

Monte Carlo Linear Clustering with Single-Point Supervision is Enough for Infrared Small Target Detection

Boyang Li¹, Yingqian Wang¹, Longguang Wang², Fei Zhang³,
Ting Liu¹, Zaiping Lin¹✉, Wei An¹, Yulan Guo¹

¹National University of Defense Technology, ²Aviation University of Air Force,

³Shanghai Jiao Tong University

{liboyang20, linzaiping}@nudt.edu.cn

Abstract

Single-frame infrared small target (SIRST) detection aims at separating small targets from clutter backgrounds on infrared images. Recently, deep learning based methods have achieved promising performance on SIRST detection, but at the cost of a large amount of training data with expensive pixel-level annotations. To reduce the annotation burden, we propose the first method to achieve SIRST detection with single-point supervision. The core idea of this work is to recover the per-pixel mask of each target from the given single point label by using clustering approaches, which looks simple but is indeed challenging since targets are always insalient and accompanied with background clutters. To handle this issue, we introduce randomness to the clustering process by adding noise to the input images, and then obtain much more reliable pseudo masks by averaging the clustered results. Thanks to this “Monte Carlo” clustering approach, our method can accurately recover pseudo masks and thus turn arbitrary fully supervised SIRST detection networks into weakly supervised ones with only single point annotation. Experiments on four datasets demonstrate that our method can be applied to existing SIRST detection networks to achieve comparable performance with their fully-supervised counterparts, which reveals that single-point supervision is strong enough for SIRST detection. Our code will be available at: <https://github.com/YeRen123455/SIRST-Single-Point-Supervision>.

1. Introduction

Single-frame infrared small target (SIRST) detection has been widely used in many applications such as marine resource utilization [23, 45], high-precision navigation [34, 13], and ecological environment monitoring [35].

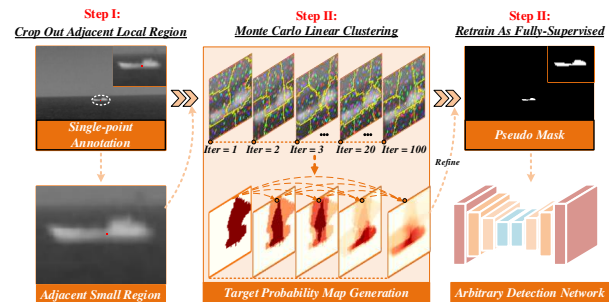


Figure 1. Pipeline of our method. The proposed Monte Carlo linear clustering along with single-point annotation can produce high-quality pseudo mask (i.e., target probability map (TPM)).

However, existing SIRST detection methods mainly rely on segmentation pipeline with pixel-level supervision. Pixel-level annotating is time-consuming and thus hinders the quick deployment of those large scale data-dependent scenarios. Therefore, it is in urgent need to alleviate the annotation burden while maintaining state-of-the-art performance in SIRST detection.

Point-level supervision, as an annotation-friendly form of supervision, can provide both category and localization information for each target. Previous works in general object segmentation utilized rich semantic information to design various training loss [4] and network modules [29, 27], and continuously expand the single-point annotation to a pixel-level pseudo mask. Although having achieved promising performance, these methods heavily rely on prior semantic information (e.g., objectness prior in [4, 50], shape prior in [29, 27, 49]), and generally need more than one point as supervision. However, infrared small targets generally occupy no more than 0.15% area of the whole image [47] and lack texture, shape and color information. Consequently, previous semantic information-dependent methods cannot well adopt to this task.

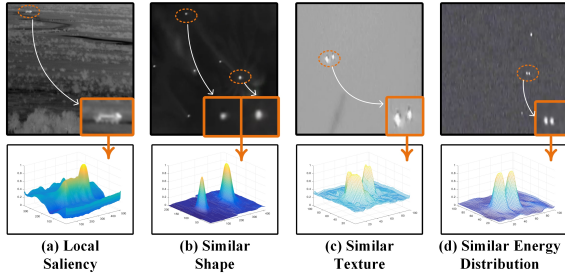


Figure 2. Some inherent characteristics of SIRST observed from the small local regions (i.e., salient in local region, similar shape, texture, and energy distribution).

We notice that although infrared small targets lack semantic information from the perspective of whole image, they are quite salient in the local small region. Specifically, as shown in Fig. 2, small targets in infrared images generally have high energy concentration and exhibit significant gradient difference in their local regions. Moreover, limited semantic information (e.g., weak texture, shape, and color distribution) makes most point-like and spot-like targets exhibit similar energy distribution (e.g., Gaussian distribution). Consequently, it is straightforward to use a simple linear clustering approach (LCA) to separate the salient regions by measuring their color and spatial distances with adjacent backgrounds.

However, LCA easily falls into local optimum and produces inaccurate segmentation results due to their fixed hyper-parameters. To handle this problem, we propose a Monte Carlo linear clustering (MCLC) method to regularize the clustering process and recover the reliable clustering result by repetitive random experiments. Specifically, we introduce randomness to the clustering process by adding noise to the input images. With the help of random noise, the distance between unexpected background region and target region can be significantly enlarged. The misclustered target regions can be pushed away from the false clustering center and thus return back the true clustering center. In this way, the single-point annotation can be gradually recovered as a much more reliable pixel-level target probability map (TPM) by averaging the clustered results. Pixels with larger values in TPM represent a higher probability of belonging to the target. Finally, we use the refined TPM to turn arbitrary fully supervised SIRST detection networks into weakly supervised ones with only single point annotation. Fig. 1 shows the overall pipeline of our method, and the main contributions are summarized as follows:

- To the best of our knowledge, this is the first single-point supervised method to achieve SIRST detection. A Monte Carlo linear clustering-based pipeline is proposed to achieve comparable performance with the fully supervised counterpart by using single-point anno-

taion.

- Inspired by the inherent characteristic of SIRST, a simple yet effective linear clustering approach with random noise-guided Monte Carlo regularization is proposed to coarsely extract and further refine the candidate target region.
- Experiments on four public SIRST datasets demonstrate the effectiveness of our method. Ablation study reveals that pixel-level labels are not necessary for SIRST detection while single-point supervision is strong enough.

2. Related Work

2.1. Infrared Small Target Detection

SIRST detection has been extensively investigated for decades. Early traditional paradigm achieves SIRST detection by measuring the discontinuity between targets and backgrounds. Typical methods include filtering-based methods [31, 14], local contrast measure-based methods [6, 18, 19, 20, 22, 36], and low rank-based methods [16, 43, 9, 44, 51, 10]. Since real scenes are much more complex with dramatic changes in target size, shape, and clutter background, it is difficult to use handcrafted features and fixed hyper-parameters to handle such variations. To address this problem, recent deep learning-based methods learn trainable features in a data-driven manner and thus achieve better performance than traditional ones.

Existing deep learning-based methods can be divided into detection based methods and segmentation based methods. Since the pixel-level classification result is essential for the subsequent recognition task in SIRST, segmentation-based methods have attracted increasing attention recently. Dai et al. [11] proposed the first segmentation-based network (i.e., ACM). They designed an asymmetric contextual module to aggregate features from shallow layers and deep layers. Then, Dai et al. [12] improved ACM by introducing a dilated local contrast measure. Specifically, a feature cyclic shift scheme was designed to achieve a trainable local contrast measure. After that, Zhang et al. [45] modeled SIRST detection as a shape detection task. A Taylor finite difference (TFD)-inspired edge block and a two-orientation attention aggregation (TOAA) block were proposed to capture precise shape of infrared targets. Recently, Li et al. [23] proposed a dense nested attention network (DNANet). A specifically-designed dense nested interactive module (DNIM) was proposed to both extract high-level information and maintain the response of small targets in deep layers.

Although the performance has been continuously improved by recent networks, existing deep learning-based methods rely on the fully-supervised training scheme with

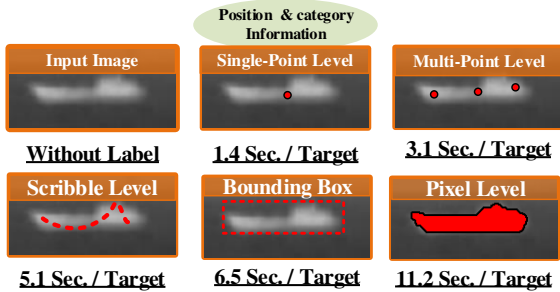


Figure 3. Sample of input images, four common weakly supervised annotations (i.e., single-point, multi-point, scribble, bounding box), and one fully supervised pixel-level annotation.

pixel-level annotations. Expensive labor cost makes existing methods hard to be deployed to large scale data-dependent tasks.

2.2. Point-based Segmentation

Point-level supervision has been widely used in many high-level tasks (e.g., object detection [7, 46], object localization [41], and semantic segmentation [4, 29, 27, 24, 49]) and yield promising performance. For semantic segmentation task, Bearman et al. [4] proposed the first single-point supervised semantic segmentation method. They incorporated the single-point supervision along with an objectness prior as the loss function to infer the extent of the object. Then, Papadopoulos et al. [29] and Maninis et al. [27] utilized four extreme points (i.e., left-most, right-most, top, and bottom pixels) as supervision to further improve the quality of the pseudo mask. After that, Austin et al. [28] followed the class activation maps (CAM) [50] pipeline and used the point annotation to refine the quality of pseudo labels. More recently, Li et al. [24] utilized the semantical consistency property of general objects with 20 randomly annotated points to achieve comparable segmentation results with fully-supervised counterpart. Moreover, in the field of cell segmentation, Zhao et al. [49] proposed the first single-point supervised segmentation method, in which two semantic prior-based training losses (i.e., divergence loss and consistency loss) were proposed to achieve high-quality cell segmentation.

Although promising performance has been achieved, existing works rely on rich semantic information (e.g., objectness prior in [4, 50], shape prior in [29, 27, 49]) and most of them need multi-point annotation. Infrared small targets generally occupy no more than 0.15% area of the whole image [47] and lack texture, shape and color information. Previous semantic information-dependent methods cannot be directly used for SIRST detection.

3. Methodology

In this section, we first introduce the motivation. Then, a detailed illustration of the proposed linear clustering ap-

proach with Monte Carlo regularization is provided. Overall architecture of proposed method is shown in Fig. 4.

3.1. Motivation

We investigate the annotation cost of four common weakly-supervised annotations (i.e., single point [5], multiple points[5, 29, 27], scribbles [25], bounding boxes[21]) by re-labelling the NUDT-SIRST [23] and NAAA-SIRST [11] datasets. Fig. 3 shows the average annotation time of four weakly-supervised and one fully-supervised approaches. Note that, single-point supervision can reduce about 87% annotation time as compared to the pixel-level annotation approach, which motivates us to apply single-point annotation to SIRST detection. Please refer to the supplementary material for more detailed analysis.

However, existing point-level supervised methods [5, 29, 27] are designed for general objects with rich semantic information (e.g., shape, texture, and color). As shown in Fig. 2, the characteristics of infrared small targets make the aforementioned semantic-based pipeline unsuitable for SIRST detection. Considering that infrared small targets are quite salient in their local regions and share similar shape, texture, and energy distribution. Intuitively, we can use simple linear clustering approaches (e.g., Super-pixel [1], K-Means [26]) to separate the salient regions by comparing their color and spatial distances among these clusters.

3.2. Linear Clustering Approach

Without loss of generality, we take the quick linear clustering method SuperPixel [1] as an example to introduce how to use LCA to coarsely separate the target region. It should be noted that other linear clustering approaches such as KMeans [26] can be also used in our method, as demonstrated in Fig. 5. The pipeline of our LCA is shown in Fig. 4 (*Step I ~ Step III*), in which the clustering center of the image is first initialized, and then adjusted according to a color and spatial distance based measurement. Finally, the clustering center is continuously updated until the pre-defined convergence requirements are met.

1) Clustering Center Initialization: We initialize the clustering center based on the pre-defined number of clusters N . Given an input image $I \in \mathbb{R}^{H \times W}$, we first divide it into N regions with equal areas and fixed grid interval $S = \sqrt{\frac{H \times W}{N}}$ (as shown in Fig. 4-*Step I*). After that, the clustering center is first initialized as the centroid of each region, and then moved to the lowest gradient position of its 3×3 neighborhood. The clustering center of the n^{th} clustering region C_n can be represented as a vector as:

$$C_n = [c_n, s_n]^T, \quad (1)$$

where $[c_n]^T$ represents the CIELAB color [32] of C_n , $[s_n]^T$ denotes the spatial position in image coordinate of C_n .

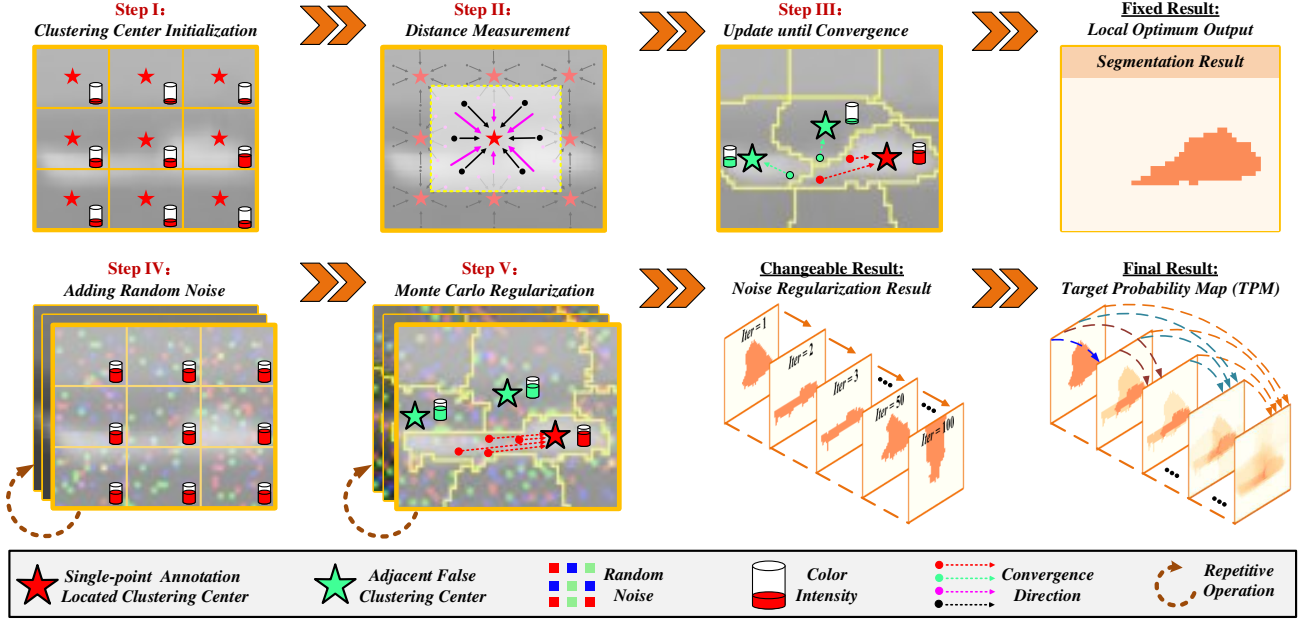


Figure 4. Illustration of our method. *Step I* ~ *Step III*: LCA can coarsely separate target from clutter background, but easily falls into local optimum and produce incomplete pseudo mask. *Step IV* ~ *Step V*: Monte Carlo regularization can increase the color distance between background and target region and thus enforce the misclustered target region away from the false clustering center. After repetitive random experiments, the accumulated clustering results are finally shown as a target probability map (TPM), where pixels with larger values represent a higher probability of belonging to the target.

2) **Distance Measurement:** After clustering center initialization, we adopt a color and spatial distance based measurement to determine the clustering center of each pixel. As shown in Fig. 4-*Step II*, for each clustering center C_n , we use L_2 distance to measure its distance with each pixel $p_i = [c_i, s_i]^T$ in its adjacent $2S \times 2S$ regions. Considering that the ranges of color value and spatial position value are inconsistent, we normalize the both distance values with normalization coefficient μ_s and μ_c . $\mathcal{D}(C_n, p_i)$ is written as:

$$\begin{aligned} \mathcal{D}(C_n, p_i) &= \sqrt{(\mathcal{D}_c)^2 + (\mathcal{D}_s)^2} \\ &= \sqrt{\frac{(c_{p_i} - c_{C_n})^2}{\mu_c^2} + \frac{(s_{p_i} - s_{C_n})^2}{\mu_s^2}}. \end{aligned} \quad (2)$$

Then, the clustering label l for each pixel p_i can be determined as:

$$l(p_i) = \operatorname{argmin}_{n \in \{1, 2, 3, \dots, N\}} [\mathcal{D}(C_n, p_i)]. \quad (3)$$

3) **Updating and Convergence:** After determining the clustering center of each pixel at the current iteration, we repetitively update the position of each clustering center until convergence, as shown in Fig. 4-*Step III*. Specifically, we first calculate the clustering centers to be the mean vector $[c_m, s_m]^T$ of all the pixels $P = \{p_1, p_2, \dots, p_m\}$ with the same clustering label $l(C'_n)$. The new clustering center C'_n can be formulated as:

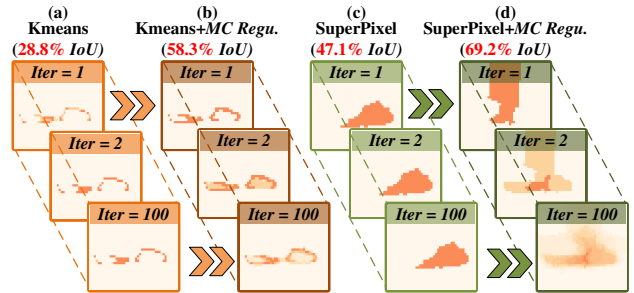


Figure 5. Visualization results achieved by (a) Kmeans, (b) Kmeans with Monte Carlo Regularization, (c) Super pixel, (d) Super pixel with Monte Carlo Regularization. Our proposed Monte Carlo regularization helps to expand informative target region.

$$C'_n(c_n, x_n, y_n) = \frac{1}{|P|} \sum_{p_m \in P} (c_{p_m}, s_{p_m}). \quad (4)$$

After that, we adopt L_2 norm to compute a distance \mathcal{E} between new clustering center locations C'_n and previous clustering center locations C_n . The updating steps can be repeated iteratively until the error is smaller than the predefined threshold T . That is $\|C'_n - C_n\|_2 < T$, where $\|\cdot\|_2$ denote the L_2 norm. Finally, we assign the single-point annotation p_{Anno} located at the n^{th} clustering region as the clustering result M_{pred} .

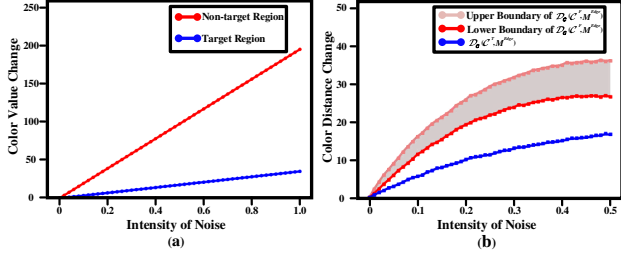


Figure 6. Quantitative analysis on (a) Sensitivity of target and non-target regions to additional noise, and (b) Color distance of target edge region with true foreground region centroid $\Delta(\mathcal{D}_c(\mathcal{C}^T, M_{pred}))$ and false background region centroid $\Delta(\mathcal{D}_c(\mathcal{C}^F, M_{pred}))$. Average results from 10 trails are reported. Readers can refer to the supplementary material for more detailed illustration.

3.3. Monte Carlo Regularization

Although LCA can coarsely separate local salient regions from the backgrounds, as shown in Fig. 5 (a) and (c), a linear clustering method easily falls into local optimum and generates incomplete or oversized results. In this subsection, we first analyze the reason of this issue, and then propose a Monte Carlo regularization to alleviate this issue, as shown in Fig. 4 (StepIV ~ StepV).

1) Inaccurate Result within a Single Clustering: Due to the varied size of small targets, LCA with fixed number of clustering centers cannot produce ideal results for all kinds of targets. That is because, both color \mathcal{D}_c and spatial distance \mathcal{D}_s determine the output of linear clustering model in Equation 2. Since edge areas $M_{pred}^{Edge} = \{m_{pred}^{Edge(1)}, m_{pred}^{Edge(2)}, \dots, m_{pred}^{Edge(Z)}\}$ of the spot and extended target (i.e., big target) are generally far from the true clustering center \mathcal{C}^T (shown as the left-most green point in Fig. 4-Step III). The spatial distance $\mathcal{D}_s(\mathcal{C}^T, M_{pred}^{Edge})$ between the true clustering center \mathcal{C}^T and edge areas M_{pred}^{Edge} is usually fater than that of the false clustering center \mathcal{C}^F . That is:

$$\mathcal{D}_s(\mathcal{C}^T, m_{pred}^{Edge(z)}) > \mathcal{D}_s(\mathcal{C}^F, m_{pred}^{Edge(z)}). \quad (5)$$

Moreover, the color value of big target generally exhibits Gaussian distribution. Those edge areas often exhibit relatively high color difference with the true clustering center, and are close to the adjacent false clustering center (shown as the right-most green point in Fig. 4 Step III). That is:

$$\mathcal{D}_c(\mathcal{C}^T, m_{pred}^{Edge(z)}) > \mathcal{D}_c(\mathcal{C}^F, m_{pred}^{Edge(z)}). \quad (6)$$

As a result, the label of edge areas l_{pred}^{Edge} may be falsely included into the adjacent false clustering center \mathcal{C}^F , resulting in incomplete clustering result.

2) Random Noise Guided Regularization: Spatial distance \mathcal{D}_s is generally fixed and hard to change in the im-

age. As shown in Fig. 6 (a), when adding random noise, target and non-target regions exhibit different sensitivity to the noise. Target regions with higher color value are more robust to the additional noise than the non-target regions. Based on this finding, we are motivated to introduce random noise to enlarge the color distance $\mathcal{D}_c(\mathcal{C}^F, M_{pred}^{Edge})$ between edge areas and false clustering center (as shown in Fig. 4-Step III), and thus make the edge areas be clustered to the true clustering center. Random noise-regularized clustering result can be formulated as follows:

$$M_{pred}^{Noise} = LCA(\text{clip}(I + \mathcal{N})), \quad (7)$$

where $\text{clip}(\cdot)$ operation represents that color values larger than 255 will be fixed at 255. \mathcal{N} is the additional random noise. The color value change caused by random noise can be formulated as:

$$\Delta(\mathcal{D}_c(\mathcal{C}, M_{pred})) = \mathcal{D}_c^{Noise} - \mathcal{D}_c^{Clean}. \quad (8)$$

Since background region has relatively low color value than foreground region. Additional noise will greatly increase the color value of adjacent false clustering centers, and thus enlarges their color distance with the edge area of target region. On the contrary, when comparing Fig. 4-Step IV with Fig. 4-Step V, we observe that foreground regions have high energy concentricity and are usually close to saturation. Additional noise cannot bring so much color change to target regions as the background regions. Benefited by the additional noise, as shown in Fig. 6 (b), the change of color distance between edge areas and false clustering center (i.e., $\Delta(\mathcal{D}_c(\mathcal{C}^F, M_{pred}))$) will be significantly larger than that between edge areas and true clustering center (i.e., $\Delta(\mathcal{D}_c(\mathcal{C}^T, M_{pred}))$), making the misclustered target region be more closed to the true clustering center.

3) Monte Carlo Process: As aforementioned, random noise can enforce the misclustered region to be gradually close to the true clustering center. Considering the randomness of noise, additional noise does not always bring such positive effect. We are motivated to adopt the Monte Carlo method to accumulate the clustering result from repetitive random experiments and gradually recover the reliable clustering result. To this end, we perform K independent clustering and formulate this process as follows:

$$M_{Pred}^{Full} = \frac{1}{K} \sum_{k=1}^K \left(LCA(\text{clip}(I + \mathcal{N}^{(k)})) \right), \quad (9)$$

where M_{Pred}^{Full} is the accumulated clustering results after K independent clustering by LCA, and finally shown as a target probability map (TPM), where pixels with larger values represent a higher probability of belonging to the target.

The optimization process finally stops when the distance between the new clustering center locations \mathcal{C}' and previous

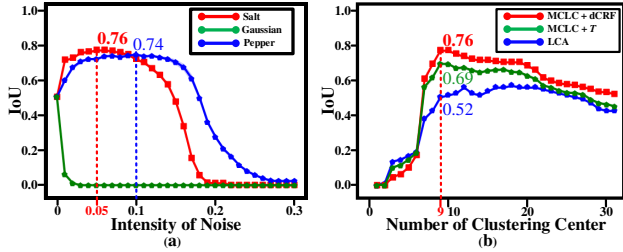


Figure 7. IoU scores achieved by our method with (a) different types and intensity of additional noise and (b) different number of clustering center on the NUAAs-SIRST dataset. Average results from 5 trails are reported. Readers can refer to the supplementary material for additional experiments on the other datasets.

Table 1. $IoU(10^{-2})$ value achieved by the different variants of our method on the training set of four representative datasets (i.e., NUAAs, IRSTD, NUDT, NUDT-sea). MC Regu. refers to Monte Carlo regularization.

| Point Label | LCA | MC Regu. | dCRF | $IoU(\%)$ |
|-------------|-----|----------|------|----------------------------------|
| ✓ | | | ✓ | 2.53 / 1.62 / 2.28 / 1.82 |
| ✓ | ✓ | | | 51.8 / 55.5 / 51.2 / 38.1 |
| ✓ | ✓ | ✓ | | 70.4 / 66.2 / 52.5 / 43.9 |
| ✓ | ✓ | ✓ | ✓ | 76.3 / 68.8 / 60.2 / 47.1 |

clustering center locations \mathcal{C} is smaller than the pre-defined threshold T^F , i.e., $\|\mathcal{C}'_n - \mathcal{C}_n\|_2 < T^F$.

4. Experiments

In this section, we first introduce our evaluation metrics and implementation details. Then, we compare our MCLC to several state-of-the-art unsupervised, point-level supervised, and fully-supervised SIRST detection methods. Finally, we present ablation studies to investigate the effectiveness of our method. Note that, we search for optimal parameters of MCLC in the training set of NUAAs-SIRST [11] and directly adopt them as default parameters to generate pseudo masks on the other three datasets[45, 23, 39].

4.1. Evaluation Metrics

Considering that both the quality of pseudo labels and the final detection results are crucial to SIRST detection, we followed previous works [37, 42, 3, 2] and propose a two-stage pipeline (i.e., pseudo mask and final detection result evaluation) to comprehensively evaluate the effectiveness of the proposed method. Firstly, we followed the same metrics in weakly-supervised general object segmentation [3, 2] and adopted intersection of union (IoU) to evaluate the quality of generated pseudo labels on the training set. For the final detection results evaluation, we followed previous works [11, 12, 23] and adopted probability of detection (P_d), and false-alarm rate F_a to evaluate the localization precision, and used IoU to evaluate the shape description ability on the test set.

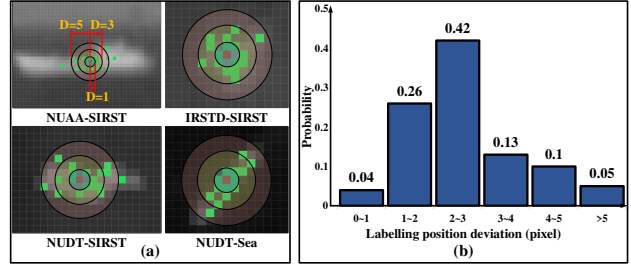


Figure 8. (a) Labelling deviation examples of small target from four datasets. (b) Label deviation distribution map generated from 105 manually re-labelled targets. Real world labelling deviation of small target is generally less than 3 pixels.

4.2. Implementation Details

1) **Datasets:** We evaluated our method on the NUAAs-SIRST [11], IRSTD-1k [45], NUDT-SIRST [23], and NUDT-SIRST-sea [39] datasets. Without specification, for the NUAAs-SIRST and NUDT-SIRST datasets, we used the same division as in [23] and set the ratio of train-val set to test set as 1. For the IRSTD-1k dataset, we followed [45] and used 80% data for training and validation, and the remaining 20% data for test. The division setting in [39] was adopted for the NUDT-SIRST-sea dataset, i.e., 41 images were used for training and validation, and the remaining 7 images were used for test. Note that, only single point-level annotation with category information (i.e., point, spot, and extend categories) are available during network training.

2) **Training Details:** In our experiments, all input images were first normalized before training. Then, these normalized images were sequentially processed by random flip, and random crop for data augmentation. Next, these images were resized to a fixed resolution. Specifically, 1024×1024 and 256×256 spatial resolutions were used for the NUDT-SIRST-sea and the remaining three datasets, respectively. Finally, we fed these resized datasets into network for both training and evaluation. Without specification, we adopted the centroid of ground truth mask as the single-point label. The influence of label position deviation is discussed in the following ablation. Moreover, all networks were trained with the Soft-IoU loss function and optimized by the Adam method [15] with the CosineAnnealingLR scheduler. The Xavier method [17] was used for all weights and bias parameters initialization. We set the learning rate and epoch number to 0.05 and 1500, respectively. The batch size was set to 16 for the NUAAs-SIRST, IRSTD-1k, NUDT-SIRST datasets but set to 4 for the NUDT-SIRST-sea dataset due to its large image size. All models were implemented in PyTorch [30] on a computer with an AMD Ryzen 9 3950X @ 2.20 GHz CPU and an Nvidia GeForce 3090 GPU.

4.3. Ablation Study

1) **Effectiveness of Proposed Components:** Table 1 shows the contribution of our proposed components on four

Table 2. $IoU(10^{-2})$ values and corresponding labelling cost with stronger supervision on the train set of four datasets.

| Annotation | Dataset | | | |
|--------------|---------|-------|------|----------|
| | NAAA | IRSTD | NUDT | NUDT-sea |
| Single-Point | 76.3 | 68.8 | 60.2 | 47.1 |
| Multi-Point | 76.6 | 69.1 | 62.6 | 53.7 |
| Scribble | 78.5 | 70.5 | 63.2 | 55.2 |
| Bounding-Box | 78.9 | 71.8 | 64.8 | 54.5 |

datasets. Compared to the initial single-point label, linear clustering approach (LCA) introduces 49.3%, 53.9%, 48.9%, 36.3% improvements for the NAAA, IRSTD-1k, NUDT, NUDT-sea datasets, respectively. That is because, LCA utilizes the inherent characteristics of infrared small targets and thus coarsely separates the targets from clustering background. By introducing random salt noise to regularize the linear clustering process, we obtain much more reliable results and achieve additional 18.6%, 10.7%, 1.3%, 5.8% improvements on the above four datasets. After that, denseCRF [8] was used as a post-processing module to refine the target probability map (TPM) generated by MCLC. It further introduces 5.9%, 2.6%, 7.7%, 3.2% improvements for the above four datasets, respectively. We visualize the Monte Carlo linear clustering process in Fig. 9. Although easily producing inaccurate results at the beginning of clustering (e.g., iteration number less than 20), MCLC can gradually recover a reliable clustering result. More visualization samples are shown in supplementary material.

2) Type and Intensity of Noise: Fig. 7 (a) shows the change trend of IoU with respect to different noise intensity under three types of common noise (i.e., salt, pepper, and Gaussian). With the increase of noise intensity, the IoU of MCLC with denseCRF under salt noise and pepper noise increases rapidly at the beginning, and reaches their peak scores of 76% and 71% at 0.05 intensity¹. After that, excessive intensity value reduces the saliency of target region and thus results in the decrease of IoU . Moreover, Gaussian noise causes huge performance decrease under any intensity. That is because, Gaussian noise randomly changes the values of all pixels in the image, which is similar to the salt and pepper with high intensity. Since all pixels in the image are influenced by the random noise, the saliency of target is greatly decreased, resulting in dramatic performance drop.

3) Number of Clustering Center: The number of clustering center determines the initial area of each clustering region. Fig. 7 (b) reports the IoU of the results generated by MCLC with denseCRF, MCLC with fixed threshold, and LCA under different number of clustering center. It shows that the IoU firstly shows a rapid increasing trend with the increase number of clustering center, and then reaches the peak score of 76% with 9 clustering centers. Afterward-

¹The intensity value for salt and pepper noise is defined as the area ratio of noise region and image region. Higher values denote that more pixels are replaced by salt or pepper pixels.

Table 3. $IoU(10^{-2})$ values achieved with single point supervision under ideal and real scene.

| | Pseudo Mask (Ideal / Real) | Final Results (Ideal / Real) | Pseudo Mask (Ideal / Real) | Final Results (Ideal / Real) |
|---------|----------------------------|------------------------------|----------------------------|------------------------------|
| ResUnet | (76.3 / 73.7) | (71.6 / 71.2) | (60.2 / 58.8) | (68.0 / 66.5) |
| DNANet | NAAA | (72.9 / 72.0) | NUDT | (70.5 / 69.1) |
| ResUnet | (68.8 / 66.7) | (64.6 / 63.8) | (47.1 / 45.7) | (39.1 / 38.2) |
| DNANet | IRSTD | (62.2 / 60.9) | NUDT-Sea | (40.2 / 38.8) |

Table 4. $IoU(10^{-2})$ values under fixed labelling time. $5.2\times$ (with search) and $8.0\times$ (w/o search) more point-level (P.) labels help to generate better results under fixed labelling time.

| | Sup. | NAAA-SIRST | IRSTD-1k |
|---------------------|------|--|---|
| | | Time Budget: 920s (427×P.) vs (82×E.) (5.2× P. labels) | Time Budget: 1400s (1000×P.) vs (125×E.) (8.0× P. labels) |
| ResUnet | Full | 56.8 | 43.8 |
| ResUnet+MCLC | P. | 71.6 ↑ 14.8 | 64.6 ↑ 20.8 |
| DNANet | Full | 58.7 | 46.1 |
| DNANet+MCLC | P. | 72.1 ↑ 13.4 | 62.2 ↑ 16.1 |

s, the quality of TPM gradually decreases when the number of clustering center further increases. The above results demonstrate that inappropriate number of clustering center will result in over-small or over-large area of initial search region, and thus introduce negative effect on MCLC.

4) Performance with Stronger Supervision: As shown in Table 2, stronger supervision can introduce additional 0.3% ~ 2.6% improvements in term of IoU on the NAAA-SIRST dataset, but at a cost of 121% ~ 700% increase of annotation time (i.e., 1.4s, 3.1s, and 11.2s for single-point, multi-point, and pixel-level annotation in Fig. 3). Similar results can be also found on the other three datasets. Therefore, we argue that single-point annotation is the most economical supervision and is strong enough for SIRST detection.

5) Labelling Position Deviation: Real-world SIRST application may suffer from labelling position deviation shown in Fig. 8 (a). To simulate the labelling position deviation in real world, we manually re-labelled 105 targets in NAAA-SIRST and IRSTD-1k dataset, and generated a distribution map of labelling position in Fig. 8 (b). Based on this distribution map, we re-produced the pseudo masks of whole dataset and generated the final detection results in Table 3. Average results from three trails demonstrate that our method can produce competitive results under practical labelling settings.

6) Time Efficiency: As aforementioned, once the optimal parameters (e.g., type and intensity of noise, number of clustering centers) are searched in one real-world dataset (e.g., NAAA-SIRST), we directly adopt these parameters to new scenes without secondary search (e.g., IRSTD in Table 4). As shown in Table 4, single-point labelling can produce about $5.2\times$ (with parameter search) ~ $8.0\times$ (w/o parameter search) annotation than the pixel-level labelling with the same annotation budget, and thus achieves much better performance. Note that, our MCLC can efficiently

Table 5. $IoU(10^{-2})$, $P_d(10^{-2})$, and $F_a(10^{-6})$ values achieved by different state-of-the-art methods on four benchmark datasets. For IoU and P_d , larger values indicate better performance. For F_a , smaller values indicate better performance. *Unsup.* refers to unsupervised methods. The best single-point supervised results are in **red** and the second best results are in **blue**.

| Method | Sup. | Dataset | | | | | | | | | | | |
|--|---------------|-----------------|--------------|--------------|---------------|--------------|--------------|-----------------|--------------|--------------|---------------------|--------------|--------------|
| | | NUAA-SIRST [11] | | | IRSTD-1k [45] | | | NUDT-SIRST [23] | | | NUDT-SIRST-sea [39] | | |
| Top-Hat [31] | <i>Unsup.</i> | 7.143 | 79.84 | 1012 | 6.222 | 55.48 | 595.6 | 20.72 | 78.41 | 166.7 | 1.17 | 2.680 | 95.37 |
| WSLCM [20] | <i>Unsup.</i> | 1.158 | 77.95 | 5446 | 0.019 | 55.79 | 31547 | 2.283 | 56.82 | 1309 | 0.60 | 10.52 | 7.330 |
| MSLSTIPT [16] | <i>Unsup.</i> | 10.30 | 82.13 | 1131 | 10.37 | 57.05 | 3707 | 8.342 | 47.40 | 888.1 | 0.33 | 0.350 | 6283 |
| NRAM [43] | <i>Unsup.</i> | 12.16 | 74.52 | 13.85 | 4.221 | 49.21 | 27.86 | 6.927 | 56.40 | 19.27 | 0.35 | 17.31 | 3.585 |
| PSTNN [44] | <i>Unsup.</i> | 22.40 | 77.95 | 29.11 | 9.871 | 51.72 | 99.75 | 14.85 | 66.13 | 44.17 | 1.50 | 13.51 | 15.44 |
| ACM [11] | <i>Full</i> | 70.33 | 93.91 | 3.728 | 60.97 | 90.58 | 21.78 | 67.08 | 95.97 | 10.18 | 47.57 | 70.46 | 21.31 |
| ACM + <i>KMeans</i> [26] | <i>Point</i> | 46.16 | 88.59 | 27.45 | 33.41 | 54.42 | 41.82 | 41.18 | 77.88 | 34.86 | 38.15 | 42.45 | 23.25 |
| ACM + <i>SuperPixel</i> [1] | <i>Point</i> | 53.84 | 90.49 | 56.18 | 49.69 | 84.69 | 42.36 | 51.76 | 92.33 | 34.60 | 40.42 | 36.70 | 16.92 |
| ACM + <i>GrabCut</i> [33] | <i>Point</i> | 57.63 | 90.87 | 60.46 | 54.03 | 81.97 | 23.99 | 55.24 | 91.65 | 6.457 | 38.76 | 52.37 | 15.54 |
| ACM + <i>Ours</i> | <i>Point</i> | 67.08 | 92.01 | 19.80 | 56.53 | 82.31 | 25.05 | 57.74 | 92.38 | 16.77 | 43.69 | 55.69 | 9.820 |
| ResUNet[39] | <i>Full</i> | 75.93 | 97.71 | 15.68 | 66.34 | 92.83 | 8.198 | 83.74 | 98.09 | 4.820 | 46.05 | 60.18 | 7.920 |
| ResUNet + <i>KMeans</i> [26] | <i>Point</i> | 22.34 | 39.92 | 156.5 | 33.86 | 80.95 | 10.32 | 42.65 | 89.20 | 42.51 | 31.11 | 51.74 | 44.63 |
| ResUNet + <i>SuperPixel</i> [1] | <i>Point</i> | 60.02 | 93.15 | 23.45 | 59.28 | 88.43 | 24.36 | 62.65 | 93.40 | 26.95 | 36.43 | 56.50 | 30.61 |
| ResUNet + <i>GrabCut</i> [33] | <i>Point</i> | 61.38 | 95.43 | 25.16 | 61.67 | 90.13 | 13.96 | 63.01 | 94.03 | 13.44 | 37.04 | 56.19 | 17.32 |
| ResUNet + <i>Ours</i> | <i>Point</i> | 71.58 | 94.67 | 15.21 | 64.59 | 90.81 | 6.223 | 68.04 | 93.86 | 25.35 | 39.06 | 59.44 | 6.979 |
| DNA _{Net} [23] | <i>Full</i> | 76.24 | 97.71 | 12.80 | 68.44 | 94.77 | 8.806 | 86.36 | 97.39 | 6.897 | 42.17 | 61.60 | 17.19 |
| DNA _{Net} + <i>KMeans</i> [26] | <i>Point</i> | 45.69 | 90.49 | 58.11 | 32.98 | 81.63 | 25.18 | 42.84 | 88.04 | 56.99 | 20.59 | 32.67 | 21.59 |
| DNA _{Net} + <i>SuperPixel</i> [1] | <i>Point</i> | 62.59 | 93.53 | 14.54 | 58.44 | 89.45 | 27.70 | 64.68 | 95.44 | 33.39 | 38.42 | 36.70 | 6.921 |
| DNA _{Net} + <i>GrabCut</i> [33] | <i>Point</i> | 61.14 | 97.71 | 20.53 | 61.00 | 91.15 | 20.87 | 64.00 | 97.03 | 40.65 | 28.08 | 49.79 | 4.816 |
| DNA _{Net} + <i>Ours</i> | <i>Point</i> | 72.86 | 96.95 | 14.43 | 62.23 | 92.13 | 24.14 | 70.52 | 95.55 | 33.20 | 40.23 | 58.32 | 11.29 |
| ISNet [40] | <i>Full</i> | 80.01 | 99.23 | 4.96 | 68.72 | 95.68 | 15.43 | 82.57 | 96.49 | 44.11 | 41.27 | 58.89 | 13.26 |
| ISNet + <i>KMeans</i> [26] | <i>Point</i> | 38.27 | 74.88 | 112.7 | 29.67 | 69.52 | 38.43 | 44.57 | 89.87 | 66.20 | 32.57 | 44.20 | 35.26 |
| ISNet + <i>SuperPixel</i> [1] | <i>Point</i> | 65.82 | 93.73 | 28.26 | 57.93 | 90.75 | 41.74 | 61.59 | 95.61 | 29.59 | 35.21 | 48.56 | 15.32 |
| ISNet + <i>GrabCut</i> [33] | <i>Point</i> | 63.72 | 96.27 | 26.32 | 62.29 | 91.37 | 19.34 | 65.53 | 94.89 | 45.31 | 33.68 | 48.24 | 23.56 |
| ISNet + <i>Ours</i> | <i>Point</i> | 75.93 | 95.44 | 11.04 | 63.19 | 92.58 | 22.23 | 66.87 | 96.22 | 48.59 | 38.29 | 52.27 | 19.97 |
| UIUNet [45] | <i>Full</i> | 76.83 | 97.64 | 12.31 | 65.27 | 92.76 | 12.40 | 75.35 | 93.33 | 27.69 | 41.88 | 55.69 | 11.56 |
| UIUNet + <i>KMeans</i> [26] | <i>Point</i> | 35.17 | 90.90 | 46.53 | 21.26 | 82.52 | 19.43 | 38.27 | 88.27 | 78.09 | 29.33 | 33.17 | 14.22 |
| UIUNet + <i>SuperPixel</i> [1] | <i>Point</i> | 64.23 | 92.11 | 49.34 | 54.42 | 87.30 | 29.82 | 51.68 | 91.32 | 39.40 | 37.83 | 47.22 | 20.09 |
| UIUNet + <i>GrabCut</i> [33] | <i>Point</i> | 63.29 | 93.86 | 33.27 | 60.37 | 89.27 | 14.24 | 64.33 | 92.57 | 33.59 | 37.03 | 51.55 | 26.12 |
| UIUNet + <i>Ours</i> | <i>Point</i> | 74.22 | 96.20 | 16.19 | 64.13 | 90.74 | 14.93 | 69.35 | 91.11 | 38.21 | 39.27 | 53.23 | 23.14 |

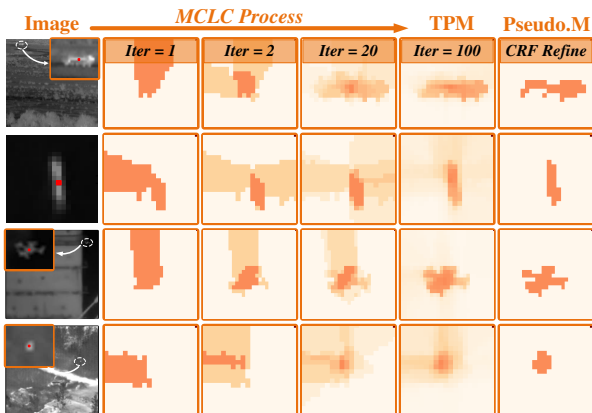


Figure 9. Examples of target probability map (TPM) and the corresponding refined pseudo masks during the MCLC process.

recover a single point annotation to a pixel-level one with only 0.075s (average result from 533 targets) on a PC-level CPU.

4.4. Comparison to the State-of-the-art Methods

We compare our method to several state-of-the-art methods on four benchmark datasets in this subsection. For fair comparison, we used the same hyper-parameters and experimental settings (e.g., special dataset division method in UIUNet[40]) as reported in their original papers when re-implementing comparing methods and retrained all the

Table 6. $IoU(10^{-2})$, $P_d(10^{-2})$, and $F_a(10^{-6})$ values achieved by existing salient object detection (SOD) methods.

| Method | Sup. | NUAA-SIRST | | | IRSTD-1k | | |
|-------------------------------|--------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | | 75.9 | 97.7 | 15.7 | 66.3 | 92.8 | 8.19 |
| ResUNet | <i>Full</i> | 75.9 | 97.7 | 15.7 | 66.3 | 92.8 | 8.19 |
| Res.+PFAN [48] | <i>Point</i> | 27.6 | 83.9 | 49.6 | 22.8 | 76.5 | 66.2 |
| Res.+F3Net [38] | <i>Point</i> | 33.2 | 81.6 | 106.2 | 26.3 | 71.3 | 82.3 |
| ResUNet+MCLC | <i>Point</i> | 71.6 | 94.7 | 15.2 | 64.6 | 90.8 | 6.22 |
| DNA _{Net} | <i>Full</i> | 76.2 | 97.7 | 12.8 | 68.4 | 94.7 | 8.8 |
| DNA.+PFAN [48] | <i>Point</i> | 29.3 | 85.2 | 66.8 | 37.6 | 78.6 | 57.7 |
| DNA.+F3Net [38] | <i>Point</i> | 33.6 | 84.9 | 78.0 | 31.2 | 77.1 | 61.2 |
| DNA_{Net}+MCLC | <i>Point</i> | 72.9 | 96.9 | 14.4 | 62.2 | 92.1 | 24.1 |

models from scratch on these four datasets.

1) **Quantitative and Qualitative Results:** Quantitative results are shown in Table 5. Our Monte Carlo linear clustering method (i.e., MCLC) achieves much better performance over those unsupervised methods. Compared to MCLC with other single-point supervised methods, MCLC can produce better pseudo masks and thus achieve better performance in term of P_d , F_a , and IoU . Compared to the fully-supervised counterparts, MCLC enables the respective networks to generate comparable result with 87% annotation time reduction. Qualitative results on four datasets are shown in Fig. 10. Compared to fully supervised methods, our MCLC can help existing detection methods to produce comparable results in a more time-efficient manner, especially for those point targets (e.g., img-1, img-4) and spot targets (e.g., img-2, img-3, img-5, img-6). Moreover, Fig. 10 (img-7, img-8) demonstrates the robustness of our method on the dense target scenarios. Readers can refer to

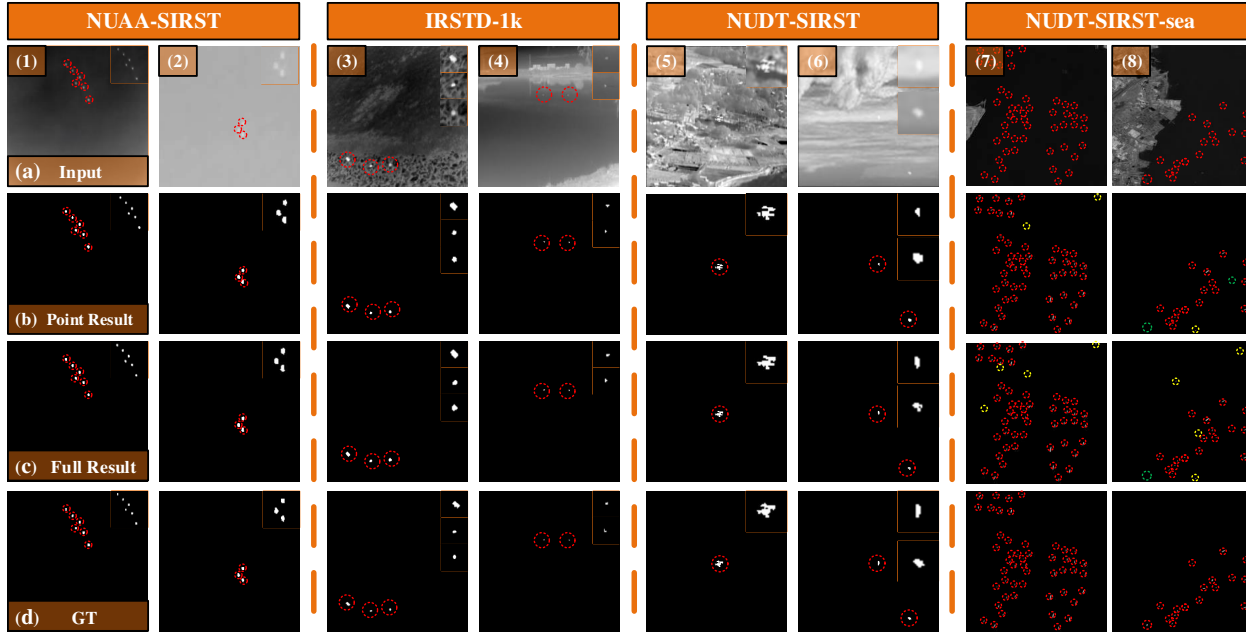


Figure 10. Qualitative results achieved by DNANet[23] under (b) point-level and (c) pixel-level supervision. The correctly detected target, false alarm, and miss detection areas are highlighted by red, yellow, and green dotted circles. More visualization results are shown in supplementary material.

the supplementary material for more visualization results.

2) **Compared with SOD Methods:** Quantitative results in Table 6 show that our MCLC performs much better than those salient object detection (SOD) methods. That is because, although the small targets are salient in local region, a single point label can not provide sufficient supervision to train high-performance SOD methods (e.g., PFAN [48], F3Net [38]). In contrast, our MCLC can make full use of single point labels to obtain higher accuracy.

5. Conclusion

In this paper, we propose a simple yet effective pipeline for single-frame infrared small target detection with single-point annotation. First, we found that SIRST is generally salient in the local small region and exhibits high energy concentricity. Then, based on this observation, we proposed a linear clustering approach (LCA) to coarsely separate the targets from clustering background by measuring the color and spatial distance with clutter background. To further refine the results, we design a Monte Carlo regularization approach to constrain the clustering process and continuously expand the the point-level annotation to a high-quality pixel-level one. Extensive experiments on four benchmark datasets demonstrate that our method achieves comparable performance with fully-supervised methods. We hope our study can draw attention to the research on weakly supervised SIRST detection.

Acknowledgement: This work was supported in part

by the Foundation for Innovative Research Groups of the National Natural Science Foundation of China under Grant 61921001 and National Natural Science Foundation of China 62001478.

References

- [1] Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE TPA-MI*, 34(11):2274–2282, 2012. 3, 8
- [2] Jiwoon Ahn, Sunghyun Cho, and Suha Kwak. Weakly supervised learning of instance segmentation with inter-pixel relations. In *CVPR*, 2019. 6
- [3] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *CVPR*, 2018. 6
- [4] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. Whats the point: Semantic segmentation with point supervision. In *ECCV*, pages 549–565. Springer, 2016. 1, 3
- [5] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. Whats the point: Semantic segmentation with point supervision. In *ECCV*, 2016. 3
- [6] CL Philip Chen, Hong Li, Yantao Wei, Tian Xia, and Yuan Yan Tang. A local contrast method for small infrared target detection. *IEEE TGRS*, 52(1):574–581, 2013. 2
- [7] Liangyu Chen, Tong Yang, Xiangyu Zhang, Wei Zhang, and Jian Sun. Points as queries: Weakly semi-supervised object detection by points. In *CVPR*, pages 8823–8832, 2021. 3
- [8] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Semantic image segmen-

- tation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*, 2014. 7
- [9] Yimian Dai and Yiquan Wu. Reweighted infrared patch-tensor model with both nonlocal and local priors for single-frame small target detection. *IEEE JSTAR*, 10(8):3752–3767, 2017. 2
- [10] Yimian Dai, Yiquan Wu, Yu Song, and Jun Guo. Non-negative infrared patch-image model: Robust target-background separation via partial sum minimization of singular values. *Infrared Physics & Technology*, 81:182–194, 2017. 2
- [11] Yimian Dai, Yiquan Wu, Fei Zhou, and Kobus Barnard. Asymmetric contextual modulation for infrared small target detection. In *WACV*, pages 950–959, 2021. 2, 3, 6, 8
- [12] Yimian Dai, Yiquan Wu, Fei Zhou, and Kobus Barnard. Attentional local contrast networks for infrared small target detection. *IEEE TGRS*, 2021. 2, 6
- [13] He Deng, Xianping Sun, Maili Liu, Chaohui Ye, and Xin Zhou. Small infrared target detection based on weighted local difference measure. *IEEE TGRS*, 54(7):4204–4214, 2016. 1
- [14] Suyog D Deshpande, Meng Hwa Er, Ronda Venkateswarlu, and Philip Chan. Max-mean and max-median filters for detection of small targets. In *Signal and Data Processing of Small Targets*, volume 3809, pages 74–83. International Society for Optics and Photonics, 1999. 2
- [15] John Duchi, Elad Hazan, and Yoram Singer. Adaptive sub-gradient methods for online learning and stochastic optimization. *JMLR*, 12(7), 2011. 6
- [16] Chenqiang Gao, Deyu Meng, Yi Yang, Yongtao Wang, Xiaofang Zhou, and Alexander G Hauptmann. Infrared patch-image model for small target detection in a single image. *IEEE TIP*, 22(12):4996–5009, 2013. 2, 8
- [17] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, pages 249–256, 2010. 6
- [18] Jinhui Han, Yong Ma, Bo Zhou, Fan Fan, Kun Liang, and Yu Fang. A robust infrared small target detection algorithm based on human visual system. *IEEE GRSL*, 11(12):2168–2172, 2014. 2
- [19] Jinhui Han, Saed Moradi, Iman Faramarzi, Chengyin Liu, Honghui Zhang, and Qian Zhao. A local contrast method for infrared small-target detection utilizing a tri-layer window. *IEEE GRSL*, 17(10):1822–1826, 2019. 2
- [20] Jinhui Han, Saed Moradi, Iman Faramarzi, Honghui Zhang, Qian Zhao, Xiaojian Zhang, and Nan Li. Infrared small target detection based on the weighted strengthened local contrast measure. *IEEE GRSL*, 2020. 2, 8
- [21] Anna Khoreva, Rodrigo Benenson, Jan Hosang, Matthias Hein, and Bernt Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *CVPR*, 2017. 3
- [22] Sungho Kim and Joohyoung Lee. Scale invariant small target detection by optimizing signal-to-clutter ratio in heterogeneous background for infrared search and track. *Pattern Recognition*, 45(1):393–406, 2012. 2
- [23] Boyang Li, Chao Xiao, Longguang Wang, Yingqian Wang, Zaiping Lin, Miao Li, Wei An, and Yulan Guo. Dense nested attention network for infrared small target detection. *IEEE TIP*, 2022. 1, 2, 3, 6, 8, 9
- [24] Yanwei Li, Hengshuang Zhao, Xiaojuan Qi, Yukang Chen, Lu Qi, Liwei Wang, Zeming Li, Jian Sun, and Jiaya Jia. Fully convolutional networks for panoptic segmentation with point-based supervision. *IEEE TPAMI*, 2022. 3
- [25] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *CVPR*, 2016. 3
- [26] J MacQueen. Classification and analysis of multivariate observations. In *5th Berkeley Symp. Math. Statist. Probability*, pages 281–297, 1967. 3, 8
- [27] Kevis-Kokitsi Maninis, Sergi Caelles, Jordi Pont-Tuset, and Luc Van Gool. Deep extreme cut: From extreme points to object segmentation. In *CVPR*, pages 616–625, 2018. 1, 3
- [28] R Austin McEver and BS Manjunath. Pcams: Weakly supervised semantic segmentation using point supervision. *arXiv preprint arXiv:2007.05615*, 2020. 3
- [29] Dim P Papadopoulos, Jasper RR Uijlings, Frank Keller, and Vittorio Ferrari. Extreme clicking for efficient object annotation. In *ICCV*, pages 4930–4939, 2017. 1, 3
- [30] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, pages 8026–8037, 2019. 6
- [31] Jean-Francois Rivest and Roger Fortin. Detection of dim targets in digital infrared imagery by morphological image processing. *Optical Engineering*, 35(7):1886–1893, 1996. 2, 8
- [32] Alan R Robertson. Historical development of cie recommended color difference equations. *Color Research & Application*, 15(3):167–170, 1990. 3
- [33] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. Grabcut interactive foreground extraction using iterated graph cuts. *ACM TOG*, 23(3):309–314, 2004. 8
- [34] Yang Sun, Jungang Yang, and Wei An. Infrared dim and small target detection via multiple subspace learning and spatial-temporal patch-tensor model. *IEEE TGRS*, 2020. 1
- [35] Huan Wang, Luping Zhou, and Lei Wang. Miss detection vs. false alarm: Adversarial learning for small object segmentation in infrared images. In *CVPR*, pages 8509–8518, 2019. 1
- [36] Xin Wang, Guofang Lv, and Lizhong Xu. Infrared dim target detection based on visual attention. *Infrared Physics & Technology*, 55(6):513–521, 2012. 2
- [37] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *CVPR*, 2020. 6
- [38] Jun Wei, Shuhui Wang, and Qingming Huang. F³net: fusion, feedback and focus for salient object detection. In *AAAI*, 2020. 8, 9
- [39] Tianhao Wu, Boyang Li, Yihang Luo, Yingqian Wang, Chao Xiao, Ting Liu, Jungang Yang, Wei An, and Yulan Guo. Mtnet: Multi-level transunet for space-based infrared tiny ship detection. *arXiv preprint arXiv:2209.13756*, 2022. 6, 8

- [40] Xin Wu, Danfeng Hong, and Jocelyn Chanussot. Uiu-net: U-net in u-net for infrared small object detection. *IEEE Transactions on Image Processing*, 32:364–376, 2022. 8
- [41] Xuehui Yu, Pengfei Chen, Di Wu, Najmul Hassan, Guorong Li, Junchi Yan, Humphrey Shi, Qixiang Ye, and Zhenjun Han. Object localization under single coarse point supervision. In *CVPR*, pages 4868–4877, 2022. 3
- [42] Fei Zhang, Chaochen Gu, Chenyue Zhang, and Yuchao Dai. Complementary patch for weakly supervised semantic segmentation. In *ICCV*, 2021. 6
- [43] Landan Zhang, Lingbing Peng, Tianfang Zhang, Siying Cao, and Zhenming Peng. Infrared small target detection via non-convex rank approximation minimization joint ℓ_2 , ℓ_1 norm. *Remote Sensing*, 10(11):1821, 2018. 2, 8
- [44] Landan Zhang and Zhenming Peng. Infrared small target detection based on partial sum of the tensor nuclear norm. *Remote Sensing*, 11(4):382, 2019. 2, 8
- [45] Mingjin Zhang, Rui Zhang, Yuxiang Yang, Haichen Bai, Jing Zhang, and Jie Guo. Isnet: Shape matters for infrared small target detection. In *CVPR*, pages 877–886, 2022. 1, 2, 6, 8
- [46] Shilong Zhang, Zhuoran Yu, Liyang Liu, Xinjiang Wang, Aojun Zhou, and Kai Chen. Group r-cnn for weakly semi-supervised object detection with points. In *CVPR*, pages 9417–9426, 2022. 3
- [47] Wei Zhang, Mingyu Cong, and Liping Wang. Algorithms for optical weak small targets detection and tracking. In *ICNNSP*, volume 1, pages 643–647. IEEE, 2003. 1, 3
- [48] Ting Zhao and Xiangqian Wu. Pyramid feature attention network for saliency detection. In *CVPR*, 2019. 8, 9
- [49] Tianyi Zhao and Zhaozheng Yin. Weakly supervised cell segmentation by point annotation. *IEEE TMI*, 40(10):2736–2747, 2020. 1, 3
- [50] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, pages 2921–2929, 2016. 1, 3
- [51] Hu Zhu, Shiming Liu, Lizhen Deng, Yansheng Li, and Fu Xiao. Infrared small target detection via low-rank tensor completion with top-hat regularization. *IEEE TGRS*, 58(2):1004–1016, 2019. 2