Contrastive Learning of Semantic and Visual Representations for Text Tracking

Zhuang Li[1*], Weijia Wu[2*], Mike Zheng Shou[3], Jiahong Li[1†], Size Li[1], Zhongyuan Wang[1], Hong Zhou[2†]

[1]Kuaishou Technology [2]Zhejiang University [3]Show Lab, National University of Singapore

Abstract

*Semantic representation is of great benefit to the Video Text Tracking (VTT) task that requires simultaneously classifying, detecting, and tracking texts in the video. Most existing approaches tackle this task by modeling the appearance similarity of texts across video frames while ignoring the cues provided by the semantics of texts. To address this issue, we propose to jointly model semantic and visual representations with a contrastive learning objective. In this paper, we present an end-to-end video text tracker with **S**emantic and **V**isual **R**epresentations (SVRep), which exploit both visual and semantic relationships between different texts in a video sequence. We conduct extensive experiments on five benchmarks. With a light-weight recognition head and contrastive head, SVRep achieves state-of-the-art performance while maintaining competitive inference speed. For example, our SVRep runs at **13.4** FPS and achieves **66.4%*** $\mathrm{ID_{F1}}$ *on ICDAR2015(video) dataset, improving the previous state-of-the-art methods by **8.2%** absolute gain.*

1. Introduction

Video text tracking(VTT) [34] is the task that requires simultaneously classifying, detecting, and tracking text instances in video. Most mainstream approaches [36, 32] all follow the *tracking-by-detection* paradigm, which firstly tackles each frame in a video sequence, and then associates the similar text instances in the adjacent frames by various appearance-based matching strategies(*i.e.,* use their IoU, transcription, and feature). IoU-based methods [32, 13] rely on image text detection models [30], where IoU of detected text boxes from two adjacent frames is higher than the given threshold are associated. Similarly, visual feature-based methods [3, 36] compute visual feature similarity from different text instances in adjacent frames, to join pairs with the same text instance.

However, for tracking text in a video sequence, it is not only important to locate the texts but more critical to *read* and *comprehend* it in the context of the visual scene. The

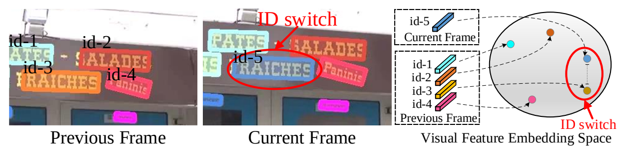

(a) The previous video text trackers [3, 36] associate texts with visual feature (*i.e.,* appearance feature).

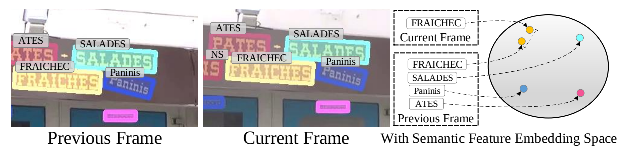

(b) SVRep tracks text by maximizing mutual information between text pairs with semantic information, not only visual feature.

Figure 1 – Existing video text tracking models cannot read, but humans and our model can. Human usually watches video by frequently reading, tracking, and comprehending the semantic information of texts in video. In this work, we try to explore tracking texts via recognizing text, relating it to other texts.

previous methods all ignore the textual cues and semantic relationship between different text instances. In contrast to the visual features of appearance, semantic features are robust cues for matching and tracking text instances in a video sequence. As shown in Figure. 1a, due to rapid motion blur, out-of-focus, and artifacts issues, the pure visual feature usually causes the ID switch of having already tracked text instances. But it is easy for humans to track the word "*FRAICHEC*" via recognizing, comprehending, and relating it to other text (*e.g., SALADES*, *Paninis*), although the low-quality video. Therefore, we try to allow the model to read text, relate it to other texts, and then track. But there are two main challenges: 1) **Reading**. How to learn efficient visual and semantic representations of text implicitly for an end-to-end video text tracking model. 2) **Relationship Modeling**. How to maximize the mutual information of visual and semantic representations between text instances of multiple frames.

To address the two challenges, we firstly propose an end-to-end video text tracker with **S**emantic and **V**isual representations (SVRep), which detects, and tracks texts by exploiting the visual and semantic relationships between different texts in a video sequence. As shown in Figure. 1b, given a video clip with texts, SVRep learns the visual and semantic representations by detecting and recognizing text, then relating it to other texts in the video clip. With the abundant semantic representations, visual representations, and mutual information, the more stable tracking performance is presented. **For the former challenge**, we propose three encoders, *i.e., positional encoder*, *visual encoder*, *semantic encoder*. Positional encoder learns the location information by embedding bounding box coordinates of each text instance. Visual encoder learns visual appearance features (*e.g.,* color, shape, texture) of each text instance by embedding the visual feature from the Masked RoI [31]. Semantic encoder learns semantic features of each text instance by embedding the sequence feature from an CTC-based recognition head. **For the later challenge**, inspired by contrastive learning [1], with semantic and visual embedding, we maximize agreement between the same texts (positive pairs) and maximizing disagreement between different texts (negative pairs) by contrasting positive pairs against negative pairs.

To achieve high efficiency, we reduce the time cost of each step by the following four designs: 1) PAN++ [31] as a lightweight architecture (*i.e.,* backbone, fpn, up-sample) is adopted to combine the proposed contrastive learning of semantic and visual representations; 2) A low-computation CTC-based recognition head with only one-layer BiLSTM is proposed to recognize text content; 3) Three lightweight encoders are employed for semantic and visual representation learning. Benefiting from the above designs, SVRep achieves a high inference speed while keeping the state-of-the-art accuracy. On the ICDAR2015(video) [14] dataset, the $\mathrm{ID_{F1}}$ of our SVRep reaches 66.4%, with up to **10%** improvements than that of Free [3] and TransVTSpotter [33], while its inference speed (*i.e.,* **13.4** FPS) is faster.

The main contributions of this work are listed below:

- We propose three light-weight encoders, *i.e.,* positional encoder, visual encoder, semantic encoder, which learn location representation, visual representation, and semantic representation, respectively. With the three encoders, the proposed model is the first one, *like human*, to track text by recognizing it.
- With contrastive learning, we firstly propose to maximize agreement between the same texts (*i.e.,* positive pair) and maximizing disagreement between different texts (*i.e.,* negative pairs) in feature embedding space for solving the video text tracking task.
- SVRep achieves the state-of-the-art performance on *five* public datasets with faster speed. Especially, SVRep achieves 66.4% $\mathrm{ID_{F1}}$ and 13.4 FPS for video text tracking task on ICDAR2015 [14], with **8.2%** improvements than the previous SOTA methods.

2. Related Work

2.1. Text Detection and Tracking

Recent methods [30, 26, 31] based on deep learning have been made tremendous progress for image-level text detection. PSENet [30] proposed the post-processing of progressive scale expansion for improving the detection accuracy. PAN++ [31] proposed an end-to-end text spotting framework, which detects and recognizes a text with a feature extractor (Masked RoI), and a lightweight attention-based recognition head. However, these image-based methods can not obtain temporal information (*i.e.,* tracking id) in the video, which is essential for other video-and-language tasks such as video understanding.

The detailed survey [34] summarizes and compares the existing text tracking, and recognition methods in a video before 2016. [29] captures spatial-temporal information by exploiting the cues of the background regions of the text. Wang *et al.* [32] links texts in the current frame and several previous frames to obtain the tracking results by hand-crafted post-processing, such as IoU-based associations. [36] tracked texts by using ConvLSTM to capture spatial structure information and motion memory, and an appearance-geometry descriptor is proposed to learn the visual representation of text instances. TransVTSpotter[33] introduced a query-based tracking branch to associate text instance with IoU matching in adjacent frame. The above methods track text by appearance similarity, ignoring the abundant semantic features. In this work, we try to allow the model to recognize text, relate it to other texts with semantic and visual features, and then track.

2.2. Semantic Representations

Many previous works have been shown the effectiveness of semantics in other tasks, such as text recognition [24, 35], cross-modal video retrieval [23], and knowledge distillation [4]. SEED [24] tries to correct the recognition mistakes via using the semantic information (*i.e.,* word embedding) from the pre-trained language model. Yu *et al.,* [35] enhance the recognition result with a global semantic reasoning module, which can simultaneously perceive the semantic information of all characters in a word or text line. For video text tracking task, the current most video text spotting works [33, 3] associate text instances with only visual feature. And few works try to explore the robustness of semantic features. ASGD [5] tries to extract semantics via encoding the classification of characters to enhance video text detection. But the features still belong to visual features for learning from classification tasks. In this paper, the proposed SVRep firstly tackles the video text tracking task via extracting strong semantic information from the sequence model of the recognition head.

2.3. Contrastive Learning

Dating back to [11], these approaches learn visual representations by measuring the similarities of sample pairs in representation space. And the contrastive loss is at the core of several recent works on unsupervised learning [1, 12] by contrasting positive pairs against negative pairs. Except for unsupervised learning, there are relations between generative adversarial networks [7] and contrastive losses (*i.e.,* noise-contrastive estimation [10]). The contrastive loss [7], as a widely successful technique for unsupervised data generation, measuring the difference between probability distributions. In this work, we firstly adopt the contrastive losses (*i.e.,* noise-contrastive estimation [10]) to model the relationship between different text pairs both in visual representation and semantic representation spaces.

3. Method

The architecture of the proposed method is shown in Figure. 2. PAN++ [31] as the base network is adopted, including backbone, FPN, Up-sample, and detection head. And the recognition head and track head are firstly proposed in the paper. For semantic and visual representations learning, we propose three encoders, *i.e.,* positional encoder, visual encoder and semantic encoder. To track text with visual and semantic relationship of text, SVRep maximizes agreement between the same text example of different frames via a contrastive loss (*i.e.,* $\mathcal{L}_{\text{con}}$) in the latent semantic and visual embedding space. Given a set $\{x_n^k\}$ (*e.g.,* 'Text', 'Mall' in Figure. 2) of text instances in a video sequences, including a positive pair (*i.e.,* the same text in different frame) of examples x_t^i and x_{t-1}^j, where t means t-th frame and i means i-th text instance in t-th frame. The *contrastive learning* aims to identify x_t^i in $\{x_n^k | k \neq j \cup n \neq t-1\}$ for a given x_{t-1}^j both in semantic and visual embedding spaces. Both in visual and semantic embedding spaces, we argue that the same text in a continuous video sequence should tend to close and different text with different identification shows a greater distance.

3.1. Semantic and Visual Representation Learning

Human usually watches videos by frequently reading, tracking, and comprehending visual and semantic information. To follow the human mechanism, we design an recognition head and three encoders (*i.e.,* positional encoder, visual encoder, and semantic encoder) for simultaneously learning the visual and semantic representations.

Recognition Head. An CTC-based recognition head [26] is proposed to replace the attention-based recognition head [31]. There are two benefits for using CTC-based: 1) much faster inference time using parallel decoding. 2) more abundant semantic representation without massive random noise. More detailed analysis and experiments for the two points can be found in the supplementary material. With Masked RoI, the visual feature patches are extracted for each text instance, and resized to a fixed-size feature ($h \times w$). Similar to PAN++ [31], in all experiments, h and w are set to 8 and 32 pixels, respectively. Then one-layer Bidirectional LSTM (BiLSTM) [9], as the sequence model, extracts a semantic sequence feature $H = \mathrm{Seq.(V)}$ from the visual feature patches. As shown in Figure. 2, the transcription module [26] produces the final output (a sequence of characters) from the input semantic sequence $H = h_1, \ldots, h_T$, where T is the sequence length. Then Connectionist Temporal Classification (CTC) [8] trains the network to optimize the summation of probabilities over all paths:

$$p(Y|H) = \sum_{\pi:\mathcal{M}(\pi)=Y} p(\pi|H), \tag{1}$$

$$\mathcal{L}_{rec} = -\log p(Y|H), \tag{2}$$

where $\mathcal{M}$ defines the operation of mapping all possible paths π to the target label. For example, it maps the path "$\mathtt{ttt} - - \mathtt{e} - \mathtt{xxx} - \mathtt{t} - -$" into "$\mathtt{text}$".

Semantic Encoder. In this encoder, the semantic sequence feature $H = h_1, \ldots, h_T$ from the output of BiLSTM in recognition head is fed into semantic encoder for learning a high-dimensional semantic representation (*i.e.,* feature) $\mathcal{R}_s$, whose size is $128 \times 1 \times 1$. To reduce the computation overhead, we only employ two 1×1 convolutions with 1 stride, two batch normalization layers, and max-pooling layer to extract the feature in the encoder, as shown in Figure 2 (Down).

Visual Encoder. Corresponding to the semantic sequence feature H from the recognition head, the fixed-size visual feature patches V of each text instance from Masked RoI is used to extract the high-dimensional visual representation $\mathcal{R}_v$. The use of V from Masked RoI has three benefits: 1) The visual feature patches V of each text instance contains abundant visual appearance representations, *e.g.,* color, shape, texture. 2) The binary mask of the rotated bounding box can eliminate the noise features caused by the background or other text lines, so as to accurately extract the visual features. 3) The reusing with recognition head reduces the time cost of feature extraction. Similar to semantic encoder, two 3×3 convolutions with 1 stride, two batch normalization layers, and max-pooling layer are used to extract the high-dimensional feature. The detailed architecture for the encoder is provided in supplementary material.

Positional Encoder. Except for the attribute (*i.e.,* visual and semantic features) of text instance, the positional information (location) is equally important for tracking. Inspired by VL-BERT [27], we extract the high-dimensional positional representation $\mathcal{R}_p$ by embedding bounding box coordinates of each text instance (*i.e.,* positional encoder). Each Masked RoI of text is characterized by a 4-d vector, as

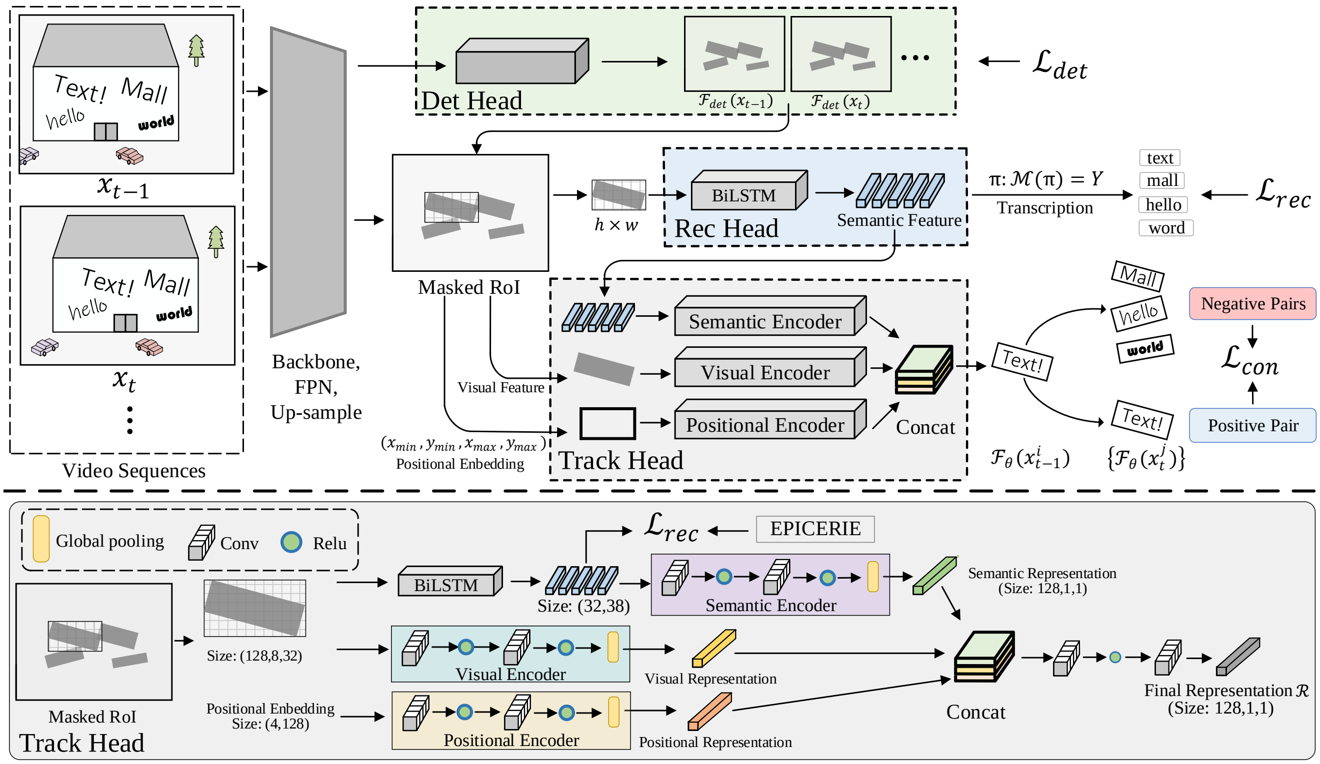

Figure 2 – The overall architecture of SVRep. Upper: It contains three main components: 1) Kernel-based Detection head [31] is used to predict the rotated bounding box; 2) Recognition head predicts text context from Masked RoI with CTC loss; 3) Track head includes three encoders, which are used to learning the high-dimensional visual and semantic embedding representations by maximizing agreement between the same text and maximizing disagreement between different texts in a video sequence. Down: The detailed architecture for track head with three encoders, *i.e.*, semantic, visual and positional encoder.

$(\frac{x_{LT}}{W}, \frac{y_{LT}}{H}, \frac{x_{RB}}{W}, \frac{y_{RB}}{H})$, where (x_{LT}, y_{LT}) and (x_{RB}, y_{RB}) denote the coordinate of the top-left and bottom-right corner respectively, and W,H are of the width and height of the input image. Then the 4-d vector is embedded into a positional feature embedding (of 4×128 in paper) by computing sine and cosine functions of different wavelengths. But the relative positional feature embedding is still a low-level feature, which can not be extended to the high-dimensional representation space. Therefore, we use two 1×1 convolutions with 1 stride, two batch normalization layers, and max-pooling layer to further convert the positional feature embedding to the final high-dimensional positional representation $\mathcal{R}_p$, whose size is $128 \times 1 \times 1$.

3.2. Contrastive Learning for Text Instance

The idea behind contrastive learning is to learn an embedding that separates (contrasts) samples from two different distributions. Given a text instance set $\{x_n^k\}$ in a video sequence, *e.g.*, 'Text', 'Mall' in figure. 2, where n means n-th frame and k means k-th text instance in n-th frame. We consider the same text in adjacent frames, *e.g.*, 'Text' x_t^i in t-th frame and 'Text' x_{t-1}^j in (t-1)-th frame, which we call *positive pair*, different text instances with different semantic and visual information (*e.g.*, 'Text' and 'Mall') in a video sequence, which we call *negative pair*.

A "critic" (a discriminating function) $\mathcal{F}_\theta(\cdot)$ is trained to maximize agreement for positive pairs and maximizing disagreement for negative pairs in semantic and visual embedding space. And for each text instance x_n^k, the corresponding positional representation $\mathcal{R}_p$, semantic representation $\mathcal{R}_s$, and visual representation $\mathcal{R}_v$ have already obtained from three encoders, and require further integration for learning the discriminating function $\mathcal{F}_\theta(\cdot)$. The three features ($\mathcal{R}_p$, $\mathcal{R}_v$, $\mathcal{R}_s$) are concatenated together directly, and fed into two 1×1 convolutions to output the final high-dimensional representation $\mathcal{R}_n^k$, whose size is $b \times c$ (c is set to 128 in paper).

For convenience, we flatten all text examples of T frames, define $\mathcal{R}_n^k$ as z_p, where $p = k + d$ (d means the sum of text examples before n-th frame). And we randomly sample a minibatch of T frames, N text examples and define the contrastive prediction task on positive pairs of the same text in different frames, resulting in $N/2$ data points. Similar to SimCLR [1], we do not sample negative examples explicitly.

Instead, given a positive pair (*i.e.,* z_p, z_q), we treat the other text pairs within a minibatch as negative examples. Let $\mathrm{sim}(\boldsymbol{u}, \boldsymbol{v}) = \boldsymbol{u}^\top \boldsymbol{v}/\|\boldsymbol{u}\|\|\boldsymbol{v}\|$ denote the cosine similarity between two vectors $\boldsymbol{u}$ and $\boldsymbol{v}$. Then the loss function for a positive pair of examples z_p and z_q is defined as:

$$\mathcal{L}_{p,q} = -\log \frac{\exp(\mathrm{sim}(\boldsymbol{z}_p, \boldsymbol{z}_q)/\tau)}{\sum_{K=1}^{N} \mathbb{1}_{[K \neq p]} \exp(\mathrm{sim}(\boldsymbol{z}_p, \boldsymbol{z}_K)/\tau)} , \quad (3)$$

where $\mathbb{1}_{[K \neq p]} \in \{0, 1\}$ is an indicator function, and τ denotes a temperature parameter. The final *NT-Xent* loss [1] (contrastive loss) is computed across all positive pairs, both (p, q) and (p, q) in a mini-batch:

$$\mathcal{L}_{track} = \underset{\{z_1, z_2, \ldots, z_N\}}{\mathbb{E}} [\mathcal{L}_{p,q}] , \quad (4)$$

where z_p, $z_q \in \{z_1, z_2, \ldots, z_N\}$ is the set of the final text representations from the discriminating function $\mathcal{F}_\theta(\cdot)$.

3.3. Loss Function of Multi-Task Learning

The proposed pipeline contains three losses, *i.e.,* detection loss, recognition loss, and contrastive loss, which belong to three different tasks. To improve learning efficiency and prediction accuracy, Multi-task learning [16] is adopted in our method to learn multiple objectives from a shared representation:

$$\mathcal{L} = \frac{1}{2\sigma_1^2}\mathcal{L}_{det} + \frac{1}{2\sigma_2^2}\mathcal{L}_{rec} + \frac{1}{2\sigma_3^2}\mathcal{L}_{track} + \log \sigma_1\sigma_2\sigma_3 , \quad (5)$$

where σ_1, σ_2, σ_3 are three learnable parameters, and the $\log \sigma_1\sigma_2\sigma_3$ is a regulariser.

3.4. Inference

In the inference phase, similar to many previous works [32, 5], SVRep obtains the final tracking result (*i.e.,* tracking trajectory) by the bipartite cosine distance matching for each text pair (z_q, z_p) in the adjacent frame and the Kuhn-Munkres(KM) algorithm [18].

4. Experiments

4.1. Dataset

ICDAR2013 Video [15] is a widely used benchmark for video text task, containing 13 videos for training and 15 videos for testing. All videos are collected and annotated from daily scenarios, and each text is labeled as a quadrangle with 8 coordinates of four corners in a clock-wise manner. **ICDAR2015 Video [14]** is the expanded version of ICDAR2013(video), which consists 50 videos with 22 additional videos. **Minetto [20]** is a small dataset, including 5 videos in outdoor scenes, and the model trained on ICDAR2015 Video is used to evaluate this dataset directly. **YouTube Video Text (YVT) [21]** dataset is harvested from YouTube, contains 30 videos, where 15 videos for training and 15 videos for testing. **BOVText [33]** is a bilingual, open-world dataset, including 2,021 videos with 1,757,598 frames. The data is collected from the worldwide user of YouTube and KuaiShou, covering various daily scenarios.

4.2. Implementation Details

All the experiments are conducted on PyTorch with 8 Tesla V100 GPUs. We use the PAN++ [31] as our basic network. Following the common practices [30, 31], we ignore the blurred text regions labeled as "DO NOT CARE" during training, and apply random scale, random horizontal flip, random rotation, and random crop on training images. Smimilar to TransVTSpotter [33], COCO-Text [28] is used as the pretrained dataset. COCO-Text is a largest scene text detection dataset with 63,686 images, which reuses the images from MS-COCO dataset [19]. For the static images from COCO-Text, we apply the random shift [37] to generate video clips with pseudo tracks. All models are optimized by using ADAM optimizer with a batch size of 48 on 8 GPUs. The initial learning rate is set to 1×10^{-3}. In the testing phase, we resize the input image to the fixed scale. All results are tested with a batch size of 1 on a V100 GPU and a 2.20GHz CPU in a single thread. In the metric, Mostly Tracked ($Tracked$) denotes the number of objects tracked for at least 80 percent of the lifespan, and Mostly Lost ($Lost$) denotes the number of objects tracked for less than 20 percent of the lifespan.

Inconsistency from Evaluation Metric. The inconsistency between the reported performance and the original paper lies in the update of the evaluation metric. From 2020, Robust Reading Competition Website [1] has updated the evaluation method, and the old metric and ground truth is not available. All works have already used the new metric after 2020.

4.3. Ablation Study

Semantic, Visual and Positional Encoder. Table. 1b presents the effect on final performance from the three encoders on ICDAR2015video. To focus on the effect of different encoders, we implement all of them in the same backbone (*i.e.,* ResNet) and loss function (*i.e.,* NT-Xent loss). When we only evaluate one encoder (*i.e.,* visual encoder), other two embedding features (*i.e.,* $\mathcal{R}_p$, $\mathcal{R}_s$) would not be used in the discriminating function $\mathcal{F}_\theta(\cdot)$ for contrastive learning as described in Sec. 3.2. As such, the comparison is solely on the three encoders. The model presents a base performance (*i.e.,* 64.7% $\mathrm{ID_{F1}}$ and 49.1% MOTA) with visual encoder. With visual and semantic features, the model obtains the $\mathrm{ID_{F1}}$ of 66.2% and MOTA of 50.1%, achieving

[1] https://rrc.cvc.uab.es/?ch=3&com=evaluation&task=1

IoU-based	Contrastive-based	ID_{F1}	MOTA	MOTP
✓		57.5	41.2	74.3
	✓	66.2	50.2	73.6
✓	✓	66.4	50.3	73.9

(a) Different text instances association methods. IoU-based and Contrastive-based refer to tracking text with IoU match and semantic and visual feature match in adjacent frame, respectively.

Visual	Semantic	Positional	ID_{F1}	MOTA	MOTP
✓			64.7	49.1	73.6
	✓		62.2	46.8	73.8
✓	✓		66.2	50.1	73.6
✓	✓	✓	66.4	50.3	73.9

(b) Effect of Semantic, Visual and Positional Encoder. 'Semantic', 'Visual', 'Positional' refers to the Semantic, Visual and Positional representations, respectively.

Name	τ or m	ID_{F1}	MOTA	MOTP
Margin Triplet	0.05	52.6	48.0	73.4
	0.1	61.0	48.1	73.4
	0.5	62.2	48.2	73.4
NT-Xent	0.05	64.1	48.5	73.7
	0.1	66.4	50.3	73.9
	0.5	65.3	49.1	73.8

(c) Effect of loss functions on ICDAR2015video. 'τ' denotes the temperature for NT-Xent loss, and 'm' refers to the margin for Margin Triplet loss.

Backbone	Shorter Side	ID_{F1}	MOTA	MOTP	FPS
ResNet18	512	60.1	44.3	73.6	23.3
ResNet18	640	65.2	48.4	74.2	18.6
ResNet18	720	65.9	49.1	73.8	16.7
ResNet50	512	60.3	44.7	73.8	19.7
ResNet50	640	65.5	48.6	73.6	15.4
ResNet50	720	66.4	50.3	73.9	13.4

(d) Effect of backbone and input image shorter side on ICDAR2015video. NT-Xent loss is used for the experiment.

Table 1 – Ablation experiments for SVRep. All models are trained and tested on ICDAR2015video.

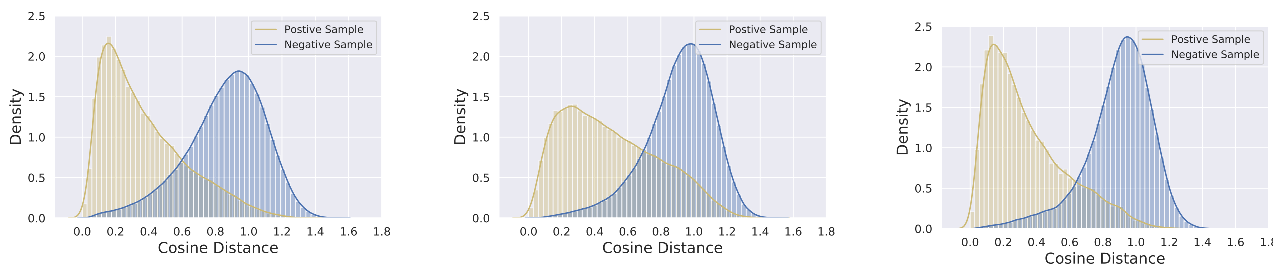

(a) Only visual representation. (ID_{F1}: 64.3, MOTA: 51.6)

(b) Only semantic representation.ID_{F1}: 60.1, MOTA: 50.0)

(c) Combination of both.(ID_{F1}: 66.3, MOTA: 53.3)

Figure 3 – Effect of semantic and visual representation on ICDAR2013video. The cosine distance of the positive sample refer to the distance of the same text (same tracking id) in adjacent frame. Negative sample refer to different texts in adjacent frame. Positional encoder is used in the experiment.

1.5% and 1.0 improvements, respectively. To further evaluate the effectiveness of semantics, we present the probability density distribution of cosine distance for different representations on ICDAR2013video, as shown in Figure 3. The combination of semantic and visual representations bring a stronger discrepancy between positive and negative sample. Note: to avoid the effect from the detection head, we directly use the bounding box GT to extract the RoI features in Figure. 3.

Loss for Contrastive Learning. We study the effect of the contrastive loss, and compare the NT-Xent loss against other commonly used contrastive loss functions, such as margin loss [25]. As shown in Table 1c, NT-Xent loss shows a better performance, around 4% improvement than the counterpart of Margin Triplet loss. τ is the temperature parameter for NT-Xent loss in Equ. 3, and we set three values to evaluate the effect from it. With the change of τ, ID_{F1} as the tracking stability metric, shows higher volatility than MOTA and MOTP, since τ affect the weight of positive text pair and negative text pairs. In this paper, we set the τ to 0.1 for the competitive performance. Beside, we also study the importance of the dimensions of the representation $\mathcal{R}$ for contrastive learning, as shown in Figure. 4b. Similar results are observed regardless of output dimension, and we set the dimension to 128 in other experiments.

Speed Analysis. Table.1d presents the time cost of SVRep with different backbones and input image shorter side. We evaluate all testing images and calculate the average speed. These results are tested with 1 batch size on one V100 GPU and one 2.20GHz CPU in a single thread. With ResNet18, 512 pixels of image shorter side, model presents the faster speed with 23.3 fps. On the contrary, model shows the best performance (*i.e.,* 66.4% ID_{F1}), while the inference speed is slow with 13.4 fps. Besides, we compare with other methods in Figure. 4a. Moreover, "SVRep+ResNet18" reaches **16.7** FPS, which is 7.7 FPS faster than that of Free [3] and TransVTSpotter [33], while its performance achieves a great improvement with ID_{F1} of around 8.0%.

Contrasive Learning $v.s$ IoU-based Matching. Contrastive Learning, as the core of the paper, is the main difference from the previous works. As shown in Figure. 1a, IoU-based bounding box match is used to compare our feature-based cosine similarity with contrastive learning. For a fair comparison, IoU-based and contrastive-based methods in this section all use the same framework (*i.e.,* the

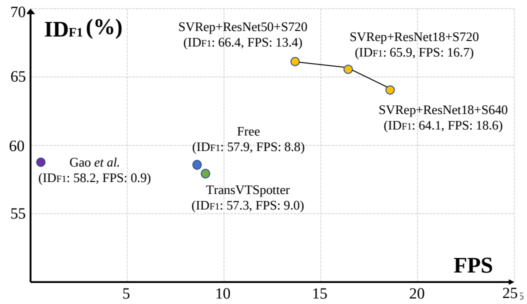

(a) The comparisons of video text tracking methods on ICDAR2015video, in terms of both accuracy and speed.

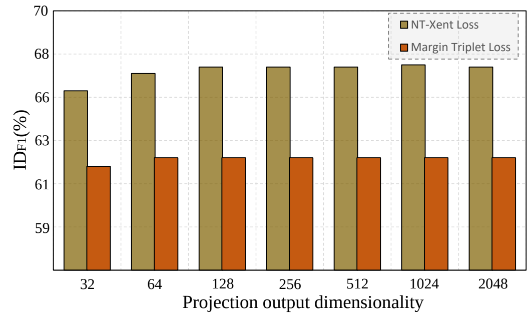

(b) Effect of contrastive loss and various dimensions of the final representation $\mathcal{R}$ on ICDAR2015video.

Dataset	Method	Video Text Tracking/%					
		$\mathrm{ID_{F1}}$	MOTA	MOTP	MostlyMatched↑	MostlyLost↓	FPS↑
ICDAR2015(video)[14]	USTB [14]	25.9	7.4	70.8	7.4	66.1	-
	StradVision [14]	25.9	7.9	70.2	6.5	70.8	-
	USTB(II-2) [14]	21.9	12.3	71.8	4.8	72.3	-
	AJOU [17]	36.1	16.4	72.7	14.1	62.0	-
	Free [3]	57.9	43.2	**76.7**	36.6	44.4	8.8
	TransVTSpotter [33]	57.3	44.1	75.8	34.3	33.7	9.0
	Gao *et al.* [6]	58.2	44.1	75.2	44.8	29.0	0.9
	SVRep+ResNet18+S(720)	65.9	49.1	73.8	44.4	27.7	**16.7**
	SVRep+ResNet50+S(720)	**66.4**	**50.3**	73.9	**45.0**	**26.5**	13.4
ICDAR2013(video)[15]	YORO†[2]	62.5	47.3	73.7	33.1	45.3	14.3
	SVRep+ResNet18+S(720)	**66.3**	**53.3**	**75.9**	**38.4**	**32.2**	**17.8**
BOVText[33]	EAST [33]	28.1	-21.6	75.8	-	-	-
	PSENet [33]	45.9	52.1	77.5	-	-	-
	DB [33]	48.3	53.2	78.3	-	-	-
	TransVTSpotter[33]	64.7	68.2	82.1	57.3	31.4	9.0
	SVRep+ResNet18+S(720)	**75.4**	**69.3**	**84.5**	**59.0**	**29.7**	**12.2**
Minetto[20]	Zuo *et al.* [38]	-	56.4	73.1	-	-	-
	Pei *et al.* [22]	-	73.1	57.7	-	-	-
	AGD&AGD[36]	-	75.6	74.7	-	-	-
	Yu *et al.*[36]	-	81.3	75.7	-	-	-
	ASGD[5]	-	83.5	76.8	-	-	-
	TransVTSpotter[33]	74.7	84.1	77.6	-	-	-
	SVRep+ResNet18+S(720)	**83.9**	**86.3**	**81.0**	**96.4**	**0**	**19.5**
YVT[21]	Free[3]	-	54.0	**78.0**	-	-	-
	TransVTSpotter[33]	64.5	53.9	75.9	-	-	9.0
	SVRep+ResNet18+S(720)	**69.1**	**54.4**	74.2	**53.8**	**29.1**	**16.2**

Table 2 – Text tracking performance on five public datasets. ‘Tracked’ and ‘Lost’ denote ‘Mostly Tracked’ and ‘Mostly Lost’, respectively. † refers to our testing performance. ‘S:’ means the shorter side of input image.

proposed SVRep), the only difference is the tracking association method. With the same condition, contrastive-based feature cosine similarity match shows a better performance than that of IoU-based, with at least 8% $\mathrm{ID_{F1}}$ improvement.

4.4. Comparison with State-of-the-arts

We compare SVRep against state-of-the-art methods for video text tracking task on *five* public benchmarks. Besides, the related experiments of video text detection task are provided in the supplementary material.

Minetto (English, small-scale dataset). Minetto, as one small dataset to evaluate the *robustness* of SVRep. Following the previous works [5, 33], we train on ICDAR2015(video) and evaluate the model on the Minetto dataset directly. Figure. 2 presents that SVRep with ResNet18 achieves a better performance (83.9% $v.s.$ 74.7% for $\mathrm{ID_{F1}}$) than TransVTSpotter [33] with a faster inference speech (19.5 fps).

ICDAR2015video and ICDAR2013video (English, medium-scale data). Both as the two most popular public benchmarks are used to evaluate our method. As shown in figure. 2, our SVRep shows a powerful performance with 66.4% $\mathrm{ID_{F1}}$ and 50.3% MOTA on ICDAR2015video, achieving 8.2% and 6.2% improvements than the previous SOTA method (*i.e.,* [6]), respectively. Besides, when the input image shorter side is 720 pixels, with ResNet18, the inference speed of our method reaches 16.7 FPS, which is faster than previous methods, while the F-measure is still very competitive (65.9%). Similarly, the proposed SVRep achieves a competitive performance and inference speed with 66.3% $\mathrm{ID_{F1}}$ and 17.8 FPS.

YVT (English, medium-scale data). Different from ICDAR2015video with only scene text, the YouTube Video Text dataset (YVT) mainly includes overlay text and scene text (*e.g.,* street signs, business signs, words on shirt). There are few reported tracking results on YVT. Similar to other datasets, SVRep achieves the state-of-the-art performance with 69.1% $\mathrm{ID_{F1}}$, while maintaining high inference (*i.e.,* 16.2 fps).

BOVText (bilingual, large-scale dataset). BOVText is a bilingual and large-scale dataset with more than 2 million video frames, which was collected from the worldwide user of YouTube and KuaiShou. With ResNet18 and 720 pixel short side, our method achieves 75.4% $\mathrm{ID_{F1}}$ for video text tracking task, at least 10% improvements than the previous works, such as TransVTSpotter. The great performance on the practice benchmark further shows the remarkable robustness and generalization of our SVRep. For the long caption text challenge (text average width-height ratio more than 6.8) on BOVText, the obvious improvement proves the robustness of SVRep for long text tracking. Besides, the proposed SVRep further presents the advantage of the inference speed, especially for the large-scale dataset, *i.e.,* BOVText, saving at least 6 hours time cost.

5. Conclusion

Unlike the previous works that track text by visual feature, we firstly propose an end-to-end video text tracker with semantic and visual representations, which tracks text by exploiting the visual and semantic relationships between texts with contrastive learning in high-dimensional embedding space. Without bells and whistles, SVRep achieves the best result with up to 8.2% $\mathrm{ID_{F1}}$ improvement and the highest speed among methods using a single model on the ICDAR2015(video) dataset. To our knowledge, our work is the first one that applies strong textual semantic information and contrastive learning to video text tracking task. We hope that the work can bring some new insights for the community, and more similar approaches with semantic knowledge can be applied to video-and-language tasks in the future.

References

[1] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607. PMLR, 2020.

[2] Zhanzhan Cheng, Jing Lu, Yi Niu, Shiliang Pu, Fei Wu, and Shuigeng Zhou. You only recognize once: Towards fast video text spotting. In *ACM MM*, pages 855–863, 2019.

[3] Zhanzhan Cheng, Jing Lu, Baorui Zou, Liang Qiao, Yunlu Xu, Shiliang Pu, Yi Niu, Fei Wu, and Shuigeng Zhou. Free: A fast and robust end-to-end video text spotter. *TIP*, 30:822–837, 2020.

[4] Ali Cheraghian, Shafin Rahman, Pengfei Fang, Soumava Kumar Roy, Lars Petersson, and Mehrtash Harandi. Semantic-aware knowledge distillation for few-shot class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2534–2543, 2021.

[5] Wei Feng, Fei Yin, Xu-Yao Zhang, and Cheng-Lin Liu. Semantic-aware video text detection. In *CVPR*, pages 1695–1705, 2021.

[6] Yuzhe Gao, Xing Li, Jiajian Zhang, Yu Zhou, Dian Jin, Jing Wang, Shenggao Zhu, and Xiang Bai. Video text tracking with a spatio-temporal complementary model. *IEEE Transactions on Image Processing*, 30:9321–9331, 2021.

[7] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *NeurIPS*, 27, 2014.

[8] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *ICML*, pages 369–376, 2006.

[9] Alex Graves, Marcus Liwicki, Santiago Fernández, Roman Bertolami, Horst Bunke, and Jürgen Schmidhuber. A novel connectionist system for unconstrained handwriting recognition. *TPAMI*, 31(5):855–868, 2008.

[10] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS*, pages 297–304, 2010.

[11] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *CVPR*, volume 2, pages 1735–1742. IEEE, 2006.

[12] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020.

[13] Xian-Sheng Hua, Pei Yin, and Hong-Jiang Zhang. Efficient video text recognition using multiple frame integration. In *ICIP*, volume 2. IEEE, 2002.

[14] Dimosthenis Karatzas, Lluis Gomez-Bigorda, Anguelos Nicolaou, Suman Ghosh, Andrew Bagdanov, Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vijay Ramaseshan Chandrasekhar, and Shijian Lu. Icdar 2015 competition on robust reading. In *ICDAR*, pages 1156–1160, 2015.

[15] Dimosthenis Karatzas, Faisal Shafait, Seiichi Uchida, Masakazu Iwamura, Lluis Gomez i Bigorda, and Sergi Robles Mestre. Icdar 2013 robust reading competition. In *ICDAR*, pages 1484–1493, 2013.

[16] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *CVPR*, pages 7482–7491, 2018.

[17] Hyung Il Koo and Duck Hoon Kim. Scene text detection via connected component clustering and nontext filtering. *TIP*, 22(6):2296–2305, 2013.

[18] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 1955.

[19] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755. Springer, 2014.

[20] Rodrigo Minetto, Nicolas Thome, Matthieu Cord, Neucimar J Leite, and Jorge Stolfi. Snoopertrack: Text detection and tracking for outdoor videos. In *ICIP*, pages 505–508, 2011.

[21] Phuc Xuan Nguyen, Kai Wang, and Serge Belongie. Video text detection and recognition: Dataset and benchmark. In *WACV*, pages 776–783, 2014.

[22] Wei-Yi Pei, Chun Yang, Li-Yu Meng, Jie-Bo Hou, Shu Tian, and Xu-Cheng Yin. Scene video text tracking with graph matching. *IEEE Access*, 6:19419–19426, 2018.

[23] Mengshi Qi, Jie Qin, Yi Yang, Yunhong Wang, and Jiebo Luo. Semantics-aware spatial-temporal binaries for cross-modal video retrieval. *IEEE Transactions on Image Processing*, 30:2989–3004, 2021.

[24] Zhi Qiao, Yu Zhou, Dongbao Yang, Yucan Zhou, and Weiping Wang. Seed: Semantics enhanced encoder-decoder framework for scene text recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13528–13537, 2020.

[25] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, pages 815–823, 2015.

[26] Baoguang Shi, Xiang Bai, and Cong Yao. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *TPAMI*, 39(11):2298–2304, 2016.

[27] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vl-bert: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530*, 2019.

[28] Andreas Veit, Tomas Matera, Lukas Neumann, Jiri Matas, and Serge Belongie. Coco-text: Dataset and benchmark for text detection and recognition in natural images. *arXiv preprint arXiv:1601.07140*, 2016.

[29] Lan Wang, Yang Wang, Susu Shan, and Feng Su. Scene text detection and tracking in video with background cues. In *ACM MM*, pages 160–168, 2018.

[30] Wenhai Wang, Enze Xie, Xiang Li, Wenbo Hou, Tong Lu, Gang Yu, and Shuai Shao. Shape robust text detection with progressive scale expansion network. In *CVPR*, pages 9336–9345, 2019.

[31] Wenhai Wang, Enze Xie, Xiang Li, Xuebo Liu, Ding Liang, Yang Zhibo, Tong Lu, and Chunhua Shen. Pan++: Towards efficient and accurate end-to-end spotting of arbitrarily-shaped text. *TPAMI*, 2021.

[32] Xiaobing Wang, Yingying Jiang, Shuli Yang, Xiangyu Zhu, Wei Li, Pei Fu, Hua Wang, and Zhenbo Luo. End-to-end scene text recognition in videos based on multi frame tracking. In *ICDAR*, volume 1, pages 1255–1260, 2017.

[33] Weijia Wu, Debing Zhang, Yuanqiang Cai, Sibo Wang, Jiahong Li, Zhuang Li, Yejun Tang, and Hong Zhou. A bilingual, openworld video text dataset and end-to-end video text spotter with transformer. In *NeurIPS*, 2021.

[34] Xu-Cheng Yin, Ze-Yu Zuo, Shu Tian, and Cheng-Lin Liu. Text detection, tracking and recognition in video: a comprehensive survey. *TIP*, 25(6):2752–2773, 2016.

[35] Deli Yu, Xuan Li, Chengquan Zhang, Tao Liu, Junyu Han, Jingtuo Liu, and Errui Ding. Towards accurate scene text recognition with semantic reasoning networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12113–12122, 2020.

[36] Hongyuan Yu, Yan Huang, Lihong Pi, Chengquan Zhang, Xuan Li, and Liang Wang. End-to-end video text detection with online tracking. *PR*, 113:107791, 2021.

[37] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking objects as points. In *ECCV*, pages 474–490, 2020.

[38] Ze-Yu Zuo, Shu Tian, Wei-yi Pei, and Xu-Cheng Yin. Multi-strategy tracking based text detection in scene videos. In *ICDAR*, pages 66–70. IEEE, 2015.