

RethinkingIRSTD: Single-Point Supervision Guided Encoder-Only Framework Is Enough for Infrared Small Target Detection

Rixiang Ni¹, Jun Chen¹, Boyang Li¹, Yonghao Li¹, Wujiao He¹, Yuji Wang¹, Feiyu Ren¹, Haoyang Yuan¹, Wei An¹

National University of Defense Technology, Changsha 410003, China
{nirixiang, chenjun11, liboyang20}@nudt.edu.cn
Code and benchmark: <https://github.com/NIRIXIANG/SPIRE-IRSTD>

Abstract. Infrared small target detection (IRSTD) aims to localize small targets in cluttered infrared scenes. Extensive research follows the pixel-level supervision-guided ‘encoder–decoder’ segmentation paradigm. Although these methods have achieved promising performance, they often overlook that infrared small targets occupy only a few pixels and are usually surrounded by blurred, low-contrast boundaries caused by cluttered backgrounds. Based on this observation, we argue that the first principle of IRSTD should be target localization rather than dense reconstruction of target regions entangled with indistinguishable background noise. In this paper, we reformulate IRSTD as a centroid regression task and propose a novel Single-Point Supervision guided Infrared Probabilistic Response Encoding method (namely, SPIRE), which is non-trivial because point-level supervision must produce detection outputs comparable to dense supervision. Specifically, we design Point-Response Prior Supervision (PRPS) to transform single-point annotations into probabilistic response maps consistent with infrared point-target response characteristics, and combine it with a High-Resolution Probabilistic Encoder (HRPE) that performs encoder-only probabilistic regression followed by lightweight peak-based inference, without decoder reconstruction. By preserving high-resolution features and increasing effective supervision density, SPIRE alleviates optimization instability under sparse target distributions. Finally, extensive experiments on two public IRSTD benchmarks, SIRST-UAVB and SIRST4, demonstrate that SPIRE achieves competitive target-level detection performance with consistently low false alarm rate (Fa) and significantly reduced computational cost.

Keywords: Infrared Small Target Detection · Probabilistic Response Modeling · Single-Point Supervision · Encoder-Only Architecture.

 Corresponding authors.

1 Introduction

Infrared Small Target Detection (IRSTD) is a fundamental component of infrared search and tracking systems [35], with wide applications in traffic monitoring [17, 40], early warning [20, 31], and maritime rescue [39, 42]. Despite its practical importance, most existing IRSTD approaches are formulated under dense prediction paradigms, typically adopting encoder–decoder segmentation architectures trained with pixel-level annotations [17, 18, 21, 25, 27–30, 32, 34–36].

Three limitations of the existing fully supervised encoder–decoder segmentation paradigm. Although this paradigm has achieved promising performance, it introduces three inherent limitations. (1) *Heavy Annotation Cost.* Pixel-wise labeling is labor-intensive and subject to boundary and scale ambiguities (see Fig. 1(a)). [18] reports that single-point annotation reduces labeling time by 87.72%, from 11.4 s to 1.4 s per image compared with pixel-wise annotation. (2) *Redundant Encoder–Decoder Framework.* To compensate for extreme foreground sparsity and stabilize optimization, multi-scale fusion and densely connected modules (e.g., nested structures, attention mechanisms) are widely employed, substantially increasing training and inference complexity (see Fig. 1(c)). (3) *Localization Uncertainty.* Instance-level localization is typically obtained through additional post-processing of segmentation masks, introducing pipeline overhead and potential error accumulation (see Fig. 1(c)).

The first principle of IRSTD & two trade-off solutions. Infrared small targets (IRSTs) are characterized by extremely limited spatial extent, Gaussian-like energy diffusion, low signal-to-noise ratios, and highly sparse spatial distributions. Considering the practical IRST search and tracking systems primarily require reliable localization rather than contour reconstruction, *we argue that the first principle of IRSTD should be centroid-oriented detection instead of dense pixel-level reconstruction.* Existing efforts to alleviate dense segmentation can be broadly categorized into two directions. (1) *Bounding-box supervised regression frameworks.* These works adopt bounding box-based detection frameworks to simplify prediction [3, 9, 19, 33]; however, they still rely on bounding-box annotations and do not explicitly model the intrinsic local response structure of infrared small targets. (2) *Single-point supervision-guided segmentation frameworks.* These works try to reduce annotation cost through point-level weak supervision with pseudo-mask generation [15, 18, 35, 36], yet they continue to depend on encoder–decoder architectures and dense mask reconstruction. Although these approaches partially ease specific burdens, they do not simultaneously resolve the coupled bottlenecks of annotation efficiency, structural complexity, and localization stability.

Combining single-point supervision with an encoder-only architecture is conceptually simple, but deriving equally informative outputs from much sparser labels is non-trivial. We observe that infrared small targets contain limited global semantic information but remain locally salient within compact neighborhoods. Motivated by this observation, we propose a novel Single-Point Supervision guided Infrared Probabilistic Response Encoding method (namely, SPIRE), a probabilistic response modeling framework that reformulates IRSTD as spatial




Fig. 1: Motivation of SPIRE forIRSTD. Existing IRSTD methods suffer from three key limitations: (a) heavy pixel-level annotation cost with extremely sparse positive samples, (b) redundant encoder-decoder architectures relying on extensive feature fusion and attention mechanisms, and (c) localization uncertainty caused by segmentation-based predictions. SPIRE reformulates IRSTD as centroid localization with single-point supervision and probabilistic response map modeling, enabling efficient annotation, a lightweight encoder-only framework, and precise peak-based localization.

probability regression centered on target localization, rather than reconstructing dense pixel regions. SPIRE models structured local probabilistic responses consistent with infrared imaging physics and centroid-oriented detection objectives. Specifically, **(1) Single-point supervision.** SPIRE first combines single-point annotations with small-target characteristic priors, such as local saliency and Gaussian-like energy distribution. It then transforms sparse point labels into probabilistic response maps via Point-Response Prior Supervision (PRPS), reformulating target learning as structured local probability modeling. The supervision encodes a Gaussian-like spatial prior capturing center-enhanced and radially decaying energy responses while incorporating contrast-aware constraints to approximate realistic radiometric intensity transitions. Consequently, the network learns structured spatial probability distributions rather than discrete pixel labels, substantially increasing effective supervision density, alleviating gradient sparsity under extreme class imbalance, stabilizing optimization, and significantly reducing annotation cost. **(2) Encoder-only and peak-extraction inference framework.** Once target discrimination is modeled through structured probabilistic responses rather than pixel-wise reconstruction, the necessity for heavy multi-scale fusion and attention-based compensation is substantially reduced. SPIRE adopts a High-Resolution Probabilistic Encoder (HRPE) that avoids aggressive down-sampling and cross-scale branches, thereby preserving the continuity and separability of local probabilistic responses. Built with




Fig. 2: Performance of SPIRE on SIRST-UAVB. (a) Performance–efficiency trade-off among representativeIRSTD methods, where the axes denote F1-score and Precision, and bubble size indicates the number of parameters. (b) Six-dimensional evaluation across Precision, Recall (Pd), F1-score, False Alarm rate, FLOPs, and parameters. SPIRE achieves a favorable balance between detection accuracy and computational efficiency.

cascaded lightweight channel-reorganization units, the encoder achieves expressive yet computationally efficient representation learning. Inference is performed by regressing probabilistic response maps and then applying fast deterministic peak extraction as lightweight post-processing, avoiding mask reconstruction and clustering-based post-processing.

The contributions of this paper can be summarized as:

- Inspired by the intrinsic characteristics and engineering objectives ofIRSTD, we rethink infrared small target detection as a centroid-oriented spatial probability regression task that can be learned using only point annotations, and accordingly establish a probabilistic response modeling paradigm termed SPIRE.
- We introduce a Point-Response Prior Supervision (PRPS) mechanism that converts sparse point labels into probabilistic response maps consistent with infrared imaging physics, inherently increasing supervision efficiency while avoiding pixel-level or box-level annotations.
- We design HRPE as a lightweight single-encoder module. It removes decoder branches and relies only on fast deterministic peak extraction as lightweight post-processing, enabling efficient centroid prediction.
- Extensive experiments demonstrate that SPIRE achieves competitive or superior detection accuracy with consistently low false alarms, while naturally maintaining lower annotation cost and computational complexity, as illustrated in Fig. 2.

2 Related Work

2.1 Infrared Small Target Detection

Infrared small targets usually occupy only a few pixels and are often embedded in cluttered backgrounds with extremely low signal-to-noise ratio (SNR), which makes stable target representation and foreground-background discrimination particularly challenging. Early model-drivenIRSTD methods [2, 4–7, 10, 13, 14] addressed this problem by exploiting handcrafted priors such as local contrast enhancement, background suppression, and morphological filtering. With the emergence of annotated datasets, deep learning-based approaches [17, 18, 21, 25, 27–30, 32, 34–36] have gradually become the dominant paradigm. CNN-based models commonly employ multi-scale feature fusion and dense cross-layer connections to preserve small target responses [17, 21, 27, 28, 30, 32], while more recent works incorporate transformer modules and attention mechanisms to model long-range contextual dependencies and suppress background interference [25, 37]. Despite these advances, mostIRSTD approaches still rely on dense pixel-wise segmentation. Such formulations struggle with extreme target sparsity and often depend on heavy architectures and feature aggregation for stable detection.

2.2 Single-Point Supervision

To alleviate the high cost and ambiguity of pixel-wise annotation, recent studies introduce single-point supervision intoIRSTD, where only target centroids are annotated. LESPS [35] first explored this idea by generating pseudo masks from point annotations to supervise a segmentation network, reducing labeling cost while maintaining compatibility with existing architectures. Building upon this idea, MCLC [18] further introduced multi-level constraints to refine pseudo-label generation and improve the quality of supervision derived from sparse annotations. More recently, FromEasy2Hard [36] proposed a curriculum-style learning strategy that progressively enhances supervision quality from coarse point guidance to refined target localization. Although these methods reduce annotation effort, they largely remain within dense segmentation frameworks. Point annotations are typically converted into pseudo masks or auxiliary supervision, without fundamentally redefining detection as sparse centroid localization.

2.3 Heatmap-Based Modeling

Heatmap-based modeling is widely used in keypoint localization tasks such as human pose estimation and landmark detection [1, 8, 23, 26]. These methods represent target locations as Gaussian heatmaps centered on annotated points and train networks to regress spatial confidence distributions. Representative approaches such as HRNet [26] and HigherHRNet [8] achieve accurate localization by maintaining high-resolution representations or aggregating multi-scale heatmaps. A few studies introduce heatmap supervision intoIRSTD. For example, P2P-HDNet [22] formulatesIRSTD as point-to-point heatmap prediction,




Fig.3: Framework of SPIRE (Single-Point Supervision guided Infrared Probabilistic Response Encoding). Given an infrared image with single-point centroid annotations, SPIRE constructs structured supervision through Point-Response Prior Supervision (PRPS), which maps centroids to the feature space and generates Gaussian-like response priors inspired by infrared PSF diffusion. A lightweight High-Resolution Probabilistic Encoder (HRPE) then regresses the probabilistic response map (heatmap). During inference, target locations are obtained by lightweight peak extraction on the predicted heatmap via local-maximum detection, thresholding, and sub-pixel centroid refinement.

but still relies heavily on feature fusion and attention-based compensation. More importantly, existing approaches rarely integrate the physical formation characteristics of infrared point targets, such as point spread function (PSF)-induced energy diffusion, into heatmap modeling or supervision design.

3 Methodology

3.1 Overall Framework

Infrared small targets occupy only a few pixels and exhibit Gaussian-like energy diffusion under the combined effects of the PSF and sensor sampling, while practical search and tracking systems primarily require reliable centroid localization rather than contour reconstruction. Conventional encoder-decoder segmentation paradigms are therefore structurally mismatched: dense pixel-wise supervision induces severe foreground-background imbalance and gradient sparsity, necessitating increasingly complex multi-scale fusion and attention compensation.

As shown in Fig. 3, SPIRE reformulates IRSTD as centroid-oriented spatial probability regression, aligning both supervision and architecture with in-

frared imaging physics and detection objectives. Given an input infrared image $\mathbf{I} \in \mathbb{R}^{H \times W \times 1}$, HRPE produces a single-channel probabilistic response map $\hat{\mathbf{P}} \in \mathbb{R}^{H' \times W'}$ at reduced resolution

$$H' = \lfloor H/s \rfloor, \quad W' = \lfloor W/s \rfloor, \quad (1)$$

where s denotes the output stride.

The ground-truth supervision map $\mathbf{P}^* \in \mathbb{R}^{H' \times W'}$ is constructed offline by PRPS from single-point annotations $\{(x_k, y_k)\}_{k=1}^K$, encoding structured local probabilistic responses consistent with infrared point-target response characteristics (Sec. 3.2). Target centroid locations are obtained via a lightweight deterministic peak-extraction procedure on $\hat{\mathbf{P}}$ (Sec. 3.4), rather than mask reconstruction or iterative post-processing.

3.2 Point-Response Prior Supervision

Under PSF-induced blur and imaging degradation, infrared small targets exhibit locally concentrated responses with radially decaying intensity rather than isolated impulses. PRPS encodes a point-response prior together with observed local radiometric structure, generating structured local probabilistic responses that capture isotropic diffusion and contrast-aware radiometric transitions.

Stride-aware coordinate mapping. Each annotated centroid (x_k, y_k) is mapped to the heatmap lattice via

$$(\tilde{x}_k, \tilde{y}_k) = \left\lfloor \left(\frac{x_k}{s}, \frac{y_k}{s} \right) + 0.5 \right\rfloor, \quad (2)$$

where s denotes the output stride. The centroid is further refined to the position of maximum intensity within a local neighborhood in the original image, yielding $(\tilde{x}_k^*, \tilde{y}_k^*)$ aligned with the true radiometric peak.

Gaussian spatial prior. An isotropic Gaussian kernel centered at $(\tilde{x}_k^*, \tilde{y}_k^*)$ encodes PSF-induced energy diffusion:

$$G_k(u, v) = \exp\left(-\frac{(u - \tilde{x}_k^*)^2 + (v - \tilde{y}_k^*)^2}{2\sigma^2}\right), \quad \|(u, v) - (\tilde{x}_k^*, \tilde{y}_k^*)\|_\infty \leq r, \quad (3)$$

where σ controls spatial spread and $r = 3\sigma$ defines the compact support radius beyond which the kernel is truncated to zero. The kernel has unit peak and is not normalized. PRPS is not intended to reconstruct target contours; instead, it provides a centroid-oriented local probabilistic response. We set $\sigma = 2$ and $r = 6$ in all experiments. This scale choice is empirically examined in Sec. 4.3 and Tab. 3.

Contrast-aware modulation. A $(2r+1) \times (2r+1)$ patch centered at $(\tilde{x}_k^*, \tilde{y}_k^*)$ is extracted from the original image and contrast-normalized:

$$C_k(u, v) = \frac{\mathbf{I}_{\text{patch}}(u, v) - \mathbf{I}_{\text{patch}}^{\min}}{\mathbf{I}_{\text{patch}}^{\max} - \mathbf{I}_{\text{patch}}^{\min}}, \quad (4)$$

where $\mathbf{I}_{\text{patch}}^{\min}$ and $\mathbf{I}_{\text{patch}}^{\max}$ denote the minimum and maximum intensities within the patch. The per-target response is then composed as

$$H_k = \text{Norm}(G_k \odot C_k), \quad (5)$$

where $\odot$ denotes element-wise multiplication and $\text{Norm}(\cdot)$ performs min-max normalization to $[0, 1]$. This preserves the center-enhanced diffusion pattern while adapting to the observed local radiometric contrast.

Multi-target aggregation. For images containing K targets, responses are aggregated into a single supervision map:

$$\mathbf{P}^*(u, v) = \mathcal{A}(\{H_k(u, v)\}_{k=1}^K), \quad (6)$$

where $\mathcal{A}$ denotes deterministic aggregation. Even when Gaussian supports overlap, centroid-oriented regression preserves distinct local maxima for each target, enabling deterministic peak extraction without mask clustering. Overlapping supports therefore do not necessarily merge nearby targets when separable local maxima are preserved. By expanding sparse point labels into spatially structured responses consistent with infrared point-target response characteristics, PRPS increases supervision density and alleviates gradient starvation under extreme foreground-background imbalance.

Training objective. The network is optimized via standard pixel-wise regression between $\hat{\mathbf{P}}$ and $\mathbf{P}^*$. Because PRPS provides dense structured supervision within each kernel support, sufficient gradient contributions are maintained to ensure stable convergence without specialized imbalance-aware losses.

3.3 High-Resolution Probabilistic Encoder

Once target discrimination is grounded in structured local probabilistic responses rather than pixel-wise mask reconstruction, heavy multi-scale fusion and decoder-based compensation become unnecessary. HRPE is a purely encoder-only, single-branch architecture that maintains high-resolution representations at a constant output stride s . This design is critical for IRSTD, as aggressive downsampling may destroy the few-pixel support of dim targets and erase sub-pixel centroid cues under low signal-to-noise ratio (SNR) and low signal-to-clutter ratio (SCR) conditions.

Stem. Two successive 3×3 convolutions with stride 2, each followed by batch normalization and ReLU, reduce the resolution to $H/s \times W/s$ while expanding channels to $C_0=64$. A bottleneck residual block and lightweight channel reorganization units further refine features at the same resolution.

Cascaded channel reorganization. After a transition layer reduces channels to $C=32$, features pass through cascaded lightweight channel reorganization units operating on a single-resolution branch. Given input $\mathbf{F} \in \mathbb{R}^{C \times H' \times W'}$, features are evenly split into $\mathbf{F}_a, \mathbf{F}_b \in \mathbb{R}^{C/2 \times H' \times W'}$. One branch is preserved as identity, while the other undergoes depthwise separable convolution and channel-wise attention reweighting:

$$\mathbf{F}' = \text{Shuffle}\left(\text{Cat}\left(\mathbf{F}_a, \phi_{\text{attn}}(\text{DWConv}(\mathbf{F}_b))\right)\right), \quad (7)$$

where DWConv denotes 3×3 depthwise convolution with batch normalization, ϕ_{attn} applies squeeze-and-excitation-style reweighting, and Shuffle interleaves channels for cross-group information exchange. Some units further include an additional depthwise convolution for local spatial refinement. This split-process-shuffle design [41] enables expressive yet efficient channel reorganization while preserving spatial continuity and separability for structured probabilistic response learning.

No multi-resolution branches, decoder pathways, or skip connections are used, yielding a strictly encoder-only architecture aligned with centroid-oriented probabilistic regression.

Prediction head. A single 1×1 convolution maps the C -channel representation to one output channel, producing $\hat{\mathbf{P}} \in \mathbb{R}^{H' \times W'}$. No activation is applied; outputs are directly regressed to the PRPS target.

3.4 Lightweight Probabilistic Inference

At inference time, target centroid locations are extracted from $\hat{\mathbf{P}}$ using fast deterministic peak extraction as lightweight post-processing, rather than mask reconstruction, connected-component analysis, or learned post-processing.

Local-maximum extraction. Following standard practice [26], a 3×3 max-pooling non-maximum suppression (NMS) with stride 1 suppresses non-maximum responses:

$$\hat{\mathbf{P}}_{\text{nms}}(u, v) = \begin{cases} \hat{\mathbf{P}}(u, v), & \hat{\mathbf{P}}(u, v) = \max_{(u', v') \in \mathcal{N}_{3 \times 3}(u, v)} \hat{\mathbf{P}}(u', v'), \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

Thresholding and candidate selection. We first collect NMS responses whose confidence exceeds τ :

$$\mathcal{M}_\tau = \{(u, v) \mid \hat{\mathbf{P}}_{\text{nms}}(u, v) > \tau\}. \quad (9)$$

Here $\mathcal{M}_\tau$ denotes the thresholded local-maximum set, $\mathcal{N}_{3 \times 3}(u, v)$ is the NMS neighborhood, and τ is the confidence threshold. Experiments show that τ has a broad stable range and does not require dataset-specific retuning.

Final detections are then selected by response-ranked top- N_{max} :

$$\mathcal{D} = \text{top-}N_{\text{max}}(\mathcal{M}_\tau; \hat{\mathbf{P}}_{\text{nms}}). \quad (10)$$

Here $\mathcal{D}$ is the predicted centroid set, and top- N_{max} returns at most N_{max} peaks with the highest NMS response scores. We set N_{max} to the maximum target count of each dataset; increasing it to several times that count leaves the metrics unchanged because extra low-score peaks are filtered by τ .

Sub-pixel centroid refinement. Each integer-coordinate peak $(u_0, v_0) \in \mathcal{D}$ is refined using the sign of the local finite-difference gradient on the pre-NMS response map [26]:

$$\begin{aligned}\Delta u &= \frac{1}{s} \text{sign}(\hat{\mathbf{P}}(u_0+1, v_0) - \hat{\mathbf{P}}(u_0-1, v_0)), \\ \Delta v &= \frac{1}{s} \text{sign}(\hat{\mathbf{P}}(u_0, v_0+1) - \hat{\mathbf{P}}(u_0, v_0-1)).\end{aligned}\tag{11}$$

Here s denotes the output stride. In this work we set $s = 4$, as validated in Sec. 4.3, achieving a favorable balance between spatial fidelity and computational efficiency. This finite-difference refinement follows standard heatmap localization and serves as a lightweight deterministic correction.

Coordinate back-mapping. Refined probabilistic response map coordinates are mapped to the original image space via the inverse affine transformation:

$$\begin{pmatrix} x_{\text{det}} \\ y_{\text{det}} \end{pmatrix} = \mathbf{T}^{-1} \begin{pmatrix} u_0 + \Delta u \\ v_0 + \Delta v \\ 1 \end{pmatrix},\tag{12}$$

where $\mathbf{T}^{-1} \in \mathbb{R}^{2 \times 3}$ restores coordinates from the stride-reduced probabilistic response map to the original infrared image grid, preserving geometric consistency between centroid-oriented spatial probability regression and the physical image coordinate system.

The complete procedure—a single forward pass through HRPE followed by fast deterministic peak extraction as lightweight post-processing—avoids mask reconstruction and iterative segmentation post-processing, directly yielding target centroid coordinates consistent with the centroid-oriented spatial probability regression formulation of SPIRE.

4 Experiment

4.1 Experimental Setup

Datasets. Experiments are conducted on two public benchmarks, SIRST-UAVB and SIRST4 [34]. SIRST-UAVB contains 3000 images and SIRST4 contains 3352 images; their train/test splits are 2400/600 and 2285/1067, respectively. SIRST4 integrates SIRSTv1 [11], SIRSTv2 [9], NUDT-SIRST [17], and IRSTD-1K [40], covering diverse scenes and target scales. SIRST-UAVB focuses on UAV-based urban scenes with clutter and low SNR. Both datasets include real and simulated targets, isolated and clustered targets, as well as empty-background images. Mask-level and point-level annotations are provided; we use centroid point annotations for supervision and retain mask annotations for fair comparison.

Implementation details. All models are implemented in PyTorch and trained on an RTX 4090 GPU. We adopt Adam with an initial learning rate of 0.005 and a ReduceLROnPlateau scheduler (factor 0.01, patience 3). The batch size is 10. Each competing method follows its original training schedule, while our model is trained for 500 epochs. Input resolutions are 512×512 for SIRST4 and 640×640 for SIRST-UAVB.

Evaluation metrics. IRSTD is evaluated at the centroid level. Segmentation outputs are converted to centroids via 8-connected clustering [17], bounding-box detectors use box centers, and probabilistic response map-based methods use peak extraction. A prediction is counted as a true positive (TP) if its Euclidean distance to a ground-truth centroid is within $\delta = 5$ pixels; otherwise it is a false positive (FP). We use this threshold as the default protocol and later verify the robustness to δ in Sec. 4.3 and Table 2. Unmatched ground-truth targets are counted as false negatives (FN).

We report Precision, Recall, and F1-score, defined as

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad (13)$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}, \quad (14)$$

where Recall is equivalent to the probability of detection (Pd) in IRSTD. The F1-score is

$$\text{F1} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}. \quad (15)$$

The false alarm rate (Fa) is defined as

$$\text{Fa} = \frac{\text{FP}}{N_{\text{pixels}}}, \quad (16)$$

where N_{pixels} denotes the total number of image pixels. We additionally report floating-point operations (FLOPs) and the number of parameters (Params) for efficiency comparison.

4.2 Comparison with State of the Art

Tab. 1 reports quantitative results on SIRST-UAVB and SIRST4, and Fig. 4 provides qualitative comparisons.

Quantitative comparison. Table 1 reports quantitative results under the unified centroid-level matching protocol. On SIRST-UAVB, our method achieves the highest Precision of 99.82 and the lowest Fa of 1.02, while maintaining F1 of 97.05. Compared with DNANet and SCTransNet, which report Fa of 15.77 and 5.09 respectively, our method consistently suppresses false alarms while using substantially fewer parameters. On SIRST4, our method achieves F1 of 94.60 while requiring over 8 times fewer FLOPs than SCTransNet and more than 16 times fewer parameters than DNANet. These results demonstrate that centroid-oriented spatial probability regression achieves competitive detection accuracy while naturally decoupling structural complexity from performance within an encoder-only framework. To complement the FLOPs and Params in Table 1, we further benchmark frames per second (FPS) for representative high-accuracy competitors, as summarized in Table 2. SPIRE reaches 261.2 FPS, running $3.23\times$ faster than DNANet, $3.89\times$ faster than SCTransNet, and $5.56\times$ faster than SD-SNet. This confirms that the encoder-only design improves practical throughput while maintaining competitive detection accuracy.




Fig. 4: Qualitative comparison of SPIRE on SIRST-UAVB and SIRST4. Red rectangles denote ground-truth targets. Red circles indicate true positives, green circles denote false positives, and orange circles correspond to false negatives under the unified centroid protocol.

Qualitative comparison. Figure 4 visualizes representative detection results. Red rectangles denote ground-truth targets. Red circles indicate true positives, green circles denote false positives, and orange circles correspond to false negatives under the unified centroid protocol. Segmentation-based methods such as DNANet and SCTransNet often produce multiple adjacent mask responses around a single target. After connected-component clustering, these redundant responses may generate extra centroids near true targets, increasing false alarms. In contrast, our structured probabilistic response regression produces a single dominant peak per target without neighboring over-activation, resulting in cleaner centroid localization and consistently lower Fa.

Overall, both quantitative and qualitative evidence validate that sparse centroid modeling with physically grounded probabilistic supervision provides a more efficient and principled alternative to dense mask reconstruction forIRSTD.

4.3 Ablation and Analysis

We first assess the robustness of the evaluation protocol, and then analyze the key design choices of SPIRE.

Robustness to matching threshold δ . Since centroid-level evaluation depends on the matching tolerance, we examine whether the target-level conclusion

Table 1: Comparison with representative methods on SIRST-UAVB and SIRST4. We report Pre(%), Rec (Pd)(%), F1(%), Fa (10^{-8}), Params (M), and FLOPs (G). Params and FLOPs are computed at 640×640 input resolution.

Method	Venue	SIRST-UAVB (2400:600)				SIRST4 (2285:1067)				FLOPs↓	Params↓
		Pre ↑	Rec ↑	F1 ↑	Fa ↓	Pre ↑	Rec ↑	F1 ↑	Fa ↓		
ACM [11]	WACV ²¹	87.01	71.16	78.29	34.04	90.17	70.38	79.05	44.08	<u>2.51</u>	0.40
ALCNet [12]	TGRS ²¹	95.06	81.11	87.53	12.72	91.45	72.41	80.82	38.90	2.36	0.43
ISTDU-Net [16]	GRSL ²²	79.28	78.08	78.67	61.54	<u>93.28</u>	72.03	81.29	29.83	49.65	2.75
RDIAN [24]	TGRS ²³	52.68	77.91	62.86	211.08	90.64	80.83	85.45	47.98	23.24	0.22
DNANet [17]	TIP ²³	94.48	89.54	91.95	15.77	93.99	81.20	87.13	<u>29.82</u>	89.13	4.70
SCTransNet [37]	TGRS ²⁴	<u>98.27</u>	<u>95.95</u>	97.09	<u>5.09</u>	81.20	86.39	83.72	114.98	63.22	11.19
MSHNet [21]	CVPR ²⁴	71.90	88.02	79.09	104.27	90.68	<u>92.86</u>	91.47	60.51	38.16	4.07
SDSNet [38]	TGRS ²⁵	97.60	96.29	96.94	7.12	87.35	88.27	87.81	73.48	42.42	2.49
L ² SKNet [28]	TGRS ²⁵	95.83	89.04	92.31	11.70	92.89	91.42	<u>92.16</u>	40.20	43.09	0.90
Ours	-	99.82	94.44	<u>97.05</u>	1.02	95.00	94.21	94.60	28.53	7.68	<u>0.29</u>

Table 2: Multi-threshold robustness and inference speed comparison on SIRST4. We report Rec (Pd)(%), F1(%), Fa (10^{-8}) under centroid-matching thresholds, together with FPS measured under the same input and hardware setting.

Method	$\delta = 3$			$\delta = 5, 8, 10$			FPS ↑
	Rec ↑	F1 ↑	Fa ↓	Rec ↑	F1 ↑	Fa ↓	
DNANet [17]	80.52	86.41	33.72	81.20	87.13	29.82	80.86
SCTransNet [37]	80.15	81.18	99.42	86.39	83.72	114.98	67.09
SDSNet [38]	86.39	87.21	67.43	88.27	87.81	73.48	46.97
Ours	92.93	94.74	18.61	94.21	94.60	28.53	261.2

is sensitive to this setting. On SIRST4, Table 2 shows that the stricter $\delta = 3$ setting reduces Recall for all methods, as the tolerance becomes comparable to the average target radius of 3.31 pixels. SPIRE remains most stable in F1 (94.74 vs. 94.60 under $\delta = 5, 8, 10$). The identical results for $\delta = 5, 8, 10$ indicate that increasing the tolerance beyond 5 does not change the evaluated matches, validating $\delta = 5$ as a reasonable default.

Impact of PRPS supervision form. We compare PRPS with single-point impulse and unconstrained Gaussian supervision. As shown in Tab. 3, impulse supervision decreases Recall from 94.44 to 90.22 and increases Fa from 1.02 to 2.85, reducing F1 from 97.05 to 94.27. Gaussian supervision improves Recall to 93.76 but still yields higher Fa of 2.44 and lower F1 than PRPS. As shown in Fig. 5, impulse supervision suffers from gradient starvation under extreme foreground-background imbalance, while isotropic Gaussian lacks the response-adaptive modulation guided by local radiometric structure, leading to less discriminative peaks in clutter. By encoding Gaussian responses consistent with infrared point-target characteristics, PRPS enables stable centroid-oriented spatial probability regression and suppresses spurious peaks.

Impact of PRPS scale. We vary σ to analyze scale sensitivity. As shown in Tab. 3, $\sigma = 2$ achieves the best F1 of 97.05. When $\sigma = 1$, Recall drops to 88.60 and Fa rises to 5.26; when $\sigma = 3$, Fa decreases to 0.67 but Recall declines to 92.50.

Table 3: Ablation study of PRPS on SIRST-UAVB. The first block compares different supervision forms, while the second block analyzes the scale sensitivity of PRPS. We report Pre(%), Rec (Pd)(%), F1(%), and Fa (10^{-8})

Variant	Pre $\uparrow$	Rec $\uparrow$	F_1 $\uparrow$	Fa $\downarrow$
<i>Supervision form</i>				
PRPS	99.82	94.44	97.05	1.02
Single-point impulse	98.71	90.22	94.27	2.85
Unconstrained Gaussian	98.93	93.76	96.28	2.44
<i>PRPS scale sensitivity</i>				
PRPS ($\sigma=2$, $r=6$)	99.82	94.44	97.05	1.02
PRPS ($\sigma=1$, $r=3$)	97.60	88.60	92.87	5.26
PRPS ($\sigma=3$, $r=9$)	99.70	92.50	95.99	0.67

Table 4: Ablation study of HRPE on SIRST-UAVB. The first block evaluates resolution preservation with different output strides s , while the second block analyzes the contribution of HRPE components. We report Pre(%), Rec (Pd)(%), F1(%), Fa (10^{-8}), Params (10^{-2} M), and FLOPs (G). Params and FLOPs are computed at 640×640 input resolution.

Variant	Pre $\uparrow$	Rec $\uparrow$	F_1 $\uparrow$	Fa $\downarrow$	FLOPs $\downarrow$	Params $\downarrow$
<i>Resolution preservation (output stride)</i>						
HRPE ($s=4$)	99.82	94.44	97.05	1.02	7.68	29.47
HRPE ($s=2$)	99.63	91.23	95.25	0.81	26.77	26.56
HRPE ($s=8$)	90.05	83.98	86.91	22.38	3.05	33.27
<i>Component ablation</i>						
HRPE	99.82	94.44	97.05	1.02	7.68	29.47
w/o channel reorganization	97.70	93.90	95.76	5.33	7.70	29.45
w/o reweighting	96.52	93.71	95.09	8.15	7.72	29.63

A small σ provides insufficient supervision density, while a large σ over-smooths probability peaks and weakens centroid separability. The moderate scale best matches the intrinsic spatial extent of infrared small targets and yields balanced probabilistic regression. This trend further clarifies that PRPS is not intended to reconstruct target contours; instead, it should provide a compact, centroid-oriented local probabilistic response for localization.

Impact of resolution preservation. We vary the output stride to evaluate high-resolution representation in HRPE. As shown in Tab. 4, stride 8 reduces Recall to 83.98 and increases Fa to 22.38, severely degrading F1. Stride 2 slightly lowers Fa to 0.81 but does not improve F1 and increases FLOPs from 7.68 to 26.77. Excessive downsampling weakens fragile small-target representations under low-SNR conditions, undermining structured probabilistic response model-




Fig. 5: Ablation study of PRPS on SIRST-UAVB. Red rectangles denote ground-truth targets. Red circles indicate true positives and orange circles correspond to false negatives under the unified centroid protocol.

ing. Overly fine resolution increases structural redundancy without clear discriminative gain. Stride 4 provides the best balance within the encoder-only framework.

Impact of HRPE components. Removing channel reorganization increases Fa to 5.33 and reduces F1 to 95.76, while removing reweighting further increases Fa to 8.15 with slight Recall degradation, as shown in Tab. 4. These results indicate that efficient channel interaction and adaptive refinement are essential for structured probabilistic response learning in clutter. Removing either component primarily weakens false-alarm control.

5 Conclusion

We revisit IRSTD from a modeling perspective and argue that dense pixel-wise segmentation is not inherently required for reliable centroid localization. By reformulating IRSTD as spatial probability regression, we propose SPIRE, which integrates PRPS for structured probabilistic supervision guided by point-response priors under extreme class imbalance and a lightweight HRPE for encoder-only probabilistic response map regression with lightweight peak-based inference rather than mask reconstruction. Experiments on SIRST-UAVB and SIRST4 demonstrate that SPIRE achieves competitive target-level detection performance with consistently low Fa and significantly reduced FLOPs and Params. These results validate that sparse centroid modeling with physically grounded probabilistic supervision provides an efficient and principled alternative to dense mask reconstruction paradigms for IRSTD.

Acknowledgements

This work was partially supported in part by the National Natural Science Foundation of China (NSFC) under No. 62101567, No.62401589, No. 62501618 and No. 4250013163.

References

1. An, X., Zhao, L., Gong, C., Wang, N., Wang, D., Yang, J.: Sharpnose: Sparse high-resolution representation for human pose estimation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 691–699 (2024)
2. Arce, G.R., McLoughlin, M.P.: Theoretical analysis of the max/median filter. *IEEE transactions on acoustics, speech, and signal processing* **35**(1), 60–69 (1987)
3. Atrash, A., Ertekin, S., Ugur, O., Moured, O., Chen, Y., Zhang, J.: Ty-ris: Tactical yolo tricks for real-time infrared small target detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2201–2210 (2025)
4. Bai, X., Bi, Y.: Derivative entropy-based contrast measure for infrared small-target detection. *IEEE Transactions on Geoscience and Remote Sensing* **56**(4), 2452–2466 (2018)
5. Bai, X., Zhou, F.: Analysis of new top-hat transformation and the application for infrared dim small target detection. *Pattern Recognition* **43**(6), 2145–2156 (2010)
6. Chen, C.P., Li, H., Wei, Y., Xia, T., Tang, Y.Y.: A local contrast method for small infrared target detection. *IEEE transactions on geoscience and remote sensing* **52**(1), 574–581 (2013)
7. Chen, Z., Luo, S., Xie, T., Liu, J., Wang, G., Lei, G.: A novel infrared small target detection method based on bemd and local inverse entropy. *Infrared Physics & Technology* **66**, 114–124 (2014)
8. Cheng, B., Xiao, B., Wang, J., Shi, H., Huang, T.S., Zhang, L.: Higherhrnet: Scale-aware representation learning for bottom-up human pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5385–5394 (2020). <https://doi.org/10.1109/cvpr42600.2020.00543>
9. Dai, Y., Li, X., Zhou, F., Qian, Y., Chen, Y., Yang, J.: One-stage cascade refinement networks for infrared small target detection. *IEEE transactions on geoscience and remote sensing* **61**, 1–17 (2023)
10. Dai, Y., Wu, Y., Song, Y., Guo, J.: Non-negative infrared patch-image model: Robust target-background separation via partial sum minimization of singular values. *Infrared Physics & Technology* **81**, 182–194 (2017)
11. Dai, Y., Wu, Y., Zhou, F., Barnard, K.: Asymmetric contextual modulation for infrared small target detection. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 950–959 (2021)
12. Dai, Y., Wu, Y., Zhou, F., Barnard, K.: Attentional local contrast networks for infrared small target detection. *IEEE transactions on geoscience and remote sensing* **59**(11), 9813–9824 (2021)
13. Deng, H., Sun, X., Liu, M., Ye, C., Zhou, X.: Small infrared target detection based on weighted local difference measure. *IEEE Transactions on Geoscience and Remote Sensing* **54**(7), 4204–4214 (2016)
14. Gao, C., Meng, D., Yang, Y., Wang, Y., Zhou, X., Hauptmann, A.G.: Infrared patch-image model for small target detection in a single image. *IEEE transactions on image processing* **22**(12), 4996–5009 (2013)
15. He, W., Liu, M., Yu, Y.: Hybrid mask generation for infrared small target detection with single-point supervision. *Neurocomputing* **658**, 131688 (2025). <https://doi.org/10.1016/j.neucom.2025.131688>
16. Hou, Q., Zhang, L., Tan, F., Xi, Y., Zheng, H., Li, N.: Istdu-net: Infrared small-target detection u-net. *IEEE Geoscience and Remote Sensing Letters* **19**, 1–5 (2022). <https://doi.org/10.1109/LGRS.2022.3141584>

17. Li, B., Xiao, C., Wang, L., Wang, Y., Lin, Z., Li, M., An, W., Guo, Y.: Dense nested attention network for infrared small target detection. *IEEE Trans Image Process* **32**, 1745–1758 (2023). <https://doi.org/10.1109/TIP.2022.3199107>
18. Li, B., Wang, Y., Wang, L., Zhang, F., Liu, T., Lin, Z., An, W., Guo, Y.: Monte carlo linear clustering with single-point supervision is enough for infrared small target detection. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 1009–1019 (2023)
19. Li, N., Wei, D.: Redetr-ristd: Real-time long-range infrared small target detection network based on the reparameterized efficient detection transformer. *Sensors* **25**(9) (2025). <https://doi.org/10.3390/s25092771>
20. Li, Y., Zhang, Y.: Robust infrared small target detection using local steering kernel reconstruction. *Pattern Recognition* **77**, 113–125 (2018)
21. Liu, Q., Liu, R., Zheng, B., Wang, H., Fu, Y.: Infrared small target detection with scale and location sensitivity. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 17490–17499 (2024)
22. Ni, R., Wu, J., Qiu, Z., Chen, L., Luo, C., Huang, F., Liu, Q., Wang, B., Li, Y., Li, Y.: Point-to-point regression: Accurate infrared small target detection with single-point annotation. *IEEE Transactions on Geoscience and Remote Sensing* **63**, 1–19 (2025). <https://doi.org/10.1109/tgrs.2025.3554025>
23. Song, I., Lee, J., Ryu, M., Lee, J.: Motion-aware heatmap regression for human pose estimation in videos. In: *PROCEEDINGS OF THE THIRTY-THIRD INTERNATIONAL JOINT CONFERENCE ON ARTIFICIAL INTELLIGENCE, IJCAI 2024*. pp. 1245–1253. *IJCAI-INT JOINT CONF ARTIF INTELL* (2024)
24. Sun, H., Bai, J., Yang, F., Bai, X.: Receptive-field and direction induced attention network for infrared dim small target detection with a large-scale dataset irdst. *IEEE Transactions on Geoscience and Remote Sensing* **61**, 1–13 (2023). <https://doi.org/10.1109/TGRS.2023.3235150>
25. Tong, X., Zuo, Z., Su, S., Wei, J., Sun, X., Wu, P., Zhao, Z.: St-trans: Spatial-temporal transformer for infrared small target detection in sequential images. *IEEE Transactions on Geoscience and Remote Sensing* **62**, 1–19 (2024). <https://doi.org/10.1109/tgrs.2024.3355947>
26. Wang, J., Sun, K., Cheng, T., Jiang, B., Deng, C., Zhao, Y., Liu, D., Mu, Y., Tan, M., Wang, X., Liu, W., Xiao, B.: Deep high-resolution representation learning for visual recognition. *IEEE Trans Pattern Anal Mach Intell* **43**(10), 3349–3364 (2021). <https://doi.org/10.1109/TPAMI.2020.2983686>
27. Wang, W., Xiao, C., Dou, H., Liang, R., Yuan, H., Zhao, G., Chen, Z., Huang, Y.: Ispanet: A pyramid self-attention network for single-frame high-resolution infrared small target detection with a large-scale dataset shr-irst. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **17**, 11146–11162 (2024). <https://doi.org/10.1109/jstars.2024.3381779>
28. Wu, F., Liu, A., Zhang, T., Zhang, L., Luo, J., Peng, Z.: Saliency at the helm: Steering infrared small target detection with learnable kernels. *IEEE Transactions on Geoscience and Remote Sensing* **63**, 1–14 (2025). <https://doi.org/10.1109/tgrs.2024.3521947>
29. Wu, J., Luo, C., Qiu, Z., Chen, L., Ni, R., Li, Y., Huang, F., Wu, J.: Dfinet: Dynamic feedback iterative network for infrared small target detection. *Pattern Recognition* **169** (2026). <https://doi.org/10.1016/j.patcog.2025.111958>
30. Wu, X., Hong, D., Chanussot, J.: Uiu-net: U-net in u-net for infrared small object detection. *IEEE Trans Image Process* **32**, 364–376 (2023). <https://doi.org/10.1109/TIP.2022.3228497>

31. Xu, H., Zhong, S., Zhang, T., Zou, X.: Multiscale multilevel residual feature fusion for real-time infrared small target detection. *IEEE Transactions on Geoscience and Remote Sensing* **61**, 1–16 (2023)
32. Xu, M., Yu, C., Li, Z., Tang, H., Hu, Y., Nie, L.: Hdnet: A hybrid domain network with multiscale high-frequency information enhancement for infrared small-target detection. *IEEE Transactions on Geoscience and Remote Sensing* **63**, 1–15 (2025). <https://doi.org/10.1109/tgrs.2025.3574962>
33. Yang, B., Zhang, X., Zhang, J., Luo, J., Zhou, M., Pi, Y.: Efnnet: Enhancing feature learning network for infrared small target detection. *IEEE Transactions on Geoscience and Remote Sensing* **62**, 1–11 (2024)
34. Yang, J., Liu, S., Wu, J., Su, X., Hai, N., Huang, X.: Pinwheel-shaped convolution and scale-based dynamic loss for infrared small target detection. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 39, pp. 9202–9210 (2025)
35. Ying, X., Liu, L., Wang, Y., Li, R., Chen, N., Lin, Z., Sheng, W., Zhou, S.: Mapping degeneration meets label evolution: Learning infrared small target detection with single point supervision. In: *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 15528–15538 (2023). <https://doi.org/10.1109/cvpr52729.2023.01490>
36. Yu, C., Zhao, J., Liu, Y., Zhao, S., Dai, Y., Yue, X.: From easy to hard: Progressive active learning framework for infrared small target detection with single point supervision. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 2588–2598 (2025)
37. Yuan, S., Qin, H., Yan, X., Akhtar, N., Mian, A.: Sctransnet: Spatial-channel cross transformer network for infrared small target detection. *IEEE Transactions on Geoscience and Remote Sensing* **62**, 1–15 (2024). <https://doi.org/10.1109/tgrs.2024.3383649>
38. Yue, T., Lu, X., Cai, J., Chen, Y., Chu, S.: Sds-net: Shallow–deep synergism-detection network for infrared small target detection. *IEEE Transactions on Geoscience and Remote Sensing* **63**, 1–13 (2025). <https://doi.org/10.1109/tgrs.2025.3588117>
39. Zhang, M., Yue, K., Zhang, J., Li, Y., Gao, X.: Exploring feature compensation and cross-level correlation for infrared small target detection. In: *Proceedings of the 30th ACM International Conference on Multimedia*. pp. 1857–1865 (2022)
40. Zhang, M., Zhang, R., Yang, Y., Bai, H., Zhang, J., Guo, J.: Isnet: Shape matters for infrared small target detection. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 877–886 (2022)
41. Zhang, X., Zhou, X., Lin, M., Sun, J.: Shufflenet: An extremely efficient convolutional neural network for mobile devices. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 6848–6856 (2018)
42. Zhao, E., Dong, L., Li, C., Ji, Y.: Infrared maritime target detection based on temporal weight and total variation regularization under strong wave interferences. *IEEE Transactions on Geoscience and Remote Sensing* **62**, 1–19 (2024)