

STARLINC: Satellite Trail Artifact Removal using Inter-Frame Correlation

Shingeon Kim¹, Hyeyoon Lee¹, Dain Kwon¹, Kanghyun Choi¹, Sunjong Park¹, Mi-Ryang Kim², Jeong-Eun Lee², and Jinho Lee¹

¹ Department of Electrical and Computer Engineering, Seoul National University

² Department of Physics and Astronomy, Seoul National University

{happyriosshs, hylee817, dain.kwon, kanghyun.choi, ryan0507, mi.ryang, lee.jeongeun, leejinho}@snu.ac.kr

Abstract. The rapid expansion of low Earth orbit satellites such as Starlink is increasingly contaminating astronomical surveys. In practice, contaminated images are often identified through inspection. However, modern surveys generate terabytes of data each night, making manual screening infeasible and necessitating reliable automated methods for satellite trail removal. Unfortunately, existing general-domain line detection methods fail to generalize to astronomical images due to domain mismatch, which are mostly grayscale with sparse bright stars and have a low signal-to-noise ratio. Moreover, training new models from scratch is impractical due to the lack of large-scale annotated astronomical datasets. To address these challenges, we introduce STARLINC, the first ML-based framework for satellite trail removal without requiring tedious pixel-level annotation of astronomical images. STARLINC combines synthetic satellite trail generation for training, inter-frame differential maps from temporally adjacent exposures to highlight transient trails, and heatmaps to provide additional localization cues for pixel-level segmentation. Extensive experiments on real-world data demonstrate substantial improvements over baselines, establishing STARLINC as a scalable solution for next-generation astronomical surveys. Code is available at <https://github.com/starioKim/STARLINC>

Keywords: satellite trail removal · astronomical imagery

1 Introduction

Satellite contamination [25] has emerged as a significant and rapidly worsening obstacle to astronomical surveys³. This challenge is driven primarily by the unprecedented growth of low Earth orbit (LEO) satellite constellations [15, 22]. With the number of active LEO satellites now exceeding 12,000 and increasing



Fig. 1: An example of contaminated images.

³ Throughout this paper, we use *astronomical survey* to denote datasets of telescopic images acquired for scientific analysis.

at an annual rate of approximately 30% [5, 23], satellite trails are no longer rare artifacts but a recurring source of image degradation across modern surveys. Previous studies estimate that they affect about 18% of astronomical survey images [23, 25], and recent SPHEREx observations report contamination in $73.3^{+1.3}_{-1.2}\%$ of images, with $2.18^{+0.11}_{-0.09}$ trails per exposure on average [7]. These contaminations occur when satellites reflect sunlight during exposure, producing bright linear trail segments (see Figure 1). Such trails introduce artificial light flux into the observations, leading to substantial errors in downstream astrophysical measurements [6, 15].

Unfortunately, the typical astronomical research pipeline still relies on manual inspection, due to the lack of a reliable automated trail removal method. As modern astronomical surveys generate terabytes of data each night, manual inspection can no longer keep pace with the volume of incoming images. Moreover, once an image is identified as containing a satellite trail, manual workflows typically exclude the entire image from the dataset, resulting in substantial loss of otherwise valuable observational data.

One might expect satellite trails would be easy to recognize for modern computer vision methods, as the trails appear as simple straight lines against a starry, mostly black sky. Contrary to this intuition, our empirical evaluation shows that existing line detection methods remain largely ineffective on real astronomical data. Classical approaches (Hough transform [9, 11], Radon-based detectors [30]) are highly sensitive to background noise. Several deep learning-based general domain line detectors [24, 29, 44] fail to reliably detect satellite trails, as they do not account for the extreme sparsity and distinct photometric statistics of astronomical imagery. Meanwhile, astronomy-specific approaches [10, 18, 27, 34] typically rely on pixel-level annotations and exhibit limited generalization. Despite strong performance on their original benchmarks, these methods degrade substantially when evaluated on unseen datasets.

To this end, we propose STARLINC, the first ML-based framework for satellite trail removal without requiring pixel-level annotation of astronomical images. STARLINC performs pixel-level segmentation to selectively mask contaminated pixels while preserving valid astronomical data.

Through empirical analyses of datasets, we identify two key characteristics:

1. The satellite trails typically appear as sharp, narrow lines that occupy only a small fraction of the image.
2. Astronomical surveys acquire multiple frames as a time series for noise reduction, which are largely identical.

For conventional machine learning methods, both characteristics can be detrimental: the high similarity among frames may confuse the model during training, while the weak signals of thin trails are easily obscured by other image structures.

In contrast, STARLINC is explicitly designed to exploit these properties, resulting in a simple yet effective framework. First, we provide a synthetic satellite trail generation method to enable effective training of a trail segmentation model without tedious manual pixel-level annotation. Secondly, STARLINC leverages temporal inter-frame differential maps across consecutive observations to focus

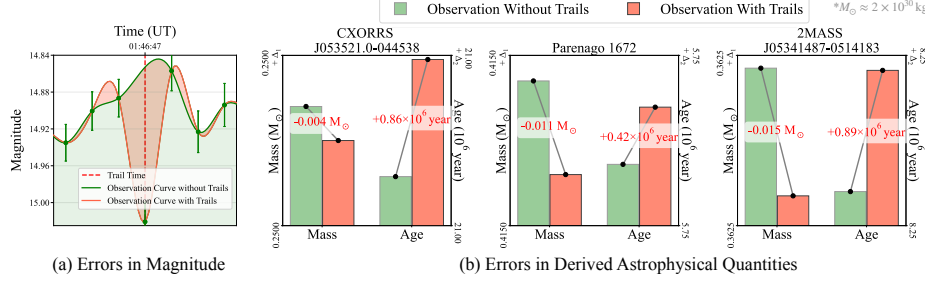


Fig. 2: Motivational study of the impact of satellite contamination. (a) A satellite trail intersecting a star perturbs its measured brightness (magnitude). (b) The resulting bias leads to errors in derived stellar properties (e.g., age and mass).

on the transient satellite trail. Lastly, we obtain heatmaps highlighting trail-relevant regions from a classifier trained to predict trail presence, and use them as additional visual cues for pixel-level segmentation. We conduct extensive experiments on real astronomical datasets from an ongoing astronomical survey [14, 20], where STARLINC consistently outperforms existing baselines by up to +102% relative mIoU improvement over satellite trail segmentation. We further evaluate cross-dataset transfer to ZTF [14] and NGC-region [20] observations, showing that a STARLINC generalizes to unseen observing conditions.

Our contributions are summarized as follows:

- We propose STARLINC, a practical satellite trail removal framework that preserves scientifically valid content by masking only contaminated pixels.
- We reduce reliance on pixel-level annotations by introducing synthetic trail generation for training and enhancing localization via temporal inter-frame differential maps and pseudo-localization via heatmaps.
- We validate STARLINC on real 7DT survey data [20], demonstrating up to +102% relative mIoU over the strongest baseline, and further show cross-dataset transfer to ZTF [14] and NGC-region [20] observations.

2 Motivation

2.1 Scientific Impact of Satellite Trail Contamination

In this section, we investigate the detrimental impact of satellite trail contamination on astronomical research. Satellite trails corrupt raw observation data, leading to systematic errors in scientific measurements. We illustrate the risks by presenting real-world cases of degradation induced by satellite trails.

The primary threat of satellite trails is the corruption of stellar photometry. A single contaminated exposure can inflate variability metrics or distort statistics (See Figure 2(a)). One datapoint affected by a satellite trail distorts the observed

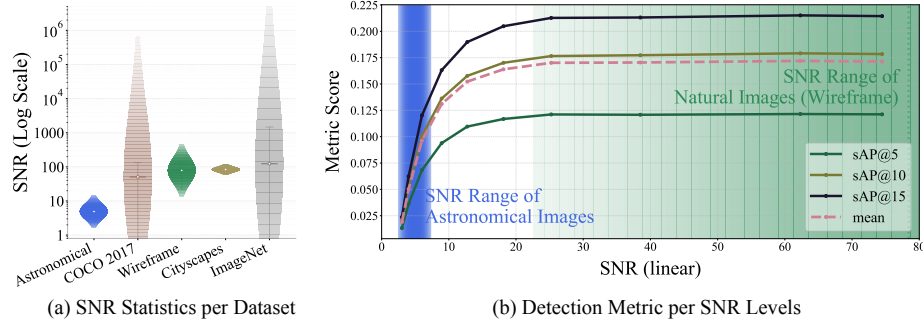


Fig. 3: (a) Visualization of SNR distribution per dataset. The astronomical images consists of very low SNR, compared to general domain datasets. (b) Line detection performance under varying signal-to-noise ratios (SNR) of images. Deep learning detectors degrade rapidly in the low-SNR region. The metric is structural Average Precision (sAP), which measures the accuracy of predicted lines against the ground truth. @5, @10, and @15 denote the pixel tolerance for matching prediction and ground-truth.

trend. This highlights that a single streak can impair the reliability of time-domain analyses. The findings call for robust mitigation of such trails.

Such errors in magnitude and brightness can propagate into significant inaccuracies in derived physical quantities (e.g., stellar mass and age). An actual demonstration is shown in Figure 2(b), with three real-world stellar sources [39]. For this illustrative analysis, we assume that the contaminated sources are main-sequence stars and use the main-sequence lifetime as an age proxy. We propagate the magnitude bias using $F \propto 10^{-0.4\Delta m}$, $L \propto F$ for fixed distance, $L \propto M^{3.5}$, and $t_{\text{MS}} \propto M/L$, with other properties fixed. The results introduce substantial bias, up to 3.9% deviation in mass, and 10% deviation in age. Thus, satellite trails remain a major unresolved source of systematic error in stellar photometry.

2.2 Limitations of Existing Line Detection Techniques

To understand the unique challenge of satellite trail detection, we identify the fundamental gap between astronomical and natural imagery. We characterize this gap using the signal-to-noise ratio (SNR), which quantifies the strength of meaningful signals relative to background noise. We further demonstrate that this discrepancy limits existing line detection methods for astronomical images.

From Figure 3(a), we observe a significant disparity in SNR between two image domains: astronomical images exhibit much lower SNR, compared to images from general vision datasets. While astronomical images exhibit SNR values of at most 16.84, general image datasets show median SNR values ranging from 35 to 116. Unlike ordinary photography, which receives abundant photon flux, astronomical observations operate in a photon-starved regime where faint signals are often dominated by instrumental and atmospheric noise. As a result, astronomical survey images operate in an inherently low-SNR regime.

This observation sheds light on how linear structures exist in each domain. In natural images, lines mostly exist as high-contrast boundaries between adjacent objects. Satellite trails, however, are isolated features independent of stellar objects or noisy backgrounds. Without surrounding contextual cues or strong contrast, these trails are faint and extremely thin signals embedded within noise.

Such features counter the assumptions underlying existing line detection methods, which are devised for high-contrast edges and rich contextual information. To demonstrate this limitation, we evaluate DeepLSD [29], a representative learning-based line detector for natural images, on Wireframe images [16] while progressively adding synthetic noise to reduce their SNR to astronomical levels. As shown in Figure 3(b), detection performance remains stable in the high-SNR regime typical of natural images ($\text{SNR} > 70$), but degrades rapidly as the SNR approaches that of astronomical imagery. In the single-digit SNR regime, detection metrics drop sharply, confirming that conventional vision-based methods are ill-suited for detecting faint linear structures under astronomical noise and motivating the need for a tailored approach. A visualized example of this noise-degradation process is presented in Figure 12 at Sec. H.

3 Related Works

3.1 Classical Methods for Line Detection

There were several early attempts to detect line structures in images through mathematical heuristics. The Hough [11] and Radon [30] transforms reformulate the task as identifying local peaks by mapping colinear pixels to a point. The LSD [13], ED [37], EDLines [2], AG3line [42], and ELSEd [35] construct line segments faster by aggregating pixels with consistent gradients from local anchors. These detectors are often paired with the Canny edge detector [9]. They often lack flexibility when pixel-intensity variations are too subtle to elicit stable responses, thus requiring extensive parameter tuning.

3.2 Deep Learning Approaches for Line Detection

To overcome the limitations of prior handcrafted heuristics, deep learning techniques have been integrated into line detection. Such methods often reframe the line segment detection problem into alternative learning tasks, such as region coloring [40] or graph inference [43] by constructing graphs for detected junction points in an image. Another strategy [36, 44] is to apply the Hough transform to CNN features, training the model in the parameter space. DeepLSD [29] refines line segment predictions by leveraging deep image gradients to produce accurate endpoints and orientations. DT-LSD [17] introduces deformable transformers that adapt receptive fields to capture geometric variations in line structure. Most learned line detectors are designed for general-domain images, where lines typically correspond to object boundaries with rich contextual information.

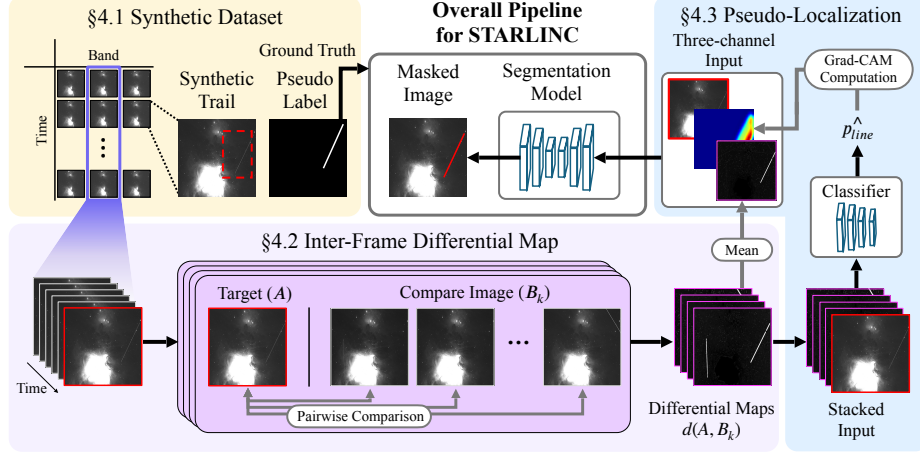


Fig. 4: Overall framework of STARLINC. The pipeline consists of three stages. Synthetic data generation (§ 4.1): Artificial satellite trails are injected into non-trail images to generate image-label pairs for segmentation model training. Inter-frame differential map (§ 4.2): Leveraging repeated exposures of the same sky region, STARLINC suppresses static celestial structures and emphasizes transient satellite trails using locally aggregated temporal differences. Pseudo-localization via heatmaps (§ 4.3): We convert an image-level trail classifier into a spatial cue by extracting Grad-CAM activation heatmaps, which provide coarse pseudo-localization of the trail and guide the downstream segmentation network without pixel-level annotations.

3.3 Automated Satellite Trail Detection in Astronomical Images

To address the challenges posed by the unique properties of astronomical trails, several works [4, 10, 12, 18, 19, 26–28, 34] have focused on detecting satellite trails in astronomical survey. Early techniques leveraged geometric algorithms such as boundary-tracing [19], and morphological filtering [4] to capture and remove trails from telescope exposures. Several methods [12, 18, 27] train CNN models or ML-based classifiers to identify satellite trails. More recent approaches, such as ASTA, MaxiMask, UNetLSD [10, 28, 34], combine segmentation networks with geometric detectors or multi-class contaminant labels to localize trails at the pixel level.

For model training, these methods require pixel-level trail annotations as supervision. However, obtaining annotations relies on expensive manual examination. This motivates the necessity for approaches that can train trail detectors without such expensive pixel-level supervision.

4 Method

Our framework (Figure 4) comprises four components that progressively generate pixel-level masks of satellite trails while preserving uncontaminated image regions. To address the lack of pixel-level annotation and class imbalance problem,

we first generate synthetic training samples by injecting simulated satellite trails into clean astronomical images (§ 4.1). We then exploit temporal redundancy in repeated observations by computing inter-frame differential maps across adjacent exposures, which suppress persistent celestial backgrounds and emphasize transient satellite artifacts (§ 4.2). These inter-frame representations are used to obtain heatmaps that provide pseudo-localization cues for the later segmentation stage, which are extracted from a ResNet-based classifier trained to predict trail presence (§ 4.3). Finally, a U-Net-based [31] segmentation model takes the original image, the corresponding inter-frame differential map, and the heatmap to produce pixel-level masks without requiring manual annotations (§ 4.4).

4.1 Synthetic Trail Generation for Training

We introduce a synthetic trail generation strategy tailored to astronomical survey to address two limitations of real survey data. First, large-scale datasets typically lack pixel-level trail annotations, and off-the-shelf line detectors generalize poorly to astronomical survey. Second, the data are highly imbalanced, with only about 20% of images containing satellite trails. To address both issues, we generate synthetic trails without manual masks for scalable training and for controlled sampling of trail occurrences.

We generate synthetic training samples by injecting simulated satellite trails into clean astronomical images. The resulting synthetic data provide paired images and pixel-level ground-truth masks, enabling direct training of the segmentation model without manual pixel-level annotations.

To maintain visual authenticity and physical fidelity, each synthetic trail is parameterized by four attributes: length L , orientation θ , thickness τ , and brightness m . The trail length, orientation, thickness, and brightness are sampled as

$$L \sim \text{clip}(\mathcal{N}(D, \sigma_L^2); \alpha, \beta), \quad (1)$$

$$\theta \sim \mathcal{U}(0, 2\pi), \quad (2)$$

$$\tau \sim \text{B}(a, b), \quad (3)$$

$$m \sim \mathcal{U}(m_{\min}, m_{\max}). \quad (4)$$

Here, $\mathcal{N}(\mu, \sigma^2)$ denotes a normal distribution with mean μ and variance σ^2 , $\mathcal{U}(u_{\min}, u_{\max})$ denotes a uniform distribution over $[u_{\min}, u_{\max}]$, and $\text{B}(a, b)$ denotes a beta distribution with shape parameters a and b .

The trail position is sampled uniformly over the image plane. To better reflect real astronomical observations, the trail thickness τ is drawn from a Beta distribution that favors smaller values. This bias increases the frequency of thin trails in the synthetic data, which is important because satellite trails often appear as faint and narrow structures under low signal-to-noise conditions. By emphasizing such thin trails during training, the model learns to detect subtle linear artifacts that are difficult to distinguish from background noise.

4.2 Inter-Frame Differential Map for Background Suppression

STARLINC leverages a key property of time-domain astronomical surveys: the same sky region is repeatedly imaged under nearly identical conditions. As a result, static celestial structures (e.g., stars and galaxies) remain highly consistent while transient artifacts such as satellite trails appear only in a subset of frames. We explicitly capture these temporal cues using inter-frame differential maps.

Given two normalized grayscale images $x_i, x_j \in \mathbb{R}^{H \times W}$, we construct a pixel-wise differential map $M_{i,j} \in \mathbb{R}^{H \times W}$. Rather than computing a purely point-wise distance, we aggregate local structural differences within a sliding window:

$$M_{i,j}(m, n) = \sum_{(u,v) \in \mathcal{N}_k(m,n)} d(x_i(u, v), x_j(u, v)), \quad (5)$$

where $d(\cdot, \cdot)$ denotes a distance function and $\mathcal{N}_k(m, n)$ is a $k \times k$ neighborhood of pixel (m, n) . This local aggregation enables the differential map to capture regional variations rather than pixel noise, thereby improving robustness.

An important consideration in selecting $d(\cdot, \cdot)$ is robustness to inherent observational noise of ground-based surveys. To this end, we adopt the Structural Similarity Index (SSIM) to define the distance as its complement:

$$d(x_i, x_j) = 1 - \text{SSIM}(x_i, x_j) = 1 - \frac{(2\mu_i\mu_j + C_1)(2\sigma_{ij} + C_2)}{(\mu_i^2 + \mu_j^2 + C_1)(\sigma_i^2 + \sigma_j^2 + C_2)}, \quad (6)$$

where μ_i and μ_j denote local means, σ_i^2 and σ_j^2 denote local variances, and σ_{ij} denotes the local covariance computed within the window. C_1 and C_2 are small constants that stabilize the computation when local statistics are close to zero. A visualized example of this process is presented in Figure 7 at Sec. C.

For each target frame, the mean differential map is computed by aggregating pairwise SSIM distances against the remaining frames. Since static background structures are shared across frames, they are largely suppressed despite natural observational noise. In contrast, a satellite trail appearing in only one frame induces a localized peak in the distance measure. These differential maps serve as the primary input representation for the subsequent segmentation stages (§4.3).

4.3 Pseudo-localization via Activation Maps

Even with temporal differential maps, the segmentation stage benefits from a spatial prior indicating where the trail is likely to appear. To provide a lightweight localization signal, we use an auxiliary trail-presence classifier as a pseudo-localizer. The resulting gradient-based activation maps [33] highlight the regions most responsible for the trail prediction and guide the segmentation network.

We train a classifier that takes input of the original grayscale frames and their corresponding differential maps. The classifier outputs a probability of trail presence, $\hat{p}_{\text{line}}$. To obtain a spatial localization cue from the trained classifier, we extract activation maps using Grad-CAM [33]. Let $\{f^k\}_{k=1}^K$ denote the K feature maps of the final convolutional layer, and let y_c denote the logit for

the trail-present class. We compute channel-wise importance weights by globally averaging the gradient of y_c with respect to each feature map.

$$\alpha_k = \frac{1}{HW} \sum_{i=1}^H \sum_{j=1}^W \frac{\partial y_c}{\partial f^k(i, j)}. \quad (7)$$

These weights are then used to produce a coarse localization heatmap via a weighted combination followed by a ReLU activation,

$$H_{GC} = \text{ReLU} \left(\sum_{k=1}^K \alpha_k f^k \right). \quad (8)$$

The resulting heatmap $H_{GC} \in \mathbb{R}^{H' \times W'}$ highlights spatial regions that contribute most to trail prediction and are directly fed into the segmentation network (§ 4.4), without requiring pixel-level supervision.

4.4 Overall Pipeline for Trail Masking

The final stage of STARLINC produces pixel-level binary trail masks using a segmentation network trained on synthetic supervision from § 4.1. For each target image, the segmentation network receives a three-channel input consisting of the original grayscale image I , the mean inter-frame differential map $\bar{M}$, and the Grad-CAM heatmap H_{GC} :

$$X = [I, \bar{M}, H_{GC}] \in \mathbb{R}^{H \times W \times 3}. \quad (9)$$

Given this composite input, the U-Net segmentation network f_θ predicts a single-channel trail probability map,

$$\hat{Y} = f_\theta(X) \in [0, 1]^{H \times W}, \quad (10)$$

which is thresholded at inference time to obtain the final binary trail mask.

Training. During training, STARLINC constructs the same three-channel representation that is used at inference. We first select a subset of target images within each temporal group and inject synthetic trails into those images. Each temporal group contains 30 repeated exposures of the same sky field. For a synthetic target image I^{syn} , the mean differential map $\bar{M}^{\text{syn}}$ is computed by comparing the target image against the remaining exposures in the group using the procedure in § 4.2. The trail-presence classifier then produces a Grad-CAM heatmap H_{GC}^{syn} as described in § 4.3. The resulting training input is

$$X^{\text{syn}} = [I^{\text{syn}}, \bar{M}^{\text{syn}}, H_{GC}^{\text{syn}}] \in \mathbb{R}^{H \times W \times 3}. \quad (11)$$

The segmentation network predicts $\hat{Y}^{\text{syn}} = f_\theta(X^{\text{syn}})$ and is supervised by the known synthetic binary trail mask Y^{syn} . Thus, pixel-level supervision is obtained without manually annotated real trail masks.

Table 1: Segmentation performance comparison.

Category	Method	Metrics					
		mIoU	Dice	Prec.	Recall	ROC AUC	PR AUC
Classical	Hough [11]	0.142	0.202	0.785	0.293	0.692	0.741
	Radon [30]	0.000	0.001	0.653	0.001	0.670	0.696
	LSD [13]	0.074	0.126	0.077	0.517	0.776	0.714
Learning-based	DeepLSD [29]	0.116	0.192	0.120	0.642	0.867	0.827
	DT-LSD [17]	0.240	0.332	0.663	0.507	0.769	0.838
Trail detection	ASTA [34]	0.148	0.217	0.399	0.236	0.730	0.511
	MaxiMask [28]	0.155	0.268	0.838	0.159	0.795	0.830
	UnetLSD [10]	0.197	0.329	0.303	0.359	0.622	0.399
Proposed	STARLINC	0.485	0.653	0.815	0.545	0.918	0.865

Inference. Given a target image I from a temporal group of repeated exposures, STARLINC computes the two auxiliary inputs: the mean differential map $\bar{M}$ and the classifier heatmap H_{GC} . The trained U-Net predicts

$$\hat{Y} = f_{\theta}([I, \bar{M}, H_{GC}]), \quad (12)$$

and converts the probability map into a binary satellite-trail mask. The predicted mask can be applied to the original astronomical image to exclude contaminated pixels while retaining uncontaminated regions.

5 Experiments

5.1 Experimental Settings

To evaluate the performance of STARLINC on real astronomical observations, we construct the primary dataset from 7DT [20] observations of the Orion Molecular Clouds region. The 7DT project is a wide-field ground-based survey that repeatedly observes the same sky field, providing multi-epoch exposures. Temporally adjacent frames are used to construct inter-frame differential maps with window size $k = 5$, and we evaluate our method on classification, segmentation, and trail removal tasks. The dataset contains 1,140 samples, which we divided into 720 training, 180 validation, and 240 test samples. For synthetic trail generation, the parameters in Eq. (1)–(4) are set as $D = 0.5 \min(H, W)$, $\sigma_L = 0.25D$, $\alpha = 0.1D$, $\beta = 0.8D$, $(a, b) = (2, 6)$, and $(m_{\min}, m_{\max}) = (0.35 \cdot 255, 0.65 \cdot 255)$. More detailed experiment settings are included in appendix Sec. A.

5.2 Evaluation on Satellite Trail Segmentation Task

In this section, we evaluate the satellite trail segmentation ability of STARLINC by measuring pixel-level localization. Table 1 reports mIoU, Dice, Precision, Recall, ROC-AUC, and PR-AUC across classical line detectors, learning-based line

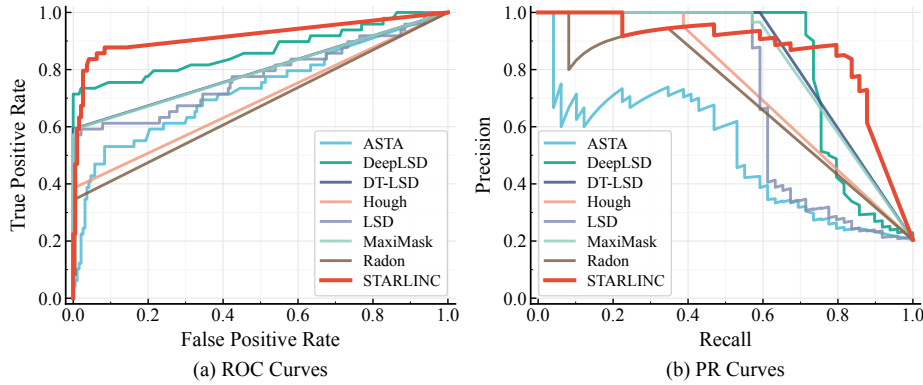


Fig. 5: Performance evaluation for segmentation model using (a) ROC curve and (b) PR curve showing performance across different decision thresholds.

detectors, and trail detection methods, together with our method. All metrics are computed against ground-truth masks manually annotated using LabelMe [32].

STARLINC outperforms all baselines on overlap-based metrics (mIoU/Dice), suggesting substantially more spatially coherent masks. Notably, several baselines show non-trivial ROC-AUC/PR-AUC yet still fail to yield meaningful pixel-level overlap, indicating that discrimination metrics alone do not imply accurate localization. For example, Hough attains high precision (0.785) but misses much of the trail extent (recall 0.293), resulting in limited overlap (mIoU 0.142), while Radon and LSD provide little to no usable masks (mIoU 0.000/0.074).

Among learning-based baselines, DeepLSD achieves the highest recall (0.642) but extremely low precision (0.120), consistent with overly dense predictions. While DT-LSD provides a better balance (mIoU 0.240, Dice 0.332) than DeepLSD, it remains far behind STARLINC. The other methods underperform: ASTA yields low scores across metrics, and MaxiMask shows a mismatch between image-level recognition and localization, with high precision (0.838) but low recall (0.159), leading to limited overlap (mIoU 0.155, Dice 0.268). In contrast, STARLINC achieves the best localization (mIoU 0.485, Dice 0.653) with high precision (0.815) and strong recall (0.545), doubling the best baseline mIoU.

In addition, Figure 5 visualizes the ROC and PR curves whose trends match the AUC metrics in Tab. 1. In particular, STARLINC maintains a more favorable trade-off, with higher ROC-AUC (0.918) and PR-AUC (0.865) than all baselines. This agreement further supports that the gains are not confined to a single operating point, but persist across a broad range of decision thresholds.

5.3 Qualitative Analysis of Satellite Trail Segmentation

Qualitative results are shown in Figure 6. The Hough transform produces short, fragmented mask segments, while DT-LSD outputs thicker, more dispersed masks than the trail. MaxiMask produces no visible mask in most cases, consistent with

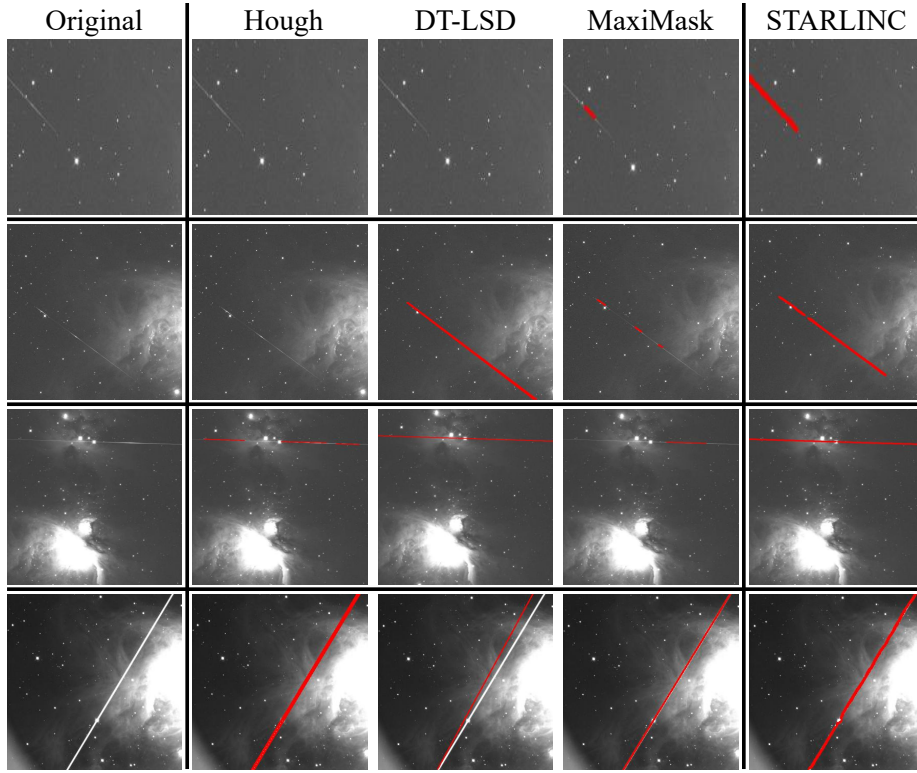


Fig. 6: Qualitative examples of satellite trail masking. Each row shows a different astronomical image containing a satellite trail. From left to right, the original observation, and trail masks predicted by Hough [11], DT-LSD [17], MaxiMask [28], and STARLINC, where detected trail locations are overlaid in red. If the model fails to detect the line, the output remains identical to the original image.

its near-zero quantitative scores. In contrast, STARLINC produces continuous masks that closely follow the true trail geometry. This performance is enabled by the three-channel input design, where the mean differential map and heatmap encode temporal structure, allowing a model trained on synthetic data to generalize effectively to real observations. These results demonstrate that reliable pixel-level segmentation can be achieved by directly extending image-level classification with synthetic supervision, without requiring costly manual annotation of real astronomical survey. A comparison among all models is provided in Sec. D.

5.4 Cross-Dataset Validation

In this section, we evaluate the cross-dataset generalizability of STARLINC. For this evaluation, we use two additional datasets: an NGC-region dataset captured with the 7DT telescope [20] and a publicly available ZTF dataset [14]. The NGC-region dataset contains 4,812 images, and the ZTF dataset contains 4,875 images.

Table 2: Cross-dataset evaluation on ZTF and NGC-region observations.

Dataset Metric		Method								
		Hough	Radon	LSD	DeepLSD [†]	DT-LSD [†]	ASTA [†]	MaxiMask [†]	UnetLSD [†]	STARLINC
ZTF	mIoU	0.166	0.029	0.288	0.321	0.019	0.067	0.365	0.163	0.469
	Dice	0.285	0.057	0.447	0.486	0.038	0.126	0.535	0.280	0.639
NGC	mIoU	0.115	0.001	0.050	0.199	0.011	0.020	0.008	0.120	0.471
	Dice	0.206	0.002	0.095	0.331	0.021	0.040	0.015	0.215	0.640

To construct evaluation ground truth, human annotators manually masked the satellite trails in these datasets. The STARLINC model is trained only on the original Orion Molecular Clouds region 7DT training set and directly applied to both target datasets without target-dataset fine-tuning. For learning-based baselines, we report their fine-tuned settings when applicable, denoted by [†].

As shown in Table 2, STARLINC achieves the best mIoU and Dice on both datasets, demonstrating that STARLINC generalizes well beyond the original Orion Molecular Clouds region. Notably, STARLINC leverages inter-frame differential maps that emphasize frame-specific satellite trails while suppressing persistent celestial structures, allowing the learned representation to transfer effectively across different observational fields and surveys.

6 Analysis

6.1 Ablation Study

We perform an ablation study to analyze the contribution of the key components in the segmentation pipeline. Tab. 3 reports the effect of the inter-frame differential map (§ 4.2) and the pseudo-localization via heatmap (§ 4.3). When both components are removed, the model fails to produce meaningful masks, resulting in zero

Table 3: Ablation results of segmentation task.

Method		Metrics			
§ 4.2	§ 4.3	mIoU	Dice	Prec.	Recall
✗	✗	0.000	0.000	1.000	0.000
✓	✗	0.410	0.582	0.585	0.578
✗	✓	0.198	0.331	0.276	0.411
✓	✓	0.485	0.653	0.815	0.545

segmentation performance. Using only the inter-frame differential map already yields strong segmentation performance, and adding the heatmap further improves the results. Their combination performs best because the differential map suppresses persistent celestial structures while the heatmap guides trail localization. Additional sensitivity analysis is provided in Sec. F.

6.2 Comparison of Inter-Frame Differential Maps

The inter-frame differential map is a central cue in STARLINC, suppressing persistent celestial content while highlighting frame-specific satellite trails. We compare the SSIM-based structural dissimilarity map with RMSE, MAE, and astronomy-oriented subtraction methods, including ZOGY [41], OIS [3], and

Table 4: Differential map comparison.

Method	mIoU
RMSE	0.469 ± 0.014
MAE	0.472 ± 0.002
ZOGY	0.392 ± 0.016
OIS	0.431 ± 0.033
DIA	0.436 ± 0.008
SSIM	0.485 ± 0.012

Table 5: Analysis of Alternative Designs.

Category	Method	mIoU	Dice	Prec.	Recall
Image-only segmentation	SAM [†] [21]	0.358	0.528	0.449	0.640
Synthetic supervision	PSF	0.539	0.700	0.748	0.658
Inter-frame representation	Otsu	0.045	0.004	0.002	0.761
	Morph.	0.158	0.270	0.172	0.630
	Hough	0.259	0.178	0.105	0.741
Localization cue	IRNet [†] [1]	0.131	0.232	0.373	0.169
	SEAM [†] [38]	0.039	0.074	0.106	0.057
Proposed	STARLINC	0.485	0.653	0.815	0.545

DIA [8]. These methods are treated as differential-map alternatives rather than stand-alone segmentation baselines, since they produce difference images rather than binary masks. All variants use the same segmentation architecture and training setup, with only the differential-map construction changed.

As shown in Tab. 4, SSIM performs best among the differential maps. RMSE and MAE are sensitive to brightness fluctuations and residual stellar cores, while ZOGY, OIS, and DIA are mainly optimized for point-source variability and normalized subtraction. Since satellite trails are line-like outliers, SSIM better captures frame-specific structural changes than pixel-wise flux differences. Thus, the SSIM-based map provides a stronger masking cue, while physical subtraction methods may still require tuning for extended artifacts. Qualitative comparisons are provided in Figure 11.

6.3 Analysis of Alternative Designs

The three-channel design of STARLINC naturally suggests several alternatives. Instead of combining the original image I , inter-frame differential map $\bar{M}$, and heatmap H_{GC} , we consider whether each cue can be used alone or replaced by an existing technique. We organize these alternatives by their target cue in Tab. 5.

Image-only segmentation. A natural baseline is to apply a general-purpose segmentation model directly to the astronomical image. We evaluate SAM [21] as an image-only segmentation alternative, without using inter-frame differential maps or classifier heatmaps. As shown in Tab. 5, SAM obtains 0.358 mIoU and 0.528 Dice, which is below STARLINC. This suggests that generic image segmentation alone is insufficient for low-SNR satellite trail masking, where trails are thin and lack the object-like appearance priors of natural images.

Synthetic supervision. We further examine whether improved synthetic-trail realism can enhance segmentation performance. For this purpose, we replace the geometric trail renderer with a PSF-aware renderer while keeping the remaining STARLINC pipeline unchanged. As shown in Tab. 5, PSF-aware rendering improves mIoU and Dice from 0.485/0.653 to 0.539/0.700, suggesting that more realistic synthetic trails can improve mask overlap. However, its lower precision

relative to the main STARLINC configuration (0.748 compared with 0.815) suggests that these gains in overlap metrics are accompanied by reduced precision. We therefore include it as an additional analysis.

Inter-frame representation. We assess whether the inter-frame differential map can serve as a direct basis for mask generation. Specifically, we apply Otsu thresholding, morphological filtering, and Hough-based line extraction to the differential map. Although these methods achieve relatively high recall (Tab. 5), their precision remains very low because residual stellar cores, detector artifacts, and background fluctuations also produce strong responses. This shows that the differential map is effective as a localization cue but requires a learned segmentation model to distinguish satellite trails from other transient structures.

Localization cue. Since the heatmap provides coarse localization of trail-like regions, we examine whether it can be used for weakly supervised segmentation. We evaluate IRNet [1] and SEAM [38], which derive masks from classifier localization cues. Both perform substantially worse than STARLINC. Their activation maps typically capture only the most discriminative trail fragments, not the full thin trail extent. Instead, STARLINC uses the heatmap only as an auxiliary cue while obtaining pixel-level supervision from synthetic trail masks.

7 Conclusion

Satellite trail contamination is becoming a persistent problem for modern astronomical surveys as low Earth orbit constellations expand. We presented STARLINC, which leverages temporal redundancy across repeated observations to isolate transient satellite artifacts. Using inter-frame differential maps, synthetic trail generation, and activation-map-based pseudo-localization, STARLINC enables pixel-level segmentation without manually annotated masks from real survey data. Experiments on the 7DT dataset show clear gains over classical and learning-based baselines, and cross-dataset evaluations further demonstrate that STARLINC maintains strong performance across survey datasets, indicating strong generalizability. By selectively masking contaminated regions instead of discarding entire frames, STARLINC offers a practical and scalable approach for handling satellite-induced artifacts in large-scale astronomical surveys.

Acknowledgements

This work has been mainly supported by the Basic Research Laboratory Program through the National Research Foundation of Korea (NRF) funded by the Korea government (MSIT) (RS-2024-00416859). Jinho Lee is the corresponding author, and he is also funded by: RS-2024-00395134, RS-2024-00347394, RS-2026-25548502, RS-2026-25495605, RS-2025-00564840, RS-2023-00256081, RS-2026-25549926

References

1. Ahn, J., Cho, S., Kwak, S.: Weakly supervised learning of instance segmentation with inter-pixel relations. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2209–2218 (2019)
2. Akinlar, C., Topal, C.: Edlines: A real-time line segment detector with a false detection control. *Pattern Recognition Letters* **32**(13), 1633–1642 (2011)
3. Alard, C., Lupton, R.H.: A method for optimal image subtraction. *The Astrophysical Journal* **503**(1), 325–331 (1998)
4. Ali, H., Lampert, C.H., Breuel, T.M.: Satellite tracks removal in astronomical images. In: Iberoamerican Congress on Pattern Recognition. pp. 892–901. Springer (2006)
5. Bassa, C., Hainaut, O., Galadí-Enríquez, D.: Analytical simulations of the effect of satellite constellations on optical and near-infrared observations. *Astronomy & Astrophysics* **657**, A75 (2022)
6. Borlaff, A.S., Marcum, P.M., Howell, S.B.: Satellite megaconstellations will threaten space-based astronomy. *Nature* (2025). <https://doi.org/10.1038/s41586-025-09759-5>, published online December 2025
7. Borlaff, A.S., Marcum, P.M., Howell, S.B., Sánchez-Alarcón, P.M., Dubois, D., McDowell, J.: Spherex confirms predictions for artificial satellite trail pollution in low earth orbit. arXiv preprint arXiv:2605.27501 (2026)
8. Bramich, D.: A new algorithm for difference image analysis. *Monthly Notices of the Royal Astronomical Society: Letters* **386**(1), L77–L81 (2008)
9. Canny, J.: A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **8**(6), 679–698 (1986)
10. Chen, X., Gu, H., Wang, C., Mu, H., Zheng, J., Du, J., Ren, J., Fan, Z., Li, J.: Artificial satellite trails detection using u-net deep neural network and line segment detector algorithm. *Publications of the Astronomical Society of the Pacific* **137**(9), 094504 (2025)
11. Duda, R.O., Hart, P.E.: Use of the hough transformation to detect lines and curves in pictures. *Communications of the ACM* **15**(1), 11–15 (1972)
12. Duev, D.A., Mahabal, A., Ye, Q., Tirumala, K., Belicki, J., Dekany, R., Frederick, S., Graham, M.J., Laher, R.R., Masci, F.J., et al.: Deepstreaks: identifying fast-moving objects in the zwicky transient facility data with deep learning. *Monthly Notices of the Royal Astronomical Society* **486**(3), 4158–4165 (2019)
13. von Gioi, R.G., Jakubowicz, J., Morel, J.M., Randall, G.: Lsd: A fast line segment detector with a false detection control. *IEEE Transactions on Pattern Analysis & Machine Intelligence* **32**(04), 722–732 (2010)
14. Graham, M.J., Kulkarni, S., Bellm, E.C., Adams, S.M., Barbarino, C., Blagorodnova, N., Bodewits, D., Bolin, B., Brady, P.R., Cenko, S.B., et al.: The zwicky transient facility: science objectives. *Publications of the Astronomical Society of the Pacific* **131**(1001), 078001 (2019)
15. Hainaut, O.R., Williams, A.P.: Impact of satellite constellations on astronomical observations with eso telescopes in the visible and infrared domains. *Astronomy & Astrophysics* **636**, A121 (2020)
16. Huang, K., Wang, Y., Zhou, Z., Ding, T., Gao, S., Ma, Y.: Learning to parse wireframes in images of man-made environments. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 626–635 (2018)
17. Janampa, S., Pattichis, M.: Dt-lsd: Deformable transformer-based line segment detection. In: 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 3477–3486. IEEE (2025)

18. Jeffries, C., Acuña, R.: Detection of streaks in astronomical images using machine learning. *Journal of Artificial Intelligence and Technology* **4**(1), 1–8 (2024)
19. Kim, D.W.: Astride: Automated streak detection for astronomical images. *Astrophysics Source Code Library* pp. ascl-1605 (2016)
20. Kim, J.H., Im, M., Lee, H., Chang, S.W., Choi, H., Paek, G.S.: Introduction to the 7-dimensional telescope: commissioning procedures and data characteristics. In: *Ground-based and Airborne Telescopes X*. vol. 13094, pp. 442–452. SPIE (2024)
21. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 4015–4026 (2023)
22. McDowell, J.C.: The low earth orbit satellite population and impacts of the spacex starlink constellation. *The Astrophysical Journal Letters* **892**(2), L36 (2020)
23. McDowell, J.C.: Starlink statistics. Jonathan’s Space Pages <https://planet4589.org/space/con/star/stats.html> (2026), accessed: 2026-02-18
24. Moon, Y., Lee, J., Mun, D., Lim, S.: Deep learning-based method to recognize line objects and flow arrows from image-format piping and instrumentation diagrams for digitization. *Applied sciences* **11**(21), 10054 (2021)
25. Mróz, P., Otarola, A., Prince, T.A., Dekany, R., Duev, D.A., Graham, M.J., Groom, S.L., Masci, F.J., Medford, M.S.: Impact of the spacex starlink satellites on the zwicky transient facility survey observations. *The Astrophysical Journal Letters* **924**(2), L30 (2022)
26. Nir, G., Zackay, B., Ofek, E.O.: Optimal and efficient streak detection in astronomical images. *The Astronomical Journal* **156**(5), 229 (2018)
27. Omat, C., Birlan, M., Nedelcu, D.A., Anghel, S.: Deep learning approach for automated detection of space objects in astronomical imaging. *Astronomy and Computing* **55**, 101081 (2026)
28. Paillassa, M., Bertin, E., Bouy, H.: Maximask and maxitrack: Two new tools for identifying contaminants in astronomical images using convolutional neural networks. *Astronomy & Astrophysics* **634**, A48 (2020)
29. Pautrat, R., Barath, D., Larsson, V., Oswald, M.R., Pollefeys, M.: Deeplsd: Line segment detection and refinement with deep image gradients. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 17327–17336 (2023)
30. Radon, J.: On the determination of functions from their integral values along certain manifolds. *IEEE transactions on medical imaging* **5**(4), 170–176 (1986)
31. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical image computing and computer-assisted intervention*. pp. 234–241. Springer (2015)
32. Russell, B.C., Torralba, A., Murphy, K.P., Freeman, W.T.: Labelme: A database and web-based tool for image annotation. *International Journal of Computer Vision* (2008)
33. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: *Proceedings of the IEEE international conference on computer vision*. pp. 618–626 (2017)
34. Stoppa, F., Groot, P., Stuik, R., Vreeswijk, P., Bloemen, S., Pieterse, D., Woudt, P.: Automated detection of satellite trails in ground-based observations using u-net and hough transform. *Astronomy & Astrophysics* **692**, A199 (2024)
35. Suárez, I., Buenaposada, J.M., Baumela, L.: Elsed: Enhanced line segment drawing. *Pattern Recognition* **127**, 108619 (2022)

36. Teplyakov, L., Kaymakov, K., Shvets, E., Nikolaev, D.: Line detection via a lightweight cnn with a hough layer. In: Thirteenth International Conference on Machine Vision. vol. 11605, pp. 376–385. SPIE (2021)
37. Topal, C., Akinlar, C.: Edge drawing: a combined real-time edge and segment detector. *Journal of Visual Communication and Image Representation* **23**(6), 862–872 (2012)
38. Wang, Y., Zhang, J., Kan, M., Shan, S., Chen, X.: Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12275–12284 (2020)
39. Wenger, M., Ochsenbein, F., Egret, D., Dubois, P., Bonnarel, F., Borde, S., Genova, F., Jasiewicz, G., Laloë, S., Lesteven, S., et al.: The simbad astronomical database—the cds reference database for astronomical objects. *Astronomy and Astrophysics Supplement Series* **143**(1), 9–22 (2000)
40. Xue, N., Bai, S., Wang, F., Xia, G.S., Wu, T., Zhang, L.: Learning attraction field representation for robust line segment detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1595–1603 (2019)
41. Zackay, B., Ofek, E.O., Gal-Yam, A.: Proper image subtraction—optimal transient detection, photometry, and hypothesis testing. *The Astrophysical Journal* **830**(1), 27 (2016)
42. Zhang, Y., Wei, D., Li, Y.: Ag3line: Active grouping and geometry-gradient combined validation for fast line segment extraction. *Pattern Recognition* **113**, 107834 (2021)
43. Zhang, Z., Li, Z., Bi, N., Zheng, J., Wang, J., Huang, K., Luo, W., Xu, Y., Gao, S.: Ppgnet: Learning point-pair graph for line segment detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7105–7114 (2019)
44. Zhao, K., Han, Q., Zhang, C.B., Xu, J., Cheng, M.M.: Deep hough transform for semantic line detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(9), 4793–4806 (2021)