

DTTNet: Improving Video Shadow Detection via Dark-Aware Guidance and Tokenized Temporal Modeling

Zhicheng Li^{1,2*}, Kunyang Sun^{1,2*}, Rui Yao^{1,2†}, Hancheng Zhu^{1,2}, Fuyuan Hu³, Jiaqi Zhao^{1,2}, Zhiwen Shao^{1,2}, Yong Zhou^{1,2}

¹School of Computer Science and Technology / School of Artificial Intelligence, China University of Mining and Technology

²Mine Digitization Engineering Research Center of the Ministry of Education, China

³School of Electronic and Information Engineering, Suzhou University of Science and Technology, China

Abstract

Video shadow detection confronts two entwined difficulties: distinguishing shadows from complex backgrounds and modeling dynamic shadow deformations under varying illumination. To address shadow-background ambiguity, we leverage linguistic priors through the proposed Vision-language Match Module (VMM) and a Dark-aware Semantic Block (DSB), extracting text-guided features to explicitly differentiate shadows from dark objects. Furthermore, we introduce adaptive mask reweighting to downweight penumbra regions during training and apply edge masks at the final decoder stage for better supervision. For temporal modeling of variable shadow shapes, we propose a Tokenized Temporal Block (TTB) that decouples spatiotemporal learning. TTB summarizes cross-frame shadow semantics into learnable temporal tokens, enabling efficient sequence encoding with minimal computation overhead. Comprehensive Experiments on multiple benchmark datasets demonstrate state-of-the-art accuracy and real-time inference efficiency.

Code — <https://github.com/city-cheng/DTTNet>

Introduction

Shadows, ubiquitous in natural imagery and video sequences, offer critical cues for vision tasks including light source conditions (Lalonde and Matthews 2014), object shapes (Karsch et al. 2011; Shao, Taff, and Walsh 2011), and depth relationships (Adams et al. 2022). Conversely, ambiguous shadows can degrade performance in object detection (Cucchiara et al. 2003; Hu et al. 2021), illumination estimation (Adams et al. 2021), and visual tracking (Chen et al. 2021). While image shadow detection focuses on spatial separation under complex lighting, video shadow detection additionally requires modeling temporal shadow deformations induced by dynamic illumination. This necessity for spatiotemporal dual modeling renders video shadow detection a substantially more challenging problem.

Recent advances in deep learning have significantly propelled video shadow detection. Unlike video object segmentation where semantics vary substantially, shadows maintain

*These authors contributed equally.

†Corresponding authors, ruiyao@cumt.edu.cn.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

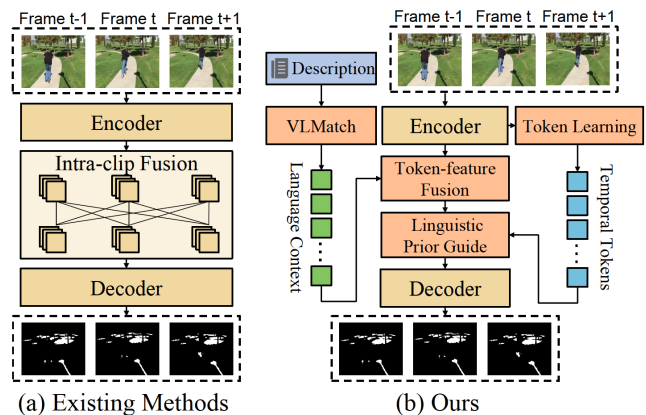


Figure 1: Comparison between existing methods and ours. Despite using intra-clip fusion, we utilize token to learn temporal characteristics and introduce textual priors to guide the model to learn from dark regions.

consistent characteristics—persistently occupying darker regions. Hence, one promising direction exploits inter-frame affinity to aggregate locally similar shadow features across video sequences. For instance, previous methods (Chen et al. 2021; Liu et al. 2023a) utilize intra-frame attention to aggregate temporally consistent features. However, frame-level feature fusion learns from predefined frame pairs, which not only produces redundant representations but also leads to fragmented temporal modeling and compromised representation coherence. To efficiently encoding temporal information, we are motivated to optimize learnable tokens to capture shadow dynamics and subsequently inject them into spatial features. This token-level fusion not only enhances the aggregation of multi-frame information but also effectively reduces computational overhead and parameter requirements. Consequently, we propose a Tokenized Temporal Block (TTB) that first encodes temporal contexts into compact tokens and then selectively transfers this knowledge to spatial pixels via a spatial matching mechanism.

In terms of spatial modeling, a mainstream shadow detection paradigm exploits shadows’ inherent low-luminance property by identifying dark regions. To achieve finer-grained shadow modeling, SSTINet (Wei et al. 2024) de-

composes shadow representations into structural and detailed components, serving targeted decoding of shadow regions. However, the lack of pixel-level annotations for dark regions confines vision-only shadow modeling to implicit representation learning, hindering explicit localization. To address this, we introduce textual priors to explicitly guide the model’s attention towards darkness as shown in Fig. 1. Specifically, to mitigate interference from non-shadow regions within dark areas, we propose distinct text descriptions for shadows and dark regions. Leveraging CLIP’s powerful zero-shot capabilities, our proposed Vision-language Match Module (VMM) then performs attention-driven cross-modal matching between these linguistic priors and image features, enhancing focus on both dark regions and shadows. Furthermore, we design the Dark-aware Semantic Block (DSB) to adaptively weight dark and shadow features.

While supervising each DSB output with the shadow mask, the inherent ambiguity of shadow edges (penumbra regions) presents a challenge: excessive focus on blurred boundaries during early training impedes learning of the primary shadow body. Therefore, unlike SSTINet (Wei et al. 2024), we apply a reweighting scheme to the supervision mask during training, assigning lower weights to penumbra regions. Additionally, edge masks are employed solely at the final decoder output, ensuring a balanced focus on both the main shadow structure and fine-grained details. Our main contributions are as follows:

- We introduce a novel framework, the Dark-aware and Temporal Tokenized Network (DTTNet), which integrates dark-aware linguistic guidance with tokenized temporal modeling to effectively capture spatial and temporal dependencies, significantly advancing the state-of-the-art in video shadow detection.
- We introduce a Tokenized Temporal Block designed to model temporal characteristics at the token level. By decoupling temporal modeling from spatial features, this module effectively captures shadow dynamics for efficient learning of coherent temporal properties.
- To resolve visual ambiguities, we leverage textual priors and introduce a Vision-language Match Module (VMM) along with a Dark-aware Semantic Block (DSB), specifically designed for shadow detection in dark regions. Furthermore, we explicitly decouple the penumbra region mask from the main shadow body to facilitate more effective model learning.

Related Work

Image Shadow Detection

Image shadow detection has advanced through diverse innovations (Khan et al. 2014; Hu et al. 2021; Wang, Li, and Yang 2018; Zhu et al. 2021). Some research use various fusion modules to enhance the ability of the encoder. For example, BDRAR (Zhu et al. 2018) combines global and local context via a bidirectional pyramidal architecture and DSD (Zheng et al. 2019) learns distraction-aware features to reduce false positives through explicit modeling of ambiguous regions. Others try to model shadow and background

separately. SDDNet (Cong et al. 2023) decomposes shadow and background features with style-guided disentanglement to mitigate background color interference. Recently, (Guan, Xu, and Lau 2024) propose a dark-region recommendation module to enhance discrimination in low-intensity regions. Although effective on static images, these methods lack mechanisms to capture temporal dependencies, making them suboptimal for videos where shadows deform dynamically under varying illumination.

Video Shadow Detection

Research on video shadow detection has gained momentum. (Chen et al. 2021) pioneered this field by constructing the first dedicated dataset and proposing a triple-flow network. Subsequent advancements have further expanded the landscape: SC-Cor (Ding et al. 2022) enhances cross-frame feature coherence for shadow regions without relying on pixel-wise labels. SCOTCH&SODA (Liu et al. 2023a) introduced transformers into the task, leveraging trajectory attention to handle deformations and a contrastive loss to learn unified shadow representations. DAS (Wang et al. 2023) employed SAM (Kirillov et al. 2023) and designed a long short-term network for subsequent frame predictions. Recently, (Duan et al. 2024) proposed a two-stage paradigm and the CVSD dataset to bring the task to complex scenarios. SSTINet (Wei et al. 2024) deployed a structure-aware module focusing on edge-region distance relations. TBGDiff (Zhou et al. 2024) uses diffusion model with the guidance of boundary and long-term frames to generate shadow masks. Most existing methods overlook shadows’ intrinsic properties, directly extracting regions from features, whereas our approach introduces linguistic priors to guide dark-region information extraction and uses tokens for temporal correlation learning, moving beyond multi-frame feature fusion.

Methodology

Overview

As shown in Fig. 2, the proposed DTTNet leverages text priors for video shadow detection. For an input video clip V , a large multimodal model processes its first frame to generate descriptive texts for shadows T_s and dark regions T_d via a predefined template. The tuple (V, T_s, T_d) is fed into the pre-trained CLIP model for feature extraction. Then the proposed VMM performs text-image matching, yielding visual-informed context for shadows and dark regions. Critically, at each encoder stage, this context is injected into our DSB to enhance encoder features with text priors, supervised by penumbra-aware mask for accuracy. Concurrently, a TTB within each encoder stage employs learnable tokens to summarize temporal information from preceding features. This tokenized temporal feature is fused via attention before feature propagation. The refined encoder features are finally decoded to output shadow masks M for all frames.

Vision-language Match Module

Shadow regions often overlap with dark areas, causing ambiguity under complex lighting conditions when relying solely on visual cues. To address this challenge, we leverage textual

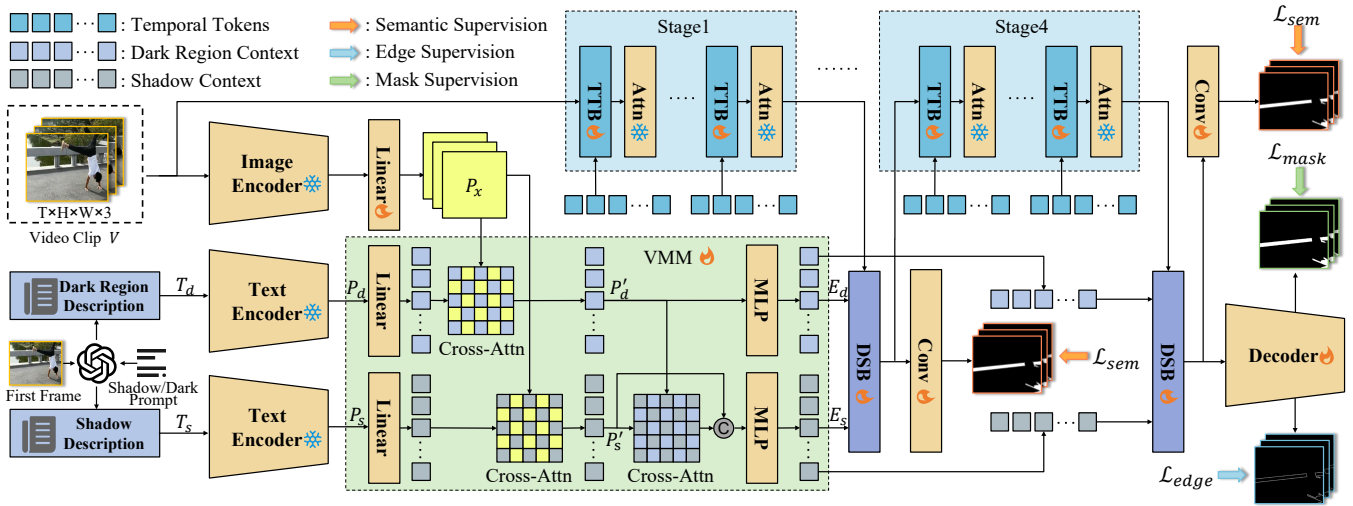


Figure 2: Architecture of Dark-aware and Temporal Tokenized Network (DTTNet). DTTNet integrates dark-aware linguistic guidance with tokenized temporal modeling to effectively capture spatial and temporal dependencies. It consists of three novel modules: the Vision-language Match Module (VMM), Dark-aware Semantic Block and the Tokenized Temporal Block. We freeze most of the parameters and update only the parameter of decoder and proposed modules.

priors to identify shadows and dark regions. For each video clip V , a vision-language large model is employed to generate structured descriptions— T_s for shadows and T_d for dark regions—based on a predefined template. The prompt used for generating these descriptions is: “Describe the OBJECT, COUNT, POSITION, and DETAIL of all the shadow/dark areas in the image. When answering COUNT, do not use numbers just say single or multiple. Your answer must be like this: COUNT shadows/dark regions, shadow/dark region of OBJECT, POSITION, DETAIL.”

Exploiting CLIP’s zero-shot power, we freeze the model parameters of CLIP and further feed both T_s, T_d and $V \in \mathbb{R}^{T \times H \times W \times 3}$ into CLIP to generate corresponding text features $P_s \in \mathbb{R}^{L_s \times C_l}, P_d \in \mathbb{R}^{L_d \times C_l}$ and image feature $P_x \in \mathbb{R}^{T \times M \times C_m}$, where L_s, L_d means length of shadow and dark region descriptions and C_l, C_m refers to channel number of CLIP’s text and image embedding. Our Vision-language Match Module then aligns these modalities, generating visual-informed contexts $E_s \in \mathbb{R}^{L_s \times C_e}$ and $E_d \in \mathbb{R}^{L_d \times C_e}$ for shadow and dark region respectively, where C_e means the embedding channel number of DTTNet. We project both modalities to a shared channel dimension via fully connected layers. For dark region context, cross-attention with P_d as query and P_x as key/value yields P'_d :

$$P'_d = \text{Attn}(P_d, P_x), \quad (1)$$

P'_d is then processed by a multi-layer perceptron (MLP) to yield the refined dark region context E_d .

$$E_d = \text{MLP}(P'_d), \quad (2)$$

For shadow region context, exploiting dark regions’ cues, we first apply cross-attention (P_s as query, P_x as key/value) for initial features P'_s . Then P'_s is fused with P'_d via another cross-attention. The concatenated outputs of both

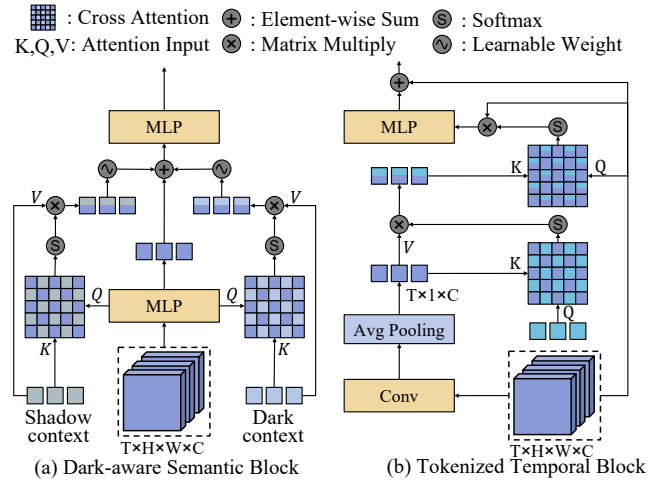


Figure 3: Details of the proposed Dark-aware Semantic Block and Tokenized Temporal Block.

cross-attention operations are fed into a MLP to generate the final shadow context E_s :

$$E_s = \text{MLP}(\text{Cat}(\text{Attn}(P_s, P_x), \text{Attn}(P'_s, P'_d))). \quad (3)$$

Here $\text{Attn}(\cdot)$ denotes the standard attention mechanism, $\text{Cat}(\cdot)$ represents concatenation along the channel dimension, and the module $\text{MLP}(\cdot)$ consists of two sequential linear layers that first compress and then restore the channel dimensionality of the input feature representations.

Dark-aware Semantic Block

Compared to the extracted image features, text features contain more accurate high-level semantics. Therefore, we leverage shadow/dark region context (*i.e.* E_s and E_d) to

perform semantic adjustment of the image features, rather than pixel-level classification (which is better suited for the decoder). We introduce a DSB before feeding the outputs of each encoder stage to the decoder. The DSB enhances the semantic expressiveness of the features. Furthermore, penumbra-aware supervision is imposed on the DSB outputs to ensure semantic accuracy.

Specifically, as shown in Fig. 3(a), the encoder feature from the i -th stage, denoted as $X_i \in \mathbb{R}^{T \times \frac{H}{16} \times \frac{W}{16} \times C_b}$, undergoes processing to integrate shadow (E_s) and dark region (E_d) contexts. First, a fully connected layer compresses the channel dimension of X_i to C_e to align with the channel dimensionality of the context embeddings. Subsequently, cross-attention mechanisms individually inject E_s and E_d into the compressed feature, yielding shadow-enhanced feature X_i^s and dark-region-enhanced feature X_i^d :

$$\begin{aligned} X_i^s &= \text{Attn}(\text{Linear}(X_i), E_s), \\ X_i^d &= \text{Attn}(\text{Linear}(X_i), E_d), \end{aligned} \quad (4)$$

Here $\text{Linear}(\cdot)$ denotes the linear layer. These two enhanced features are then dynamically fused using learned weighting parameters (*i.e.* α and β). The fused representation is passed through a 1×1 convolutional layer to generate a soft shadow mask \tilde{M}_i , which is supervised by the processed ground-truth shadow annotation \tilde{M} . Concurrently, the fused representation is directly added to the original encoder feature X_i , resulting in the final enhanced feature output X_i^f :

$$X_i^f = X_i + \alpha * X_i^s + \beta * X_i^d, \quad (5)$$

where α and β are learnable weights that allow the model to adapt the importance between shadow and dark regions.

Tokenized Temporal Block

Existing approaches primarily encode temporal information through frame-level feature fusion, neglecting learning coherent temporal representations. To address this, we introduce the Tokenized Temporal Block (TTB). This module employs a collection of learnable tokens to acquire a universal temporal representation applicable across all shadow videos after training. This representation enables more efficient extraction of consistent shadow features across consecutive frames. As shown in Fig. 2, we introduce a TTB before each layer of the encoder to capture temporal information prior to feature processing. Fig. 3(b) shows the structure of TTB. For the encoder feature X_{j-1} output from the $(j-1)$ -th layer, its channel dimension is first aligned via a 1×1 convolutional layer. Subsequently, spatial information is compressed using average pooling, yielding a temporal feature $Z_j \in \mathbb{R}^{T \times 1 \times C_e}$:

$$Z_j = \text{AvgPool}(\text{Conv}(X_{j-1})), \quad (6)$$

where $\text{Conv}(\cdot)$ refers to the 1×1 convolution. A set of learnable tokens $K_j \in \mathbb{R}^{L_k \times C_e}$ is then used as queries, with Z_j serving as keys and values in a cross-attention operation, enabling the tokens to assimilate relevant temporal information from the temporal sequence:

$$K_j = \text{Attn}(K_j, Z_j). \quad (7)$$

Following this, the original feature X_{j-1} is spatially flattened to $N \times T \times C_e$ and undergoes cross-attention with the learned tokens. This process allows each frame to effectively incorporate information from the tokens while preserving its original spatial structure.

Finally, a 1×1 convolutional layer is used to restore the channel dimension, preparing the feature X_j for processing in the j -th encoder layer:

$$X_j = \text{Conv}(\text{Attn}(K_j, X_{j-1})). \quad (8)$$

Decoder

For features extracted at each encoder stage, we first apply a MLP to model channel-wise dependencies. The resulting features from all stages are then integrated through a convolutional block comprising sequential convolutional, batch normalization, ReLU activation, and convolutional layer. As outlined in Fig. 2, we design a decoder that gradually upsamples features from the compressed feature size of $\frac{H}{16} \times \frac{W}{16}$ to $\frac{H}{4} \times \frac{W}{4}$. To mitigate information loss and noise amplification during upsampling, we perform convolutions before interpolations and employ DySample (Liu et al. 2023b) as the upsampling method. After two successive upsampling stages, the decoder employs an additional convolutional layer to produce preliminary mask features. Subsequently, two parallel 1×1 convolutional layers generate the final shadow mask prediction \tilde{M}_s and edge mask prediction \tilde{M}_e from the above mask features.

Loss Function

Our loss function consists of three components: penumbra-aware semantic loss \mathcal{L}_{sem} , shadow edge loss \mathcal{L}_{edge} , and shadow mask loss \mathcal{L}_{mask} . Given that encoder features at low resolutions prioritize semantic accuracy over pixel-level classification fidelity, we formulate the semantic loss as a regression task using mean absolute error to supervise the features of the DSB. To address ambiguity in shadow boundaries (penumbra regions), we attenuate edge values in proportion to their distance from the center while preserving values of main regions. This strategy focuses the DSB on capturing definitive shadow structures while mitigating interference from uncertain boundaries, implemented as:

$$\hat{M}(u, v) = \begin{cases} 1, & \text{if } \text{Eros}(M(u, v)) > 0 \\ \text{Dist}(M(u, v)), & \text{else} \end{cases} \quad (9)$$

$$\mathcal{L}_{sem} = \sum_i L_{mse}(\tilde{M}_i, \hat{M}), \quad (10)$$

where (u, v) denotes the location on the reprocessed mask \hat{M} and L_{mse} refers to the Mean Squared Error loss. \tilde{M}_i is the auxiliary output of DSB in the i -th stage. $\text{Eros}(\cdot)$ refers to the morphological erosion operation with a kernel size of 3. $\text{Dist}(\cdot)$ represents the distance transform function. To supervise the final shadow mask, we employ a dual-branch supervision strategy where the original shadow mask and extracted edge components serve as distinct supervisory signals for the decoder's output features. For edge components, we calculate edge loss as follows:

$$M_e = \mathbb{I}(1 - \hat{M}), \quad (11)$$

Methods			Evaluation Metrics					
Tasks	Techniques	Year	MAE ↓	F _β ↑	IoU ↑	BER ↓	S-BER ↓	N-BER ↓
IOS	FPN (Lin et al. 2017)	2017	0.044	0.707	0.512	19.49	36.59	2.40
	R ³ Net (Deng et al. 2018)	2018	0.043	0.710	0.502	20.40	37.37	3.55
	Segformer (Xie et al. 2021)	2021	0.030	0.773	0.601	11.56	21.39	1.73
	DDP (Ji et al. 2023)	2023	0.038	0.771	0.608	10.74	18.90	2.57
VOS	STM (Oh et al. 2019)	2019	0.069	0.598	0.408	25.69	47.44	3.95
	COSNet (Lu et al. 2019)	2019	0.040	0.706	0.515	20.51	39.22	1.79
	FEELVOS (Voigtlaender et al. 2019)	2019	0.043	0.710	0.512	19.76	37.27	2.26
	STCN (Cheng, Tai, and Tang 2021)	2021	0.048	0.684	0.528	12.42	21.36	3.48
	Pix2Seq (Chen et al. 2023)	2023	0.034	0.775	0.618	10.63	19.13	2.14
ISD	BDRAR (Zhu et al. 2018)	2018	0.050	0.695	0.484	21.30	40.28	2.32
	DSD (Zheng et al. 2019)	2019	0.044	0.702	0.519	19.89	37.88	1.89
	MTMT (Chen et al. 2020)	2020	0.043	0.729	0.517	20.29	38.71	1.86
	FSD (Hu et al. 2021)	2021	0.057	0.671	0.486	20.57	38.06	3.06
	SDDNet (Cong et al. 2023)	2023	0.040	0.754	0.548	14.05	26.10	1.61
	SILT (Yang et al. 2023)	2023	0.031	0.796	0.606	12.80	24.29	1.29
VSD	TVSD (Chen et al. 2021)	2021	0.033	0.757	0.565	17.70	33.96	1.44
	STICT (Lu et al. 2022)	2022	0.046	0.702	0.545	16.60	29.58	3.59
	SC-Cor (Ding et al. 2022)	2022	0.042	0.762	0.615	13.61	24.31	2.91
	SCOTCH & SODA (Liu et al. 2023a)	2023	0.029	0.793	0.640	9.06	16.26	1.44
	DAS (Wang et al. 2023)	2023	0.034	0.754	0.575	12.58	23.60	1.57
	TBGDiff (Zhou et al. 2024)	2024	0.023	0.797	0.667	8.58	16.00	1.15
	TSVSD (Duan et al. 2024)	2024	0.027	0.801	0.684	8.96	-	-
	SSTINet (Wei et al. 2024)	2024	0.017	0.866	0.746	6.48	<u>12.32</u>	<u>0.65</u>
	DTTNet (Ours)	-	0.016	<u>0.849</u>	<u>0.718</u>	6.45	12.29	0.61

Table 1: Quantitative comparisons between our proposed method and SOTA methods on the ViSha (Chen et al. 2021) dataset. ”↑” denotes the higher the value, the better the performance, and ”↓” means the lower the value, the better the performance. We compare with recent methods from Image Object Segmentation (IOS), Image Shadow Detection (ISD), Video Object Segmentation (VOS), and Video Shadow Detection (VSD). The best values are highlighted in **bold**, while the second best values are underlined.

$$\mathcal{L}_{edge} = L_{bce}(\tilde{M}_e, M_e) + L_{dice}(\tilde{M}_e, M_e), \quad (12)$$

where M_e is defined as the ground truth edge mask and $\mathbb{I}(\cdot)$ means instruction function. Meanwhile, we calculate mask loss for shadow mask:

$$\mathcal{L}_{mask} = L_{bce}(\tilde{M}_s, M) + L_{dice}(\tilde{M}_s, M), \quad (13)$$

where L_{bce} means Binary Cross-Entropy with Logits loss and L_{dice} refers to the Dice loss (Milletari, Navab, and Ahmadi 2016).

The final loss function is composed of the weighted sum of these three losses:

$$\mathcal{L} = \lambda_1 * \mathcal{L}_{sem} + \lambda_2 * \mathcal{L}_{edge} + \lambda_3 * \mathcal{L}_{mask}, \quad (14)$$

λ_1 , λ_2 and λ_3 are hyper-parameters that balance the losses.

Experiments

Dataset and Evaluation Metrics

Datasets. The proposed DTTNet is validated on the Visha (Chen et al. 2021) and CVSD (Duan et al. 2024) datasets. Following previous studies (Wang et al. 2023; Ding et al. 2022; Liu et al. 2023a), we employ the Visha dataset to evaluate the model’s performance. Visha comprises 120 videos with diverse content and varying lengths, with more than half of the clips derived from standard video tracking

benchmarks. In addition, we conduct experiments on the more challenging CVSD dataset (Duan et al. 2024), which was recently introduced and includes 196 video clips and a total of 19,757 frames featuring complex shadow patterns.

Evaluation Metrics. To facilitate a fair and thorough performance comparison, we adopt the methodology established in prior research and compute six distinct evaluation metrics. These include Mean Absolute Error (MAE), F-measure, and Intersection over Union (IoU), as well as Balanced Error Rate (BER). Additionally, we incorporate the S-BER score, which is specifically designed for shadow regions, and the N-BER score tailored to non-shadow regions.

Implementation Details

The proposed model is implemented using the MMSegmentation codebase (Contributors 2020). The backbone network of DTTNet is the pre-trained DINOv2 (Oquab et al. 2023). All parameters of the backbone are kept frozen, while the trainable parameters are exclusively derived from four new components: the Vision-language Match Module, Tokenized Temporal Block, Dark-aware Semantic Block, and the decoder. For optimizing these parameters, the AdamW is employed, configured with the learning rate of $5 * 10^{-5}$ and the weight decay of 0.01. Training is conducted with a batch size of 2 and each batch includes 5 frames of the video. We set

Methods		Evaluation Metrics			
Techniques	Year	MAE ↓	F_β ↑	IoU ↑	BER ↓
TVSD	2021	0.099	0.539	0.369	27.28
STICT	2022	0.073	0.608	0.447	23.27
SC-Cor	2022	0.070	0.573	0.476	19.94
SCOTCH & SODA	2023	0.082	0.585	0.426	23.27
DAS	2023	0.087	0.561	0.435	19.15
TBGDiff	2024	0.057	<u>0.663</u>	0.445	23.01
TSVSD	2024	<u>0.046</u>	0.638	<u>0.515</u>	18.32
DTTNet (Ours)	-	0.042	0.766	0.548	23.32

Table 2: Quantitative comparisons between our proposed method and SOTA methods on CVSD (Duan et al. 2024), a more complex dataset.

λ_1 , λ_2 , and λ_3 to 1, 0.5, and 1, respectively. All experiments are carried out at a resolution of 512×512, and Random Horizontal Flip is adopted in the training phase.

Comparison with State-of-the-art Methods

Compared Methods. We compare our network with 21 methods across related tasks, including IOS methods: FPN (Lin et al. 2017), R^3 Net (Deng et al. 2018), Segformer (Xie et al. 2021), DDP (Ji et al. 2023); VOS methods: STM (Oh et al. 2019), COSNet (Lu et al. 2019), FEELVOS (Voigtlaender et al. 2019), STCN (Cheng, Tai, and Tang 2021), Pix2Seq (Chen et al. 2023); ISD methods BDRAR (Zhu et al. 2018), DSD (Zheng et al. 2019), MTMT (Chen et al. 2020), FSD (Hu et al. 2021), SDDNet (Cong et al. 2023), SILT (Yang et al. 2023); and video shadow detection (VSD) TVSD (Chen et al. 2021), STICT (Lu et al. 2022), SC-Cor (Ding et al. 2022), SCOTCH & SODA (Liu et al. 2023a), DAS (Wang et al. 2023), TBGDiff (Zhou et al. 2024), TSVSD (Duan et al. 2024) and SSTINet (Wei et al. 2024).

Quantitative Comparisons. Quantitative results for our approach alongside other methods are presented in Table 1. Since ISD methods are purpose-built for shadow detection, they outperform VOS and IOS networks in this task. Nevertheless, these ISD methods lack the integration of temporal information, resulting in inferior performance compared to VSD techniques. Among all existing approaches, SSTINet (Wei et al. 2024) delivers the strongest results, yet our proposed method exceeds current state-of-the-art techniques in terms of MAE, BER, S-BER, and N-BER metrics. Specifically, our method reduces the MAE from 0.017 to 0.016, lowers the BER from 6.48 to 6.41, improves the S-BER from 12.32 to 12.29, and decreases the N-BER from 0.65 to 0.53. Additionally, on the CVSD dataset (Duan et al. 2024)—which has not been included the evaluation scope of most existing studies—our approach attains the top performance in MAE, F_β , and IoU metrics, with detailed results provided in Table 2.

Qualitative Comparisons. We present in Fig 4 the shadow masks generated by DTTNet and state-of-the-art methods. The first two rows illustrate scenarios with interference from other dark regions: we observe that existing

Methods	Params(M)	FPS	IoU↑	BER↓
SCOTCH & SODA	211.8	11.4	0.640	9.07
DAS	101.3	12.1	0.667	8.58
TBGDiff	<u>102.3</u>	<u>13.5</u>	0.640	9.07
SSTINet	338.3	6.06	0.746	6.48
DTTNet (Ours)	46.6 (502.1)	26.63	<u>0.714</u>	6.45

Table 3: Parameters and efficiency on ViSha (Chen et al. 2021) dataset. Since we freeze most of the parameters, the total number of parameters is provided in parentheses.

methods struggle to focus on shadow regions and are consistently distracted by dark areas, whereas our method can minimize such interference to accurately localize shadows. For instance, in the first row, dark streaks formed at the road edge above tree shade cause other methods to mistakenly classify them as shadows. Furthermore, the black baffle surrounding the sandpit in the second row also disrupts other approaches, while DTTNet effectively identifies shadow regions while excluding these distractors—demonstrating that our network does not simply equate dark regions with shadows. In the third row, the gray rough ground texture is highly misleading: TBGDiff (Zhou et al. 2024) fails to classify it as background, whereas our results exhibit the least noise. The last row shows a blurry high-speed captured scene where the athlete’s shadow is nearly indistinguishable from the track and occupies a tiny area, leading other methods to miss the shadow entirely. In contrast, our network successfully discriminates the blurry shadow with the highest accuracy, showcasing its segmentation performance.

Ablation Studies

In this section, we first conduct ablation studies on the ViSha and CVSD dataset to demonstrate the effectiveness of our proposed modules. We remove all proposed modules and retaining only the decoder for generating masks to build the baseline model. The experimental results of the baseline model are presented in Table 4(a).

Comparisons on Efficiency and Parameter As shown in Table 3, we compare the parameter count and speed of recent video shadow detection networks on the Visha dataset. Since we freeze most of the parameters, the total number of parameters is provided in parentheses, with the number of learnable parameters outside the parentheses. DTTNet is able to reach the state-of-the-art performance and the faster inference speed. Notably, when using Automatic mixed precision (Micikevicius et al. 2017), it can get real-time efficiency up to 26.63 fps.

Effectiveness of Vision-language Match Module The Vision-language Match Module (VMM) is designed to align CLIP’s text and image features, bridging the semantic gap between textual priors and visual content for shadow and dark regions. As shown in Table 4(b), removing VMM leads to a significant performance drop: IoU decreases by 0.015 and F1-score drops by 0.01 compared to the the model with both VMM and DSB. This validates the necessity of VMM

Index	Components			ViSha (Chen et al. 2021)				CVSD (Duan et al. 2024)			
	VMM	DSB	TTB	MAE↓	F_β ↑	IoU↑	BER↓	MAE↓	F_β ↑	IoU↑	BER↓
(a)	-	-	-	0.027	0.775	0.640	10.01	0.048	0.720	0.501	25.85
(b)	-	✓	-	0.023	0.807	0.671	8.47	0.045	0.731	0.522	24.77
(c)	✓	✓	-	0.021	0.817	0.686	7.98	0.044	0.754	0.534	22.61
(d)	-	-	✓	0.017	0.833	0.712	7.27	0.043	0.761	0.545	23.68
(e)	✓	✓	✓	0.016	0.849	0.718	6.45	0.042	0.766	0.548	23.32

Table 4: Ablation study on components of DTTNet. The best values are highlighted in bold.




Figure 4: Qualitative comparison results of state-of-the-art methods. In comparison to other methods, our results exhibit less noise, and the predictions for shadow boundaries are more accurate. (b-d) are the best methods in IOS, ISD, and VOS in Table 1, and (e-g) are the latest networks in VSD.

in leveraging CLIP’s zero-shot capability to introduce semantic priors. The degradation confirms the critical role of VMM in enabling cross-modal knowledge transfer. Without VMM, the model loses the ability to effectively fuse text features (P_{ts}, P_{td}) with image features (P_x), resulting in misalignment between semantic guidance and visual content. The cross-attention mechanism of VMM can match text features to corresponding visual information, thereby reducing the differences between modalities and enabling the DSB to better fuse text priors into the visual features of the encoder.

Effectiveness of Dark-aware Semantic Block The Dark-aware Semantic Block (DSB) integrates textual context into visual features and employs penumbra-aware supervision to prioritize shadow semantic learning. It works with VMM to help the model to learn more shadow information from linguistic priors. Removing DSB and VMM from the full model results in a 0.82 BER increase and 0.016 F1-score drop, as Table 4(c) shows, demonstrating its key role in refining semantic representation. Since the DSB is designed to acquire semantics from priors rather than refine the original image features for primary pixels and the decreased supervision in the edge (the penumbra area), its impact on the IoU score is relatively minor.

Effectiveness of Tokenized Temporal Block. The Tokenized Temporal Block (TTB) decouples temporal modeling from spatial features via learnable tokens, enabling efficient aggregation of cross-frame shadow patterns. As shown in Table 4(d), removing TTB results in a significant performance drop, with a 0.032 IoU decrease and 3.9% F1-score

reduction. This highlights the importance of temporal feature modeling in Video Shadow Detection. Without TTB, the baseline model relies on naive frame-wise feature concatenation, which struggles to capture temporal correlations. When adding TTB to baseline model, the performance can be enhanced greatly. TTB’s token-based summarization effectively distills temporal invariants into compact tokens, which are then spatially injected to enhance frame-level features. As it avoids redundant cross-frame attention operations, the computation cost is relatively small. As Table 3 shows, although we employ TTB in every layer of the encoder, the model can get over 26.63 fps in a single RTX 3090 when using automatic mixed precision.

Conclusion

In this paper, we present DTTNet, a novel framework that integrates dark-aware linguistic guidance with tokenized temporal modeling. Specifically, the Tokenized Temporal Block (TTB) efficiently capture shadow dynamics via compact tokens to enhance cross-frame coherence. The Vision-language Match Module (VMM) leverages textual priors to explicitly guide attention toward shadows within dark regions, complemented by the Dark-aware Semantic Block (DSB) for adaptive feature weighting. Additionally, we reweight the penumbra regions for supervision to help the network focus on shadow body in the early stage. Experiments show that each component is effective and our approach can reach state-of-the-art results.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant Nos. 62172417, 62272461, 62472424 and 62476189.

References

- Adams, H.; Stefanucci, J.; Creem-Regehr, S.; and Bodenheimer, B. 2022. Depth perception in augmented reality: The effects of display, shadow, and position. In *2022 IEEE conference on virtual reality and 3D user interfaces (VR)*, 792–801. IEEE.
- Adams, H.; Stefanucci, J.; Creem-Regehr, S.; Pointon, G.; Thompson, W.; and Bodenheimer, B. 2021. Shedding light on cast shadows: An investigation of perceived ground contact in ar and vr. *IEEE transactions on visualization and computer graphics*, 28(12): 4624–4639.
- Chen, T.; Li, L.; Saxena, S.; Hinton, G.; and Fleet, D. J. 2023. A generalist framework for panoptic segmentation of images and videos. In *ICCV*, 909–919.
- Chen, Z.; Wan, L.; Zhu, L.; Shen, J.; Fu, H.; Liu, W.; and Qin, J. 2021. Triple-cooperative video shadow detection. In *CVPR*, 2715–2724.
- Chen, Z.; Zhu, L.; Wan, L.; Wang, S.; Feng, W.; and Heng, P.-A. 2020. A multi-task mean teacher for semi-supervised shadow detection. In *CVPR*, 5611–5620.
- Cheng, H. K.; Tai, Y.-W.; and Tang, C.-K. 2021. Rethinking space-time networks with improved memory coverage for efficient video object segmentation. *NeurIPS*, 34: 11781–11794.
- Cong, R.; Guan, Y.; Chen, J.; Zhang, W.; Zhao, Y.; and Kwong, S. 2023. Sddnet: Style-guided dual-layer disentanglement network for shadow detection. In *ACM MM*, 1202–1211.
- Contributors, M. 2020. MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark.
- Cucchiara, R.; Grana, C.; Piccardi, M.; and Prati, A. 2003. Detecting moving objects, ghosts, and shadows in video streams. *IEEE transactions on pattern analysis and machine intelligence*, 25(10): 1337–1342.
- Deng, Z.; Hu, X.; Zhu, L.; Xu, X.; Qin, J.; Han, G.; and Heng, P.-A. 2018. R3net: Recurrent residual refinement network for saliency detection. In *IJCAI*, volume 684690. AAAI Press Menlo Park, CA, USA.
- Ding, X.; Yang, J.; Hu, X.; and Li, X. 2022. Learning shadow correspondence for video shadow detection. In *ECCV*, 705–722. Springer.
- Duan, X.; Cao, Y.; Zhu, L.; Fu, G.; Wang, X.; Zhang, R.; and Li, P. 2024. Two-stage video shadow detection via temporal-spatial adaptation. In *European Conference on Computer Vision*, 196–214. Springer.
- Guan, H.; Xu, K.; and Lau, R. W. 2024. Delving into dark regions for robust shadow detection. *arXiv preprint arXiv:2402.13631*.
- Hu, X.; Wang, T.; Fu, C.-W.; Jiang, Y.; Wang, Q.; and Heng, P.-A. 2021. Revisiting shadow detection: A new benchmark dataset for complex world. *IEEE TIP*, 30: 1925–1934.
- Ji, Y.; Chen, Z.; Xie, E.; Hong, L.; Liu, X.; Liu, Z.; Lu, T.; Li, Z.; and Luo, P. 2023. Ddp: Diffusion model for dense visual prediction. In *ICCV*, 21741–21752.
- Karsch, K.; Hedau, V.; Forsyth, D.; and Hoiem, D. 2011. Rendering synthetic objects into legacy photographs. *ACM Transactions on graphics (TOG)*, 30(6): 1–12.
- Khan, S. H.; Bennamoun, M.; Sohel, F.; and Togneri, R. 2014. Automatic feature learning for robust shadow detection. In *CVPR*, 1939–1946. IEEE.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *ICCV*, 4015–4026.
- Lalonde, J.-F.; and Matthews, I. 2014. Lighting estimation in outdoor image collections. In *2014 2nd international conference on 3D vision*, volume 1, 131–138. IEEE.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *CVPR*, 2117–2125.
- Liu, L.; Prost, J.; Zhu, L.; Papadakis, N.; Liò, P.; Schönlieb, C.-B.; and Aviles-Rivero, A. I. 2023a. Scotch and soda: A transformer video shadow detection framework. In *CVPR*, 10449–10458.
- Liu, W.; Lu, H.; Fu, H.; and Cao, Z. 2023b. Learning to upsample by learning to sample. In *ICCV*, 6027–6037.
- Lu, X.; Cao, Y.; Liu, S.; Long, C.; Chen, Z.; Zhou, X.; Yang, Y.; and Xiao, C. 2022. Video shadow detection via spatio-temporal interpolation consistency training. In *CVPR*, 3116–3125.
- Lu, X.; Wang, W.; Ma, C.; Shen, J.; Shao, L.; and Porikli, F. 2019. See more, know more: Unsupervised video object segmentation with co-attention siamese networks. In *CVPR*, 3623–3632.
- Micikevicius, P.; Narang, S.; Alben, J.; Diamos, G.; Elsen, E.; Garcia, D.; Ginsburg, B.; Houston, M.; Kuchaiev, O.; Venkatesh, G.; et al. 2017. Mixed precision training. *arXiv preprint arXiv:1710.03740*.
- Milletari, F.; Navab, N.; and Ahmadi, S.-A. 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, 565–571. Ieee.
- Oh, S. W.; Lee, J.-Y.; Xu, N.; and Kim, S. J. 2019. Video object segmentation using space-time memory networks. In *ICCV*, 9226–9235.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Shao, Y.; Taff, G. N.; and Walsh, S. J. 2011. Shadow detection and building-height estimation using IKONOS data. *International journal of remote sensing*, 32(22): 6929–6944.
- Voigtlaender, P.; Chai, Y.; Schroff, F.; Adam, H.; Leibe, B.; and Chen, L.-C. 2019. Feelvos: Fast end-to-end embedding learning for video object segmentation. In *CVPR*, 9481–9490.

- Wang, J.; Li, X.; and Yang, J. 2018. Stacked conditional generative adversarial networks for jointly learning shadow detection and shadow removal. In *CVPR*, 1788–1797.
- Wang, Y.; Zhou, W.; Mao, Y.; and Li, H. 2023. Detect any shadow: Segment anything for video shadow detection. *IEEE TCSVT*.
- Wei, H.; Xing, G.; Liao, J.; Zhang, Y.; and Liu, Y. 2024. Structure-aware spatial-temporal interaction network for video shadow detection. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24, International Joint Conferences on Artificial Intelligence Organization*, 1425–1433.
- Xie, E.; Wang, W.; Yu, Z.; Anandkumar, A.; Alvarez, J. M.; and Luo, P. 2021. SegFormer: Simple and efficient design for semantic segmentation with transformers. *NeurIPS*, 34: 12077–12090.
- Yang, H.; Wang, T.; Hu, X.; and Fu, C.-W. 2023. SILT: Shadow-aware Iterative Label Tuning for Learning to Detect Shadows from Noisy Labels. In *ICCV*, 12687–12698.
- Zheng, Q.; Qiao, X.; Cao, Y.; and Lau, R. W. 2019. Distraction-aware shadow detection. In *CVPR*, 5167–5176.
- Zhou, H.; Wang, H.; Ye, T.; Xing, Z.; Ma, J.; Li, P.; Wang, Q.; and Zhu, L. 2024. Timeline and Boundary Guided Diffusion Network for Video Shadow Detection. In *ACM MM*, 166–175.
- Zhu, L.; Deng, Z.; Hu, X.; Fu, C.-W.; Xu, X.; Qin, J.; and Heng, P.-A. 2018. Bidirectional feature pyramid network with recurrent attention residual modules for shadow detection. In *ECCV*, 121–136.
- Zhu, L.; Xu, K.; Ke, Z.; and Lau, R. W. 2021. Mitigating intensity bias in shadow detection via feature decomposition and reweighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4702–4711.