits and preprocessing are provided in Appendix B.1, evaluation metrics in Appendix B.2, baseline configurations in Appendix B.3, and implementation and hyperparameter settings in Appendix B.4. Complete proofs of the theoretical results are provided in Appendix A, with additional experimental results and analyses in Appendix C. Anonymous source code is available at `https://anonymous.4open.science/r/cm_RFFI`.

REFERENCES

Aristotelis Ballas and Christos Diou. Gradient-guided annealing for domain generalization. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 20558–20568. IEEE, 2025.

Metehan Cekic, Soorya Gopalakrishnan, and Upamanyu Madhow. Wireless fingerprinting via deep learning: The impact of confounding factors. In *2021 55th Asilomar Conference on Signals, Systems, and Computers*, pp. 677–684. IEEE, 2021.

Wanshi Chen, Lingjia Liu, Xiaofeng Liu, Erik G Larsson, Mohamed El Jaafari, and Tony QS Quek. Toward standardization of 6g and nextg: Key technologies to enable fundamental enhancements. *IEEE Journal on Selected Areas in Communications*, 44:4333–4365, 2026.

Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4690–4699, 2019.

Zehua Han, Jing Xiao, Qirui Zhao, Zhexuan Cui, Yufeng Wang, Duona Zhang, and Wenrui Ding. Open-world radio frequency fingerprint identification via augmented semi-supervised learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 264–272, 2025.

Samer Hanna, Samurdhi Karunaratne, and Danijela Cabric. Open set wireless transmitter authorization: Deep learning approaches and dataset considerations. *IEEE Transactions on Cognitive Communications and Networking*, 7(1):59–72, 2020.

Samer Hanna, Samurdhi Karunaratne, and Danijela Cabric. Wisig: A large-scale wifi signal dataset for receiver and channel agnostic rf fingerprinting. *IEEE Access*, 10:22808–22818, 2022.

David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghuai Zhang, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (rex). In *International conference on machine learning*, pp. 5815–5826, 2021.

Aodi Li, Liansheng Zhuang, Xiao Long, Minghong Yao, and Shafei Wang. Seeking consistent flat minima for better domain generalization via refining loss landscapes. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15349–15359. IEEE, 2025a.

Kunling Li, Jiazhong Bao, Xin Xie, Jianan Hong, and Cunqing Hua. Receiver-agnostic radio frequency fingerprint identification for zero-trust wireless networks. *IEEE Journal on Selected Areas in Communications*, 43(6):1981–1997, 2025b.

Kunling Li, Cunqing Hua, Hongyu Zhu, Tianjie Ju, and Pengwenlong Gu. Source-free open-world rf fingerprint identification. In *Forty-third International Conference on Machine Learning*, 2026.

Bingyuan Liu, Ismail Ben Ayed, Adrian Galdran, and Jose Dolz. The devil is in the margin: Margin-based label smoothing for network calibration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 80–88, 2022.

Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphereface: Deep hypersphere embedding for face recognition. In *2017 IEEE conference on computer vision and pattern recognition (CVPR)*, pp. 6738–6746. IEEE, 2017.

Jie Ma, Junqing Zhang, Guanxiong Shen, Linning Peng, and Alan Marshall. Towards channel-robust and receiver-independent radio frequency fingerprint identification. *IEEE Transactions on Information Forensics and Security*, 20:12112–12125, 2025.

Mazda Moayeri, Michael Rabbat, Mark Ibrahim, and Diane Bouchacourt. Embracing diversity: Interpretable zero-shot classification beyond one vector per class. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '24, pp. 2302–2321, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704505.

Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International conference on machine learning*, pp. 10–18, 2013.

Dileepa Pitawela, Gustavo Carneiro, and Hsiang-Ting Chen. Cloc: Contrastive learning for ordinal classification with multi-margin n-pair loss. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15538–15548. IEEE, 2025.

Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. *arXiv preprint arXiv:1911.08731*, 2019.

Olawale Salaudeen and Sanmi Koyejo. Causally inspired regularization enables domain general representations. In Sanjoy Dasgupta, Stephan Mandt, and Yingzhen Li (eds.), *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 3124–3132. PMLR, 02–04 May 2024.

Guanxiong Shen, Junqing Zhang, Alan Marshall, and Joseph R Cavallaro. Towards scalable and channel-robust radio frequency fingerprint identification for lora. *IEEE Transactions on Information Forensics and Security*, 17:774–787, 2022.

Feng Shi, Xixi Zhang, Tiantian Tang, Yu Wang, Tomoaki Ohtsuki, and Guan Gui. Robust specific emitter identification across modulation domains via domain-invariant variational autoencoding. *IEEE Transactions on Information Forensics and Security*, 2026.

Yanru Sun, Zongxia Xie, Emadeldeen Eldele, Dongyue Chen, Qinghua Hu, and Min Wu. Learning pattern-specific experts for time series forecasting under patch-level distribution shift. *Advances in Neural Information Processing Systems*, 38:91810–91844, 2025.

Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.

Peng Tang, Guoru Ding, Yitao Xu, Yutao Jiao, Yehui Song, and Guofeng Wei. Causal learning for robust specific emitter identification over unknown channel statistics. *IEEE Transactions on Information Forensics and Security*, 19:5316–5329, 2024.

Yuteng Wang, Zhenxin Cai, Tiantian Tang, Tomoaki Ohtsuki, and Guan Gui. Toward robust radio frequency fingerprint identification: A federated learning framework with feature alignment. *IEEE Transactions on Information Forensics and Security*, 21:2641–2651, 2026.

Hongxin Wei, Renchunzi Xie, Hao Cheng, Lei Feng, Bo An, and Yixuan Li. Mitigating neural network overconfidence with logit normalization. In *International conference on machine learning*, pp. 23631–23644. PMLR, 2022.

Lingnan Xie, Linning Peng, and Junqing Zhang. Towards robust rf fingerprint identification using spectral regrowth and carrier frequency offset. In *IEEE INFOCOM 2025 - IEEE Conference on Computer Communications*, pp. 1–10, 2025.

Haoran Zha, Hanhong Wang, Yu Wang, Zhi Sun, Guan Gui, and Yun Lin. Enhancing security in 5g nr with channel-robust rf fingerprinting leveraging srs for cross-domain stability. *IEEE Transactions on Information Forensics and Security*, 20:3429–3444, 2025.

Di Zhao, Jingfeng Zhang, Hongsheng Hu, Philippe Fournier-Viger, Gillian Dobbie, and Yun Sing Koh. Unlearning during training: Domain-specific gradient ascent for domain generalization. In *The Fourteenth International Conference on Learning Representations*, 2026.

Tianya Zhao, Xuyu Wang, and Shiwen Mao. Cross-domain, scalable, and interpretable rf device fingerprinting. In *IEEE INFOCOM 2024 - IEEE Conference on Computer Communications*, pp. 2099–2108, 2024.

Beier Zhu, Kesen Zhao, Jiequan Cui, Qianru Sun, Yuan Zhou, Xun Yang, and Hanwang Zhang. Reducing class-wise performance disparity via margin regularization. *arXiv preprint arXiv:2602.00205*, 2026.

APPENDIX

In appendix, we provide detailed theoretical proofs, comprehensive implementation details, and additional experimental results and corresponding analyses to support the main paper. The content is organized as follows:

- **Appendix A: Proofs**
  - A.1. Proof of Decision-Complete Margins (Proposition 3.1)
  - A.2. Proof of Maximal Admissible Cell (Theorem 3.2)
  - A.3. Proof of Evidence-Release Threshold (Theorem 3.3)
  - A.4. Proof of Bottleneck-Risk Transfer (Theorem 3.4)
- **Appendix B: Experimental Setup**
  - B.1. Datasets and Preprocessing
  - B.2. Evaluation Metrics
  - B.3. Baseline Configurations
  - B.4. Implementation Details
- **Appendix C: Detailed Experimental Results**
  - C.1. Curvature-Dependent Evidence Release
  - C.2. Multi-Seed Performance Comparison
  - C.3. Plug-in Compatibility Performance
  - C.4. Source-Target Evidence Geometry
  - C.5. Dataset-wise Component Ablation

A PROOFS

A.1 PROOF OF PROPOSITION 3.1

Proposition 3.1 (Decision-Complete Margins). For a fixed candidate identity y, the margin vector $m_y = (m_{yk})_{k \neq y} \in \mathbb{R}^{K-1}$ uniquely determines the softmax probabilities and the centered logits. Moreover, if ties are counted as errors, $\hat{y} = y$ if and only if $m_{yk} > 0$ for every $k \neq y$.

Proof. Fix y and write $m_k := m_{yk} = z_y - z_k$ for $k \neq y$. Then every competing logit can be expressed relative to z_y as

$$z_k = z_y - m_k, \qquad k \neq y. \tag{13}$$

Substituting this relation into the softmax denominator yields

$$\sum_{j=1}^{K} e^{z_j} = e^{z_y} + \sum_{k \neq y} e^{z_y - m_k} \tag{14}$$

$$= e^{z_y} \left(1 + \sum_{k \neq y} e^{-m_k} \right). \tag{15}$$

The common factor e^{z_y} cancels, giving

$$p_y = \frac{1}{1 + \sum_{k \neq y} e^{-m_k}}, \qquad p_k = \frac{e^{-m_k}}{1 + \sum_{j \neq y} e^{-m_j}}. \tag{16}$$

Hence the $K - 1$ margins uniquely determine the complete softmax probability vector.

The raw logits are identifiable only up to a common additive constant, since replacing z by $z + c\mathbf{1}$ leaves every margin unchanged. Let $\bar{z} := K^{-1} \sum_j z_j$ and define the centered logits $\tilde{z}_j := z_j - \bar{z}$.

Using equation 13, we have

$$\bar{z} = \frac{1}{K}\left(z_y + \sum_{k\neq y}(z_y - m_k)\right) \tag{17}$$

$$= z_y - \frac{1}{K}\sum_{k\neq y} m_k. \tag{18}$$

Therefore

$$\tilde{z}_y = \frac{1}{K}\sum_{j\neq y} m_j, \qquad \tilde{z}_k = -m_k + \frac{1}{K}\sum_{j\neq y} m_j. \tag{19}$$

Thus the centered logits are explicit functions of m_y. Conversely, if two logit vectors have the same margins with respect to y, equation 13 shows that they can differ only by one common additive constant, which centering removes.

Finally, with ties counted as errors, y is the unique predicted class exactly when $z_y > z_k$ for every $k \neq y$. Since $m_{yk} = z_y - z_k$, this is equivalent to $m_{yk} > 0$ for every rival k. Hence m_y is decision-complete. □

A.2 Proof of Theorem 3.2

Theorem 3.2 (Maximal Admissible Cell). Let $d = K - 1$ and $\tau = 1 - \kappa - \rho > 0$. If a centrally symmetric set $\mathcal{A} = -\mathcal{A}$ satisfies $M_{yk}(1 + r_k + \delta_k) \geq \kappa M_{yk}$ for every $r \in \mathcal{A}$, every $\|\delta\|_\infty \leq \rho$, and every rival k, then $\mathcal{A} \subseteq \mathcal{C}_y(\tau) := [-\tau, \tau]^d$. Conversely, $\mathcal{C}_y(\tau)$ itself satisfies the reserve condition.

Proof. Write $M_k := M_{yk} > 0$. Since M_k is positive, the reserve condition is equivalent, coordinate-wise, to

$$1 + r_k + \delta_k \geq \kappa. \tag{20}$$

Fix any $r \in \mathcal{A}$ and any coordinate k. Because equation 20 must hold for every perturbation with $\|\delta\|_\infty \leq \rho$, it must hold in particular for the worst negative displacement in that coordinate, $\delta_k = -\rho$. Hence

$$1 + r_k - \rho \geq \kappa, \tag{21}$$

which is equivalent to

$$r_k \geq \kappa + \rho - 1 = -\tau. \tag{22}$$

Thus every admissible point is bounded below by $-\tau$ in each coordinate.

The reserve condition itself is one-sided and does not directly upper-bound positive evidence. The upper bound follows from central symmetry. Since $r \in \mathcal{A}$ implies $-r \in \mathcal{A}$, applying equation 22 to $-r$ gives

$$-r_k \geq -\tau, \tag{23}$$

so $r_k \leq \tau$. Combining both inequalities gives

$$|r_k| \leq \tau \quad \forall k, \qquad \Longleftrightarrow \qquad r \in [-\tau, \tau]^d = \mathcal{C}_y(\tau). \tag{24}$$

Because $r \in \mathcal{A}$ is arbitrary,

$$\mathcal{A} \subseteq \mathcal{C}_y(\tau). \tag{25}$$

It remains to verify that the box itself is admissible. Take any $r \in \mathcal{C}_y(\tau)$ and any $\|\delta\|_\infty \leq \rho$. Then $r_k \geq -\tau$ and $\delta_k \geq -\rho$, so

$$1 + r_k + \delta_k \geq 1 - \tau - \rho \tag{26}$$

$$= 1 - (1 - \kappa - \rho) - \rho = \kappa. \tag{27}$$

Multiplying by $M_k > 0$ recovers the required reserve condition. Therefore $\mathcal{C}_y(\tau)$ is centrally symmetric and admissible, while every other centrally symmetric admissible set is contained in it. Hence it is the largest such set under set inclusion. □

Remark (Role of Central Symmetry). The reserve condition is intrinsically one-sided: without the constraint $A = -A$, it only imposes the lower bound $r_k \geq -\tau$, while positive evidence deviations remain unbounded. Central symmetry is therefore not implied by classification correctness; rather, we impose it as an explicit geometric design choice to obtain a finite, two-sided evidence budget. This choice is also consistent with the mean-preserving minibatch coordinates in Eq. (8): for each class–rival pair, positive and negative relative-evidence deviations are balanced around the shared reference, so excessively large positive deviations must be compensated by negative ones, which may in turn consume decision reserve. Under this centrally symmetric design family, $\mathcal{C}_y(\tau) = [-\tau, \tau]^{K-1}$ is the largest set satisfying the prescribed reserve condition.

A.3 Proof of Theorem 3.3

Theorem 3.3 (Evidence-Release Threshold). Under the mean-preserving constraint $n^{-1}\sum_i r_i = 0$, consider $\Omega(R) = n^{-1}\sum_i \omega(r_i)$, where ω is twice differentiable, $\nabla\omega(0) = 0$, and $\nabla^2\omega(0) = -2\mu I$ with $\mu > 0$. Let J_α be cross-entropy plus $\alpha\Omega$. Then the critical strength separating strict local point stability from instability is $\alpha_c = \lambda_{\min}(C_y)/(2\mu)$, and a unit direction v is released once $2\alpha\mu > v^\top C_y v$.

Proof. Let

$$\phi(m) = \log\left(1 + \sum_{k=1}^{d} e^{-m_k}\right), \qquad J_\alpha(R) = \frac{1}{n}\sum_{i=1}^{n}\phi(M + D_M r_i) + \frac{\alpha}{n}\sum_{i=1}^{n}\omega(r_i), \tag{28}$$

where $D_M = \mathrm{Diag}(M)$ and $R = (r_1, \ldots, r_n)$. The mean-preserving constraint defines the linear subspace

$$\mathcal{T} := \left\{ R : \sum_{i=1}^{n} r_i = 0 \right\}, \tag{29}$$

and the point-evidence state is $R = \mathbf{0} \in \mathcal{T}$.

We first compute the local curvature contributed by cross-entropy. At the reference margin M, define

$$a_k := \frac{e^{-M_k}}{1 + \sum_j e^{-M_j}}, \qquad k = 1, \ldots, d. \tag{30}$$

Differentiating ϕ with respect to m_k gives

$$\frac{\partial\phi}{\partial m_k}(M) = -\frac{e^{-M_k}}{1 + \sum_j e^{-M_j}} = -a_k. \tag{31}$$

Differentiating again, for any k, ℓ,

$$\frac{\partial^2\phi}{\partial m_k \partial m_\ell}(M) = -\frac{\partial a_k}{\partial m_\ell} \tag{32}$$

$$= a_k \mathbf{1}\{k = \ell\} - a_k a_\ell. \tag{33}$$

Hence

$$\nabla_m^2\phi(M) = \mathrm{Diag}(a) - aa^\top. \tag{34}$$

This Hessian is positive definite. Indeed, for any nonzero $u \in \mathbb{R}^d$,

$$u^\top\left[\mathrm{Diag}(a) - aa^\top\right]u = \sum_k a_k u_k^2 - \left(\sum_k a_k u_k\right)^2 \tag{35}$$

$$\geq \left(1 - \sum_k a_k\right)\sum_k a_k u_k^2 > 0, \tag{36}$$

where the inequality is weighted Cauchy-Schwarz, and $1 - \sum_k a_k = p_y(M) > 0$.

Since $m = M + D_M r$ is affine in r with Jacobian D_M, the Hessian in relative-evidence coordinates follows directly from the chain rule:

$$\nabla_r^2\phi(M + D_M r)\big|_{r=0} = D_M\left[\mathrm{Diag}(a) - aa^\top\right]D_M = C_y. \tag{37}$$

Because $D_M \succ 0$ and the raw-margin Hessian is positive definite, $C_y \succ 0$ as well.

Now let $H = (h_1, \dots, h_n)$ be any feasible perturbation direction in $\mathcal{T}$, so $\sum_i h_i = 0$. The first-order variation of the CE term at $R = \mathbf{0}$ is

$$\mathrm{d}J_{\mathrm{CE}}(\mathbf{0})[H] = \frac{1}{n} \sum_i \nabla\phi(M)^\top D_M h_i \tag{38}$$

$$= \nabla\phi(M)^\top D_M \left(\frac{1}{n} \sum_i h_i \right) = 0. \tag{39}$$

The regularizer also contributes no linear term because $\nabla\omega(0) = 0$:

$$\mathrm{d}\Omega(\mathbf{0})[H] = \frac{1}{n} \sum_i \nabla\omega(0)^\top h_i = 0. \tag{40}$$

Therefore $R = \mathbf{0}$ is stationary on the mean-preserving subspace.

Using equation 37 together with $\nabla^2\omega(0) = -2\mu I$, the second directional derivative is

$$\mathrm{d}^2 J_\alpha(\mathbf{0})[H, H] = \frac{1}{n} \sum_i h_i^\top C_y h_i + \frac{\alpha}{n} \sum_i h_i^\top (-2\mu I) h_i \tag{41}$$

$$= \frac{1}{n} \sum_i h_i^\top \big(C_y - 2\alpha\mu I \big) h_i. \tag{42}$$

Equivalently, along the feasible path $R(t) = tH$,

$$J_\alpha(tH) - J_\alpha(\mathbf{0}) = \frac{t^2}{2n} \sum_i h_i^\top \big(C_y - 2\alpha\mu I \big) h_i + o(t^2), \qquad t \to 0. \tag{43}$$

Thus, after the mean-preserving constraint removes the first-order variation, local point stability is governed at second order by $C_y - 2\alpha\mu I$.

The smallest eigenvalue of this matrix is

$$\lambda_{\min}(C_y) - 2\alpha\mu, \tag{44}$$

which changes sign at

$$\alpha_c = \frac{\lambda_{\min}(C_y)}{2\mu}. \tag{45}$$

If $\alpha < \alpha_c$, then $C_y - 2\alpha\mu I \succ 0$, and every nonzero feasible H satisfies

$$\mathrm{d}^2 J_\alpha(\mathbf{0})[H, H] \geq \big(\lambda_{\min}(C_y) - 2\alpha\mu \big) \frac{1}{n} \sum_i \|h_i\|_2^2 > 0. \tag{46}$$

Together with stationarity, the second-order sufficient condition implies that $R = \mathbf{0}$ is a strict local minimum of J_α on $\mathcal{T}$.

If $\alpha > \alpha_c$, let v be a unit eigenvector of C_y associated with $\lambda_{\min}(C_y)$. Because $n \geq 2$, the perturbation $h_1 = v$, $h_2 = -v$, and $h_i = 0$ for $i > 2$ belongs to $\mathcal{T}$. Substituting it into equation 42 gives

$$\mathrm{d}^2 J_\alpha(\mathbf{0})[H, H] = \frac{2}{n} \big(\lambda_{\min}(C_y) - 2\alpha\mu \big) < 0. \tag{47}$$

By equation 43, $J_\alpha(tH) < J_\alpha(\mathbf{0})$ for all sufficiently small nonzero t, so the point-evidence state is no longer a local minimum.

The same construction gives the directional criterion. For any unit v, choosing $h_1 = v$ and $h_2 = -v$ yields

$$\mathrm{d}^2 J_\alpha(\mathbf{0})[H, H] = \frac{2}{n} \left(v^\top C_y v - 2\alpha\mu \right). \tag{48}$$

Hence direction v acquires negative local curvature exactly when $2\alpha\mu > v^\top C_y v$. Lower-curvature evidence modes therefore cross the release threshold at smaller α.

At $\alpha = \alpha_c$, the second-order form is only positive semidefinite, so higher-order terms may determine the behavior along the zero-curvature eigenspace; no stronger conclusion is implied at equality. □

Remark (Automatic Mean Preservation and Adaptive Evidence Geometry). The mean-preserving condition in Theorem 3.3 is automatically satisfied by the practical construction in Eq. (8). For every eligible class y and valid rival k,

$$\frac{1}{n_y}\sum_{i\in B_y} r_{ik} = \frac{\sum_{i\in B_y}(m_{ik} - M_{yk})}{n_y(M_{yk}+\epsilon)} = 0.$$

Hence, the minibatch relative-evidence coordinates lie exactly in the mean-preserving subspace analyzed in Theorem 3.3. The stabilizer ϵ does not affect this property because the denominator is shared within each class-rival pair.

The fixed reference M_y in Theorem 3.3 should therefore be understood as describing the local decision geometry around an instantaneous reference state. In the practical MarginCell construction, M_{yk} is recomputed dynamically during training; consequently, the class–rival margins, rival probabilities, and the corresponding CE curvature C_y evolve jointly with the classifier. MarginCell thus realizes the local release principle adaptively: the dispersion reward supplies an isotropic anti-contraction force, while the current rival-wise CE geometry determines which evidence directions are released first. This adaptive construction avoids enforcing a fixed evidence prototype while preserving controlled rival-wise variation.

A.4 Proof of Theorem 3.4

Theorem 3.4 (Bottleneck-Risk Transfer). For any $\rho \geq 0$, a sample remains strictly correct under every relative-evidence shift $\|\delta\|_\infty \leq \rho$ if and only if $s(x,y) > \rho$, where $s(x,y) = \min_{k\neq y} m_{yk}(x)/M_{yk} = 1 + \min_k r_{yk}(x)$. For a correctly classified sample, $s(x,y)$ is therefore its supremal ℓ_∞ robustness radius in relative-evidence space. Moreover, suppose the source and target evidences are expressed with respect to the same fixed positive reference M_y, and there exists a label-preserving class-balanced source-target coupling under which the probability of a displacement larger than ρ is at most η. Then $R_T^{\mathrm{bal}} \leq \eta + \Pr_S^{\mathrm{bal}}[s(X,Y) \leq \rho]$.

Proof. Write $M_k := M_{yk} > 0$. From the definition of the relative coordinate,

$$m_k = M_k(1 + r_k). \tag{49}$$

After an additive relative-evidence shift δ, the perturbed coordinate is $r'_k = r_k + \delta_k$, so

$$m'_k = M_k(1 + r_k + \delta_k). \tag{50}$$

Because $M_k > 0$, the perturbed sample is strictly classified as y if and only if

$$1 + r_k + \delta_k > 0 \qquad \text{for every } k. \tag{51}$$

Thus strict correctness is controlled by the smallest perturbed normalized margin.

Recall

$$s(r) := 1 + \min_k r_k = \min_k(1 + r_k). \tag{52}$$

For any $\|\delta\|_\infty \leq \rho$, each coordinate satisfies $\delta_k \geq -\rho$, and therefore

$$\min_k(1 + r_k + \delta_k) \geq \min_k(1 + r_k - \rho) \tag{53}$$

$$= s(r) - \rho. \tag{54}$$

The bound is attained. If $k^\star \in \arg\min_k(1 + r_k)$, choose $\delta_{k^\star} = -\rho$ and, for example, $\delta_k = 0$ for $k \neq k^\star$. Then $\|\delta\|_\infty = \rho$ and

$$1 + r_{k^\star} + \delta_{k^\star} = s(r) - \rho. \tag{55}$$

Consequently,

$$\inf_{\|\delta\|_\infty \leq \rho} \min_k(1 + r_k + \delta_k) = s(r) - \rho. \tag{56}$$

Combining equation 51 and equation 56, the sample remains strictly correct for every admissible shift if and only if

$$s(r) - \rho > 0 \qquad \Longleftrightarrow \qquad s(x,y) > \rho. \tag{57}$$

For a correctly classified sample $s(x, y) > 0$, define

$$\rho_\star(x, y) := \sup \{\rho \geq 0 : \hat{y}(r + \delta) = y \text{ for all } \|\delta\|_\infty \leq \rho\}. \tag{58}$$

Equation 57 gives $\rho_\star(x, y) = s(x, y)$. The supremum, rather than a closed maximum, reflects the convention that a tie on the decision boundary is counted as an error.

To connect this sample-wise radius to transfer, first note that s is 1-Lipschitz under ℓ_∞. For any two relative-evidence vectors r and r',

$$s(r') = 1 + \min_k r'_k \tag{59}$$

$$\geq 1 + \min_k r_k - \|r' - r\|_\infty \tag{60}$$

$$= s(r) - \|r' - r\|_\infty. \tag{61}$$

Interchanging r and r' gives the reverse inequality, hence

$$|s(r') - s(r)| \leq \|r' - r\|_\infty. \tag{62}$$

Let Π^{bal} be the assumed label-preserving class-balanced coupling. A coupled draw is denoted as (X_S, X_T, Y), with the source and target relative-evidence vectors r^S and r^T computed using the same reference M_Y. Define $\Delta := r^T - r^S$. By assumption,

$$\Pr_{\Pi}^{\text{bal}}\big(\|\Delta\|_\infty > \rho\big) \leq \eta. \tag{63}$$

Under the condition $\|\Delta\|_\infty \leq \rho$ and $s(X_S, Y) > \rho$, equation 61 gives

$$s(X_T, Y) \geq s(X_S, Y) - \|\Delta\|_\infty > \rho - \rho = 0. \tag{64}$$

The coupled target sample is therefore strictly correct. Equivalently,

$$\{\hat{Y}_T \neq Y\} \subseteq \{\|\Delta\|_\infty > \rho\} \cup \{s(X_S, Y) \leq \rho\}. \tag{65}$$

Taking class-balanced probability and applying the union bound,

$$R_T^{\text{bal}} = \Pr_{\Pi}^{\text{bal}}(\hat{Y}_T \neq Y) \tag{66}$$

$$\leq \Pr_{\Pi}^{\text{bal}}(\|\Delta\|_\infty > \rho) + \Pr_{\Pi}^{\text{bal}}(s(X_S, Y) \leq \rho) \tag{67}$$

$$\leq \eta + \Pr_{S}^{\text{bal}}[s(X, Y) \leq \rho], \tag{68}$$

where the last term is the source marginal of the coupling. This proves the target-risk bound. □

Corollary (cell reserve). If $r \in \mathcal{C}_y(\tau)$, then $r_k \geq -\tau$ for every k, and therefore

$$s(x, y) = 1 + \min_k r_k \geq 1 - \tau. \tag{69}$$

Thus the admissible cell permits nonzero evidence variation while guaranteeing a minimum bottleneck reserve of $1 - \tau$.

B EXPERIMENTAL SETUP

B.1 DATASETS AND PREPROCESSING

We evaluate MarginCell on three representative public RFFI benchmarks covering distinct wireless protocols and acquisition conditions, each reflecting a distinct challenge in cross-environment RFFI: cross-day channel variation under multi-receiver acquisition (WiSig), large-scale temporal fingerprint drift across commercial WiFi devices (UCLA-CORES), and heterogeneous location- and mobility-induced propagation shifts (LoRa).

For each dataset, we split source and target domains by physical acquisition environments while keeping the transmitter identity space unchanged across domains. Target environments are held out during training and used only for evaluation, while environment labels are used solely to construct the source-target partitions and are not required by MarginCell during optimization. This provides a controlled source-only domain generalization protocol for evaluating robustness to unseen environmental shifts. Detailed dataset statistics, domain partitions, and exact preprocessing configurations are summarized in Table 3.

Table 3: **Dataset statistics and preprocessing configurations.** Setting denotes the acquisition variability. L_{IQ} represents the length of complex IQ segments. Split indicates the source-to-target environment partition. All network inputs are normalized and stacked into 3 channels.

Dataset	Setting	Protocol	Segment	K	Split (Source : Target)	N_{total}	L_{IQ}	Net Input ($H \times W \times C$)
WiSig	Multi-Day/Multi-Rx	WiFi 802.11a/g	Preamble	6	2 days : 2 days	24,000	256	$32 \times 32 \times 3$
UCLA-CORES	Multi-Day	WiFi 802.11g	Preamble	58	3 days : 2 days	58,000	256	$32 \times 32 \times 3$
LoRa	Multi-Scenario	LoRa	Preamble	25	6 env. : 2 env.	25,000	8192	$102 \times 62 \times 3$

- **WiSig (Hanna et al., 2022).** WiSig has been widely regarded as one of the foundational benchmarks for RFFI, owing to its high signal quality, large scale, and systematic coverage across multiple receivers and acquisition days. Its multi-receiver, multi-day design exposes both receiver heterogeneity and temporal channel drift under otherwise controlled WiFi transmissions, making it a strong reference for cross-environment generalization. We adopt the *ManySig* subset collected on the ORBIT testbed, which contains signals from 6 identical Atheros 5212 WiFi transmitters captured by 12 USRP B210/N210 receivers over four days. For each packet, we retain a fixed 256-sample preamble segment as the fingerprint-bearing input.
- **UCLA-CORES (Hanna et al., 2020).** UCLA-CORES is one of the largest publicly available multi-day RFFI collections built from commercial WiFi hardware, originally comprising 163 consumer devices repeatedly measured across multiple acquisition days. Its combination of a large device population and repeated over-the-air measurements makes it particularly valuable for studying whether fine-grained RF fingerprints remain identifiable under realistic day-to-day channel and acquisition variations. The dataset was collected on the ORBIT testbed using commercial Atheros 5212, 9220, and 9280 WiFi modules communicating with an IEEE 802.11g access point, with signals captured by a USRP N210 receiver. We use the 58-device cohort consistently observed across five acquisition days, taking the first three days as source environments and the remaining two as unseen targets, and retain the first 256 IQ samples of each packet preamble.
- **LoRa (Shen et al., 2022).** LoRa is a representative benchmark for scalable and channel-robust RFFI in large-scale IoT deployments, featuring off-the-shelf commercial devices measured under substantially more diverse propagation conditions. Its defining characteristic is environment-driven channel variability: measurements span different locations and mobility conditions, introducing strong multipath and mobility-induced distortions that complement the predominantly cross-day shifts in the two WiFi benchmarks. The signals are captured using a USRP N210 receiver from commercial LoRa transmitters under real over-the-air propagation. We use a 25-device subset organized into six source and two held-out target environments, and extract an 8192-sample payload-independent preamble segment from each packet for fingerprint identification.

Domain partitions. We use one fixed source-target partition for each benchmark throughout all experiments and random seeds. Specifically, WiSig uses two acquisition days as source domains and the remaining two days as unseen target domains; UCLA-CORES uses the first three acquisition days for training and the remaining two days for evaluation; and LoRa uses six acquisition environments as source domains and two held-out environments as target domains. The same transmitter cohorts and domain partitions are used for all compared methods. Target environments are never accessed during training, hyperparameter selection, or model selection.

Preprocessing. To initially suppress channel-induced nuisance while retaining device-specific distortions, we transform raw IQ sequences into log-power spectrograms. Specifically, we apply STFT with a Hamming window of size $N_{\text{FFT}} = 128$ and 50% overlap, compute the power spectrogram $S(f, t)$, and use adjacent-frame ratioing $S(:, t)/S(:, t-1)$ to attenuate slowly varying channel components, followed by logarithmic scaling and normalization. For fair comparison, all baselines and MarginCell use exactly the same preprocessing pipeline and network inputs.

B.2 Evaluation Metrics

Following prior robustness-oriented DG evaluation protocols (Salaudeen & Koyejo, 2024; Moayeri et al., 2024), we assess cross-environment RFFI from three complementary perspectives: overall recognition performance, worst-environment robustness, and hard-identity tail robustness.

For any sample set $\mathcal{S}$, define its recognition accuracy as $\mathcal{A}(\mathcal{S}) = \frac{1}{|\mathcal{S}|} \sum_{(x_i, y_i) \in \mathcal{S}} \mathbb{1}[\hat{y}_i = y_i]$.

- **Target Accuracy (Acc):** Measures the aggregate recognition performance across all unseen target samples $\mathcal{D}_T$, defined as $\text{Acc} = \mathcal{A}(\mathcal{D}_T)$.
- **Worst-Environment Accuracy (WEnv):** Measures robustness to the most challenging unseen environment $e \in \mathcal{E}_T$, computed as $\text{WEnv} = \min_{e \in \mathcal{E}_T} \mathcal{A}(\mathcal{D}_T^e)$.
- **Tail-Identity Accuracy ($\text{TIA}_{@20}$):** Measures tail robustness over hard identities. Letting $\mathcal{H}_{20}$ denote the lowest-performing 20% of identities ranked by class-wise accuracy $\mathcal{A}(\mathcal{D}_{T,c})$, it is computed as $\text{TIA}_{@20} = \frac{1}{|\mathcal{H}_{20}|} \sum_{c \in \mathcal{H}_{20}} \mathcal{A}(\mathcal{D}_{T,c})$.

B.3 Baseline Configurations

To systematically evaluate MarginCell under source-only RF domain generalization, we compare against ERM and three complementary groups of baselines. Our selection covers different mechanisms for handling environmental shifts:

- **Domain generalization (DG):** methods covering worst-group optimization (GroupDRO), risk equalization (V-REx), distribution alignment (MMD-DG), and pattern specialization (TFPS).
- **Robust RFFI:** methods covering signal-side channel suppression (CIS and CKR-RFF) and representation-side stabilization through causal purification (CLPS) and prototype-based metric learning (CDSI-RFF).
- **Output regularization:** methods covering soft-target confidence control (Label Smoothing), inequality-based logit-gap constraints (MbLS), and logit-scale normalization (LogitNorm).

Together, these comparisons assess whether explicitly releasing decision-safe evidence variation while preserving rival-wise reserve provides benefits beyond generic DG objectives, RF-specific nuisance suppression, and conventional output regularization. Below, we summarize the defining mechanism and comparative role of each baseline.

ERM reference. ERM pools labeled source samples and minimizes standard cross-entropy without an additional DG or output-regularization objective. It provides the common reference for evaluating the contribution of explicit robustness mechanisms beyond supervised source fitting.

DG baselines. We include representative risk- and representation-based DG objectives, together with a pattern-specialization approach to distribution shift:

- **GroupDRO** (Sagawa et al., 2019) *(worst-group risk optimization).* A group-robust optimization baseline that minimizes the worst-case training risk over predefined groups, adaptively emphasizing high-loss groups rather than optimizing only average performance. It provides a risk-level robustness reference for assessing whether protecting difficult source groups is sufficient for generalization to unseen environments.
- **V-REx** (Krueger et al., 2021) *(risk-variance regularization).* A risk-invariance baseline that augments average source risk with a penalty on the variance of risks across source environments. Unlike GroupDRO's emphasis on the worst group, it encourages uniformly low risks across environments. Its inclusion contrasts environment-level risk equalization with MarginCell's regulation of within-identity decision evidence.
- **MMD-DG** (Muandet et al., 2013) *(kernel-based distribution alignment).* A kernel mean-matching baseline following the invariant-representation principle. Maximum mean discrepancy measures differences between source-domain feature distributions through their kernel mean embeddings; reducing this discrepancy encourages statistical alignment while source supervision preserves identity discrimination. It serves as a distribution-level invariance reference for evaluating decision-level evidence regularization.
- **TFPS** (Sun et al., 2025) *(pattern-specific expert modeling).* A pattern-specialized architecture originally developed for time-series forecasting under patch-level distribution shift. It combines time- and frequency-domain encoding with subspace-based pattern identification and routes patches to pattern-specific experts. It represents architectural accommodation of heterogeneous patterns, providing a specialization-based comparator rather than a canonical DG objective.

Robust RFFI baselines. Following the distinction between signal-side suppression and representation-side invariance, we include RF-specific approaches that reduce channel interference before recognition or stabilize identity-discriminative features during learning:

- **CIS** (Shen et al., 2022) *(channel-independent spectral representation).* A physics-motivated approach that constructs channel-independent spectrograms from ratios of adjacent STFT frames, exploiting local channel stationarity to attenuate shared channel responses. The original framework additionally combines channel augmentation with triplet-based metric learning. It provides a signal-side invariance reference, contrasting physical nuisance cancellation with direct control of decision evidence.
- **CKR-RFF** (Zha et al., 2025) *(channel-knowledge-guided fingerprinting).* A protocol-informed framework that uses 5G NR sounding reference signals to obtain channel state information and construct relative RF fingerprints. Its original design also incorporates training regularization to reduce overfitting and domain bias. It represents channel-informed feature construction, providing a physically grounded counterpart to learning tolerance in decision-evidence space.
- **CLPS** (Tang et al., 2024) *(causality-driven feature purification).* The original method combines centralized logarithmic power spectrum features with a causal purification network. Spectral processing attenuates channel effects, while feature decorrelation and purification weaken spurious channel-identity associations. It bridges signal-side suppression and representation-side stabilization, serving as a reference for removing environment-related shortcuts rather than explicitly accommodating residual evidence variation.
- **CDSI-RFF** (Zhao et al., 2024) *(prototype-based metric learning).* A prototype-driven framework that learns fingerprint embeddings through classification, prototype-similarity, and discriminability losses, and identifies devices by cosine similarity to their prototypes. The original pipeline further uses explanation-guided augmentation during few-shot fine-tuning. It provides a metric-learning reference that stabilizes identity representations around prototypes, contrasting this representation-space prior with bounded variation in decision evidence.

Output-regularization baselines. These methods act directly on training targets or classifier logits, providing close comparisons for determining whether MarginCell's gains can be explained by conventional confidence, margin, or scale control:

- **Label Smoothing (LS)** (Szegedy et al., 2016) *(soft-target confidence control).* A widely adopted regularizer that replaces one-hot labels with a mixture of the ground-truth distribution and a uniform distribution over classes. Training against these softened targets discourages excessively confident predictions. It provides a confidence-control reference for assessing whether generic softening alone accounts for improved cross-environment robustness.
- **MbLS** (Liu et al., 2022) *(inequality-based margin regularization).* A margin-based calibration method that penalizes logit gaps $\max_j z_j - z_k$ only when they exceed a prescribed margin, leaving smaller gaps unpenalized by the regularizer. It is a particularly relevant comparator because it permits output slack; the comparison examines whether this allowance suffices without actively rewarding within-identity evidence dispersion.
- **LogitNorm** (Wei et al., 2022) *(logit-scale normalization).* A scale-control method that applies ℓ_2 normalization and a fixed temperature to each logit vector before computing cross-entropy, decoupling optimization from unrestricted logit-norm growth. It provides a normalization-based reference for assessing whether controlling overall output magnitude suffices, rather than regulating class-relative, rival-wise evidence geometry.

Unified implementation protocol. For fair comparison, all methods use the same preprocessing, input representation, and ResNet-18 backbone unless required otherwise by the original method. For baselines from different protocols or tasks, only incompatible components are adapted, while preserving the core objective and using no target-domain information.

B.4 Implementation Details

Hardware and software. All experiments and evaluations are conducted on a server equipped with a 13th Gen Intel(R) Core(TM) i9-13900K CPU, an NVIDIA GeForce RTX 4090 GPU, and 128 GB of RAM. MarginCell and all baselines are implemented in Python 3.8.20 using PyTorch 2.4.1.

Training schedule. For fair comparison, we employ a uniform ResNet-18 backbone with a 512-dimensional feature representation for all experiments. Models are trained for 100 epochs using the Adam optimizer with a batch size of 200 and an initial learning rate of 5×10^{-4}, which is decayed with a cosine scheduler. We use a weight decay of 10^{-4} and a 5-epoch warm-up stage before fully activating the MarginCell regularization. All compared methods follow the same training budget and backbone configuration unless their original formulation requires method-specific optimization. For all methods, we report the final checkpoint; target-domain information is never used for model selection or early stopping.

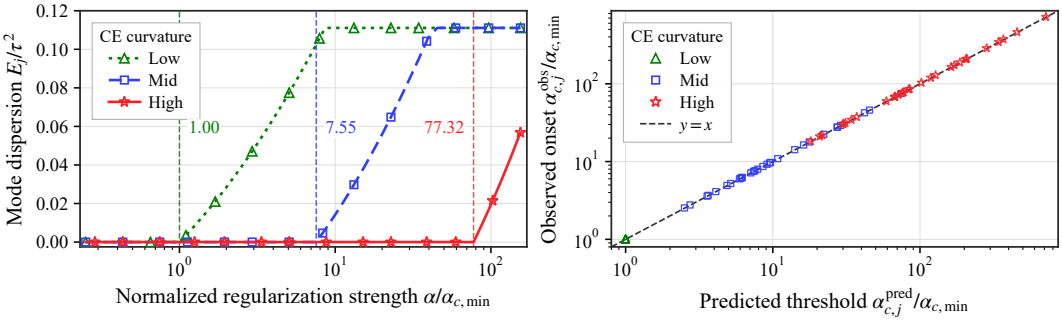
Hyperparameter settings. For MarginCell, we set the cell half-width to $\tau = 0.5$, the barrier strength to $\beta = 6$, and the overall regularization weight to $\lambda = 0.1$. The validity threshold for positive class-rival references is set to $\gamma = 0.1$, with $\epsilon = 10^{-6}$ used for numerical stability in relative-evidence normalization. These parameters are kept fixed across WiSig, UCLA-CORES, and LoRa, as well as in the plug-in experiments, unless otherwise specified. For the sensitivity analysis in Section 5.5, we evaluate all 5×4 combinations of $\tau \in \{0.10, 0.25, 0.50, 0.75, 0.90\}$ and $\beta \in \{0, 3, 6, 9\}$. For the PointAlign controls in Section 5.3, relative-evidence coordinates are computed minibatch-wise as in MarginCell and regularized toward the class reference $r = 0$, with coefficients 1 and 16 for PointAlign1 and PointAlign16, respectively.

Remark. MarginCell operates directly on source-label-conditioned classifier logits and requires neither target-domain observations nor environment labels during training. It introduces no additional trainable parameters or inference branch: the relative-evidence references and regularization terms are used only during optimization, while inference follows the standard ResNet classifier. Consequently, MarginCell can be integrated into existing DG or robust RFFI objectives without modifying their deployment pipeline.

C Detailed Experimental Results

C.1 Curvature-Dependent Evidence Release

Figure 7 examines the curvature-dependent release predicted by Theorem 3.3 under the fixed-class-mean setting of Section 3.3. The two panels evaluate the order of release across evidence directions and the quantitative accuracy of the predicted thresholds.



(a) Direction-specific release **(b) Threshold prediction**

Figure 7: **Curvature-dependent evidence release.** (a) Mode dispersion versus regularization strength; vertical dashed lines mark predicted release thresholds. (b) Observed versus predicted thresholds, with $y = x$ indicating exact agreement.

Direction-specific release. Figure 7(a) plots the normalized mode dispersion E_j/τ^2 against $\alpha/\alpha_{c,\min}$, where $\alpha_{c,\min}$ is the smallest predicted release threshold. Under weak regularization, all modes remain concentrated near zero; as α increases, the low-, mid-, and high-curvature modes are released successively around the predicted normalized thresholds 1.00, 7.55, and 77.32. **This ordering is consistent with the curvature competition characterized by Theorem 3.3: lower-curvature directions require less anti-contraction to escape CE-induced point stability.**

More Importantly, MarginCell supplies an isotropic anti-contraction force, while the anisotropic CE geometry determines which directions are released first. The resulting release is therefore geometry-adaptive: greater evidence freedom is first allocated to directions that are less strongly constrained by the current decision geometry, whereas higher-curvature directions remain contracted until stronger regularization is applied. This selective mechanism provides additional capacity to absorb environment-induced evidence variation without indiscriminately relaxing the entire identity-discriminative structure, which is the type of controlled flexibility desired for domain generalization.

Agreement with theory. Figure 7(b) compares the observed and predicted release onsets, both normalized by $\alpha_{c,\min}$. **The points closely follow the diagonal $y = x$, indicating agreement in the magnitude of the release thresholds rather than merely their ordering across curvature modes.** This numerically supports the curvature-dependent threshold predicted by Theorem 3.3 and complements Figure 7(a), which visualizes the corresponding mode-wise transition from contraction to dispersion.

Together, the two panels reinforce a central insight of MarginCell: admissible evidence freedom is not automatically realized under CE optimization, but must be actively released; meanwhile, the classifier's own decision geometry naturally determines where that freedom emerges first. The dispersion reward provides this release force, while the boundary barrier subsequently prevents the released variation from eroding decision reserve.

C.2 Multi-Seed Performance Comparison

Table 4 reports the mean and standard deviation of the main performance comparison over five independent runs with different random seeds. MarginCell maintains cross-environment recognition accuracy above 93% on both WiSig and LoRa, while exhibiting consistently strong performance across different datasets and evaluation metrics. In particular, on $\text{TIA}_{@20}$, which focuses on hard

Table 4: MarginCell performance comparison with other methods across datasets over 5 random seeds (mean ± std). (**Bold**: Best; Underline: Second-best).

Type	Method	WiSig (2→2 d.)			UCLA-CORES (3→2 d.)			LoRa (6→2 env.)			*Overall Avg.*		
		Acc	WEnv	$TIA_{@20}$	Acc	WEnv	$TIA_{@20}$	Acc	WEnv	$TIA_{@20}$	Acc	WEnv	$TIA_{@20}$
Reference	ERM	83.50 ± 5.42	83.00 ± 5.36	51.00 ± 16.26	65.37 ± 3.79	60.94 ± 5.60	31.61 ± 10.23	87.93 ± 3.75	82.60 ± 4.49	63.00 ± 10.78	78.93 ± 2.53	75.51 ± 2.99	48.54 ± 7.34
DG Methods	TFPS	91.75 ± 1.92	86.30 ± 3.15	76.35 ± 5.66	70.82 ± 2.28	67.36 ± 2.34	32.84 ± 11.75	79.65 ± 0.39	66.58 ± 0.46	54.63 ± 0.53	80.74 ± 1.00	73.41 ± 1.32	54.61 ± 4.35
	GroupDRO	88.15 ± 6.54	86.73 ± 6.97	64.80 ± 19.68	69.24 ± 5.66	65.99 ± 6.51	44.77 ± 3.48	89.26 ± 1.26	83.10 ± 1.06	62.19 ± 2.92	82.22 ± 2.91	78.61 ± 3.20	57.25 ± 6.73
	MMD-DG	89.02 ± 3.73	86.47 ± 2.69	67.40 ± 11.05	71.84 ± 1.71	68.11 ± 2.28	46.65 ± 5.80	92.75 ± 3.04	88.85 ± 4.53	76.75 ± 11.67	84.54 ± 1.70	81.14 ± 1.91	63.60 ± 5.70
	V-REx	89.92 ± 5.83	88.13 ± 6.81	70.10 ± 17.58	69.95 ± 3.87	68.91 ± 2.91	41.87 ± 9.97	90.11 ± 0.94	84.88 ± 0.46	65.44 ± 5.75	83.33 ± 2.35	80.64 ± 2.47	59.14 ± 7.00
Avg. (DG)		89.71 ± 2.43	86.91 ± 2.65	69.66 ± 7.29	70.46 ± 1.86	67.59 ± 1.96	41.53 ± 4.21	87.94 ± 0.86	80.85 ± 1.17	64.75 ± 3.34	82.70 ± 1.06	78.45 ± 1.17	58.65 ± 3.02
Robust RFFI	CDSI-RFF	89.92 ± 1.86	83.03 ± 0.07	70.05 ± 5.64	70.23 ± 2.15	68.20 ± 2.68	21.41 ± 5.21	39.80 ± 3.54	31.15 ± 4.10	14.00 ± 3.36	66.65 ± 1.51	60.79 ± 1.63	35.15 ± 2.79
	CLPS	94.33 ± 1.65	89.47 ± 3.18	83.00 ± 4.96	73.33 ± 2.41	70.91 ± 2.34	42.64 ± 9.20	88.18 ± 7.46	83.05 ± 6.93	58.44 ± 25.54	85.28 ± 2.67	81.14 ± 2.66	61.36 ± 9.20
	CIS	93.22 ± 1.85	**91.43** ± 2.74	79.95 ± 5.85	68.20 ± 2.92	67.32 ± 3.79	29.32 ± 2.77	89.05 ± 0.49	82.75 ± 0.07	63.44 ± 1.15	83.49 ± 1.16	80.50 ± 1.56	57.57 ± 2.19
	CKR-RFF	88.90 ± 5.18	85.83 ± 4.19	67.00 ± 15.62	70.75 ± 1.29	68.38 ± 2.04	36.68 ± 4.18	90.94 ± 2.46	84.13 ± 4.49	71.00 ± 12.20	83.53 ± 1.96	79.45 ± 2.16	58.23 ± 6.75
Avg. (Robust RFFI)		91.59 ± 1.51	87.44 ± 1.48	75.00 ± 4.57	70.63 ± 1.14	68.70 ± 1.40	32.51 ± 2.93	76.99 ± 2.16	70.27 ± 2.30	51.72 ± 7.13	79.74 ± 0.96	75.47 ± 1.03	53.08 ± 2.99
Regularization	LS	90.90 ± 5.51	87.93 ± 6.21	73.00 ± 16.68	73.55 ± 1.86	71.85 ± 3.69	47.83 ± 7.67	91.88 ± 2.86	87.18 ± 4.07	69.88 ± 14.67	85.44 ± 2.16	82.32 ± 2.76	63.57 ± 7.83
	MbLS	89.75 ± 5.28	87.10 ± 4.77	69.45 ± 16.00	72.58 ± 2.06	70.95 ± 2.78	47.78 ± 8.65	91.46 ± 3.41	87.70 ± 3.32	72.13 ± 9.90	84.60 ± 2.20	81.92 ± 2.15	63.12 ± 6.90
	LogitNorm	91.38 ± 4.67	88.77 ± 5.50	74.55 ± 14.04	70.78 ± 1.60	67.25 ± 2.33	41.22 ± 1.46	91.63 ± 1.34	84.93 ± 3.08	72.00 ± 3.89	84.60 ± 1.71	80.32 ± 2.24	62.59 ± 4.88
Avg. (Reg.)		90.68 ± 2.98	87.93 ± 3.19	72.33 ± 9.01	72.30 ± 1.07	70.02 ± 1.72	45.61 ± 3.88	91.66 ± 1.55	86.60 ± 2.03	71.34 ± 6.04	84.88 ± 1.18	81.52 ± 1.39	63.09 ± 3.84
-	**MarginCell**	**95.15** ± 2.52	91.33 ± 3.21	**85.70** ± 1.34	**76.74** ± 2.86	**75.17** ± 1.21	**53.44** ± 3.21	**93.34** ± 1.95	**89.23** ± 3.78	**79.19** ± 2.23	**88.41** ± 1.43	**85.24** ± 1.70	**72.78** ± 1.31
Improv.	*vs. Best*	+0.82	-0.10	**+2.70**	+3.19	+3.32	**+5.61**	+0.59	+0.38	**+2.44**	+2.97	+2.92	**+9.18**
	vs. Mean	+4.49	+3.95	**+13.37**	+5.72	+6.51	**+14.08**	+8.37	+10.66	**+17.38**	+6.19	+7.04	**+14.94**

identities, MarginCell improves the overall average performance by 9.18% over the best competing method, further demonstrating its robustness to fragile class competition and difficult identities.

Across different random seeds, MarginCell maintains stable and consistent cross-environment performance, indicating favorable training stability. Overall, these results further validate that MarginCell can release controlled evidence variation while reliably preserving the decision reserve against the weakest rivals, thereby improving cross-environment generalization robustness.

C.3 Plug-in Compatibility Performance

Table 5 reports the plug-in performance of MarginCell with representative DG and robust-RFFI objectives on Acc, WEnv, and $TIA_{@20}$ across benchmarks. MarginCell consistently improves Acc and WEnv across all baseline-dataset combinations, with average gains of 3.94% and 4.67%, respectively. **Such consistent improvements across datasets and training paradigms demonstrate that MarginCell can serve as a broadly complementary regularizer for existing cross-environment learning objectives.**

The gains are even more pronounced on hard identities. MarginCell improves $TIA_{@20}$ in most settings, with an average gain of 12.72%. **This larger tail-identity improvement shows that regulating true-vs-rival evidence is particularly effective at alleviating the fragile class competition faced by difficult identities under environment shifts.** By releasing bounded intra-identity variation while preserving the weakest-rival reserve, MarginCell strengthens the part of the decision geometry that is most vulnerable to cross-environment degradation.

C.4 Source-Target Evidence Geometry

Figure 8 visualizes how different evidence-control strategies transfer from source to target in a two-dimensional reserve plane. To expose the high-dimensional rival-wise evidence geometry, each sample is projected by the minimum normalized true-vs-rival reserve over rivals 3 and 7 on the

Table 5: MarginCell complements diverse domain-generalization objectives across datasets. Positive Δ values indicate performance improvements. (**Bold**: Best; Underline: Second-best).

Dataset	Base Objective	Accuracy (%)		WEnv (%)		$TIA_{@20}$ (%)	
		+MC	Δ	+MC	Δ	+MC	Δ
WiSig	GroupDRO	**97.92**	**+9.77**	**96.17**	**+9.44**	**94.00**	**+29.20**
	V-REx	91.17	+1.25	89.33	+1.20	73.75	+3.65
	MMD-DG	97.17	+8.15	94.50	+8.03	91.75	+24.35
	CLPS	96.17	+1.84	92.33	+2.86	88.50	+5.50
	CIS-RFF	95.22	+2.00	92.17	+0.74	79.25	-0.70
UCLA-CORES	GroupDRO	77.38	**+8.14**	**75.85**	**+9.86**	55.46	+10.69
	V-REx	70.29	+0.34	70.28	+1.37	42.37	+0.50
	MMD-DG	74.25	+2.41	73.06	+4.95	46.60	-0.05
	CLPS	**79.68**	+6.35	74.19	+3.28	50.33	+7.69
	CIS-RFF	71.17	+2.97	70.90	+3.58	**56.13**	**+26.81**
LoRa	GroupDRO	93.12	+3.86	88.80	+5.70	77.38	+15.19
	V-REx	93.27	+3.16	89.05	+4.17	73.62	+8.18
	MMD-DG	**95.60**	+2.85	**92.55**	+3.70	**89.75**	+13.00
	CLPS	93.88	**+5.70**	89.10	**+6.05**	78.38	+19.94
	CIS-RFF	89.30	+0.25	87.85	+5.10	88.25	**+24.81**

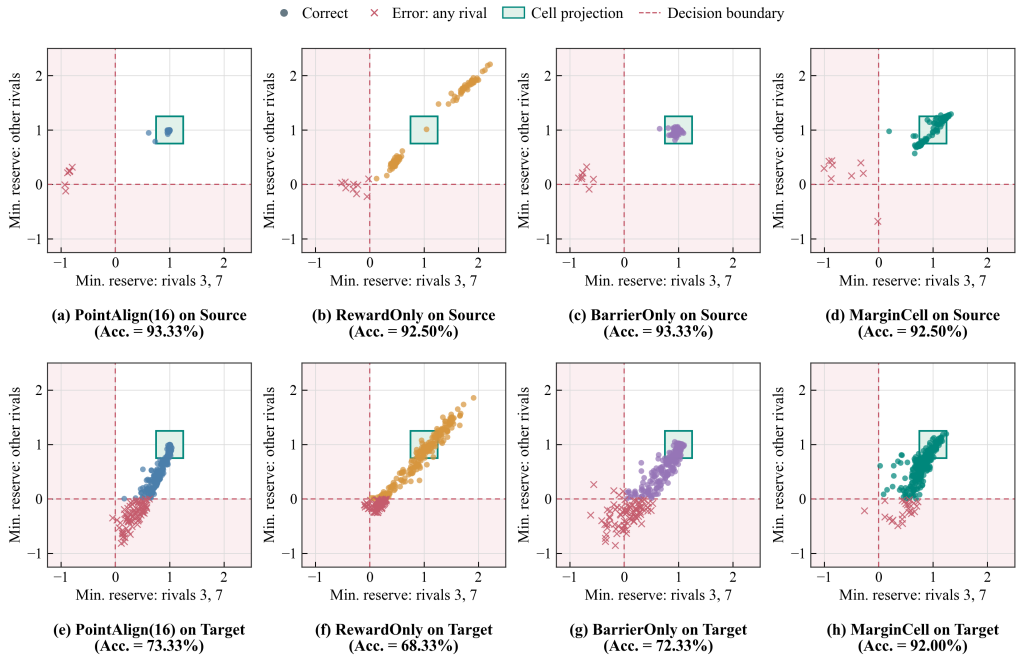


Figure 8: Visualization of Source-Target Decision-Evidence Geometry.

horizontal axis, and the corresponding minimum over all remaining rivals on the vertical axis. The green box denotes the projection of the admissible evidence cell, while the dashed lines mark zero-reserve decision boundaries; crossing either boundary therefore indicates that at least one rival has exhausted its decision reserve. The top and bottom rows show the resulting source and target geometries, respectively.

A clear source-target contrast emerges. On the source domain, PointAlign_{16}, Barrier-only, and MarginCell all achieve similarly high accuracy, indicating that source fitting alone does not distin-

guish whether the learned evidence geometry will remain robust. After transferring to new environment, however, PointAlign$_{16}$ collapses evidence tightly around the source reference but exhibits substantial boundary crossing on the target domain, reducing accuracy to 73.33%. This shows that enforcing stronger source-side evidence consistency does not create the reserve needed to absorb unseen shifts. Reward-only exhibits the opposite failure mode: although it releases evidence variation, the absence of boundary control allows the target distribution to expand along reserve-eroding directions, resulting in only 68.33% accuracy. Barrier-only limits this expansion, yet many target samples still approach or cross the zero-reserve boundaries, resulting in 72.33% accuracy.

In contrast, MarginCell preserves a finite spread of source evidence while keeping most target samples on the positive-reserve side of both decision boundaries, reaching 92.00% target accuracy despite comparable source accuracy. The visualization therefore makes the role of release-and-bound explicit: the dispersion reward prevents evidence from collapsing to a single source state, while the boundary barrier prevents the released variation from consuming the weakest-rival reserve. **Robust transfer thus comes neither from evidence equality nor from unconstrained dispersion, but from preserving bounded decision-equivalent freedom across environments, directly supporting the central "Same Identity, Relaxed Evidence" insight.**

C.5 Dataset-wise Component Ablation

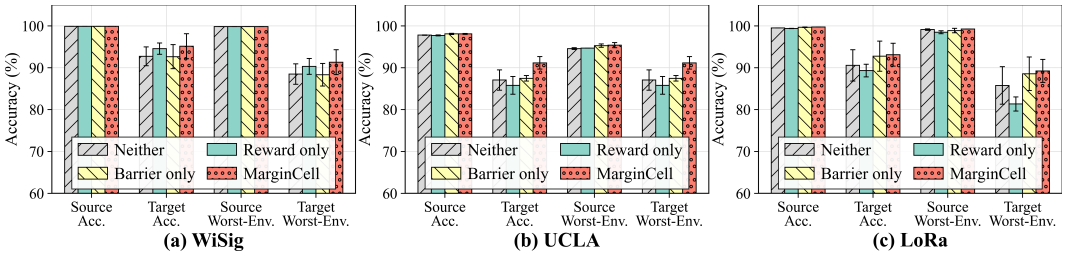


Figure 9: Ablation on Reward and Barrier Components.

Figure 9 presents the component ablation of MarginCell across WiSig, UCLA-CORES, and LoRa in terms of source and target Acc/WEnv. It can be observed that source performance remains largely unchanged across variants, whereas the target metrics are considerably more sensitive to the two components. **Across all three benchmarks, the full MarginCell achieves the strongest overall target performance**, showing that the improvement mainly stems from enhanced cross-environment robustness rather than better source fitting.

The dataset-wise results further reveal the complementary roles of the two components. On WiSig, Reward-only already improves target performance, indicating that releasing CE-contracted evidence freedom can be beneficial under relatively mild shifts. On UCLA-CORES and LoRa, however, Reward-only becomes less reliable, suggesting that unconstrained release can introduce reserve-eroding deviations. Barrier-only provides more conservative control but does not actively exploit admissible evidence variation. By combining both effects, MarginCell releases useful within-identity variation while constraining unsafe excursions, yielding more consistent robustness across different environment shifts.