

# Siamese Natural Language Tracker: Tracking by Natural Language Descriptions with Siamese Trackers

Qi Feng  
Boston University  
fung@bu.edu

Vitaly Ablavsky  
University of Washington  
vxa@uw.edu

Qinxun Bai  
Horizon Robotics  
qinxun.bai@gmail.com

Stan Sclaroff  
Boston University  
sclaroff@bu.edu

## Abstract

We propose a novel Siamese Natural Language Tracker (SNLT), which brings the advancements in visual tracking to the tracking by natural language (NL) descriptions task. The proposed SNLT is applicable to a wide range of Siamese trackers, providing a new class of baselines for the tracking by NL task and promising future improvements from the advancements of Siamese trackers. The carefully designed architecture of the Siamese Natural Language Region Proposal Network (SNL-RPN), together with the Dynamic Aggregation of vision and language modalities, is introduced to perform the tracking by NL task. Empirical results over tracking benchmarks with NL annotations show that the proposed SNLT improves Siamese trackers by 3 to 7 percentage points with a slight tradeoff of speed. The proposed SNLT outperforms all NL trackers to-date and is competitive among state-of-the-art real-time trackers on LaSOT benchmarks while running at 50 frames per second on a single GPU. Code for this work is available at <https://github.com/fredfung007/snlt>.

## 1. Introduction

Visual and language recognition skills evolve jointly in children from a young age. For example, it was observed [33] that children at the age of twenty months whose vocabulary size lags behind their peers have difficulty recognizing objects with sparse features (*i.e.*, stylized versions of real-world objects). Conversely, a child’s ability to engage in the play activity called *object substitution* tends to be a predictor of healthy language development [34].

By contrast, in computer vision, particularly in object tracking, appearance-based methods [24, 25] and tracking via natural-language (NL) descriptions [14, 26] evolve independently, without benefiting each other.

In this paper, we derive a formulation that links




Figure 1: The proposed Siamese Natural Language Tracker (SNLT) improves Siamese trackers by leveraging predictions from two modalities: vision and language. Our SNLT implementation runs at 50 frames per second on an NVIDIA 2080 Ti GPU, and outperforms best published natural language trackers, *i.e.* Feng [14] and Li [26], on the LaSOT [13] test set.

the vision and language modalities in such a way that improvements in appearance-based tracking yield improvements in language-based tracking. Our formulation applies to Siamese trackers, a broad family of trackers that includes SiamFC [1], SiamRPN [25], SiamRPN++ [24], *etc.* Siamese trackers have proven to be successful in many tracking scenarios and have achieved state-of-the-art performance. We show that by bringing the advancements of these Siamese trackers to the tracking by NL task, we can build NL trackers that outperform all prior NL trackers and promise to see improvements over time with the advancements of Siamese trackers.

Hence, we present the first practical and general so-

lution to the challenge of tracking with NL descriptions in real-time. Firstly, we propose a Siamese Natural Language Region Proposal Network (SNL-RPN) that transforms an NL description into a convolutional kernel and shares feature extraction layers with Siamese trackers; the combined network can be trained end-to-end. Secondly, we propose a novel formulation to dynamically aggregate the predictions of our SNL-RPN from two modalities: vision and language, which turns the SNL-RPN from a visual-language detector into a real-time Siamese Natural Language Tracker (SNLT). The overview of a realization of the proposed SNLT is shown in Fig. 2.

We plot the frames per second (FPS) v.s. the precision on LaSOT [13] for recent Siamese trackers and our SNLT in Fig. 1<sup>1</sup>. Our proposed SNLT consistently improves the performance of SiamFC [1], SiamRPN [25], and SiamRPN++ [24] with a slight trade-off of speed. It also outperforms all NL trackers to date. This demonstrates that the SNLT brings the advancement of visual tracking models to the tracking by NL task and provides a wide range of state-of-the-art NL trackers.

Contributions of this paper are threefold:

1. A novel and universal Siamese Natural Language Region Proposal Network (SNL-RPN) is proposed for all Siamese trackers, providing a wide class of strong tracking by NL descriptions baselines.
2. A Dynamic Aggregation of predictions from vision and language modalities, is proposed to transform our SNL-RPN into a real-time Siamese Natural Language Tracker (SNLT). Prior to this work, we are only aware of two NL trackers [14, 26].
3. Empirical results over tracking benchmarks with NL annotations show that the proposed SNLT improves Siamese trackers by 3 to 7 percentage points. The SNLT outperforms all NL trackers and is competitive with state-of-the-art real-time trackers on LaSOT benchmarks while running at over 50 frames per second.

## 2. Related Works

### 2.1. Visual Object Tracking

In the past two decades, tracking by detection models [3, 20] and Bayesian filtering based algorithms [4, 21] have been thoroughly studied in the field of visual

<sup>1</sup>We use the code and weights from the original authors for Li [26]’s tracker. As the language query “dictionary” used in the original work is different from LaSOT [13], the performance reported here is sub-optimal as training code is not available.

object tracking. Some deep learning based models [1, 7, 29, 35] have been introduced in recent years, and are argued to perform better when handling occlusion and appearance change. ECO [7] applies convolutional filters on convolution feature maps to obtain satisfactory performance on multiple tracking datasets. ECO still suffers from efficiency issues [18], though its efficiency is improved from the original convolution filter based tracker, C-COT [9]. These trackers maintain appearance and motion models explicitly by maintaining the visual features over time. ATOM [6] introduced a classification module that is updated online to better handle scenarios where multiple similar targets exist. DiMP [2] aims to train an online model to distinguish the background and foreground of the exemplar during tracking, which further pushes the performance of convolutional filter based trackers. PrDiMP [8] introduced probabilistic regression to further improve the DiMP tracker. KL-divergence based loss is first introduced to train the regression network of the PrDiMP.

On the other hand, a series of Siamese trackers are introduced by exploiting a siamese convolutional neural network architecture for tracking by detection. SiamFC [1] conducts a local search for regions with similar regional visual features obtained by a CNN in every frame. SiamRPN [25] and SiamRPN++ [24] performs tracking as one shot detection using the Siamese network as a region proposal network. However, these Siamese trackers do not model the temporal appearance variations of the target and therefore suffer from model drift problems. SiamRCNN [38] is the most recent Siamese tracker that produces state-of-the-art tracking performance on several benchmark datasets by performing a global search via re-detection while trading off the speed. SiamRCNN runs at only 4 frames per second.

### 2.2. Natural Language Processing in Vision Tasks

In the past decade, researchers have started to look into exploiting natural language understanding in vision tasks. These models usually combine two components: a language model and an appearance model to learn a new feature space that is shared between both NL and appearance [19, 37]. More recent object detection and vision grounding models [17, 40] jointly exploit vision and NL using Siamese networks and depth-wise convolutional neural networks between the NL representations and visual representations.

Li *et al.* define two tracking by NL descriptions problems [26]. Feng *et al.* formalize the tracking by NL in a tracking by detection framework with a Bayesian detection formulation [14]. In their work, however, an assumption is made that appearances and the NL

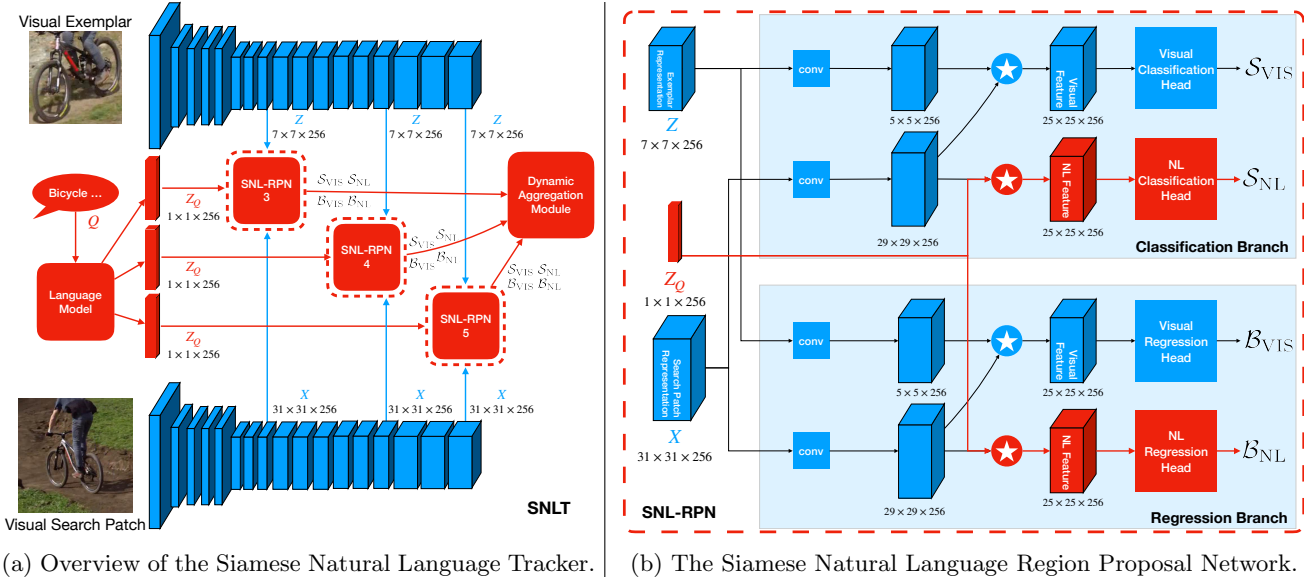


Figure 2: **2a** shows an overview of the proposed Siamese Natural Language Tracker (SNLT) and **2b** shows its key component: the Siamese Natural Language Region Proposal Network (SNL-RPN). Without loss of generality, we use SiamRPN++ as the backbone for this example realization of the SNLT. Novel modules of our proposed architecture are highlighted in **red**. The Language Model predicts representations of the input NL description for the SNL-RPN. The three SNL-RPN modules in **2a** are identical except for their inputs. The Aggregation Module, which dynamically combines predictions from both visual and language modalities based on the entropy of predictions, is described in Sec. 3.3. As shown in **2b**, the SNL-RPN consists of two branches: a Regression branch and a Classification branch. The SNL-RPN takes the convolution feature maps of the template  $Z$ , the convolution feature maps of the search patch  $X_t$ , and the sentence embedding  $Z_Q$  as inputs and predicts classification scores and regressions for each of the predefined anchor boxes in the SNL-RPN. The star operator, depends on the backbone Siamese tracker, can be cross-correlation, depth-wise cross correlation, *etc.*

description are conditionally independent given the bounding boxes. By directly measuring the joint conditional probability between the language network and the visual network, in this paper, we derive a fully convolutional neural network (CNN) that performs tracking by NL description. Following Li *et al.*, the NL description is defined as a declarative sentence of arbitrary length for the target. Similar to Li *et al.*'s work [25], we formulate the tracking with NL description problem as one-shot detection.

### 3. Siamese Natural Language Tracker

In this section, we present the Siamese Natural Language Tracker (SNLT), which works in conjunction with a wide range of Siamese trackers including SiamFC [1], SiamRPN [25], DaSiamRPN [41], SiamRPN++ [24], and more recent Siam R-CNN [38].

#### 3.1. Overview

In Fig. 2a, we present in detail how SNLT incorporates and enhances Siamese trackers with NL de-

scriptions of the target, using SiamRPN++ [24] as an example backbone. Realization of the SNLT for other Siamese trackers can be derived in a similar way.

The SNLT takes three inputs for each frame, a visual exemplar, a visual search patch, and a language query  $Q$ . We use convolutional neural networks (CNNs), *e.g.* AlexNet [23] and ResNet-50 [16], to extract visual representations of the visual exemplar and visual search patch, denoted as  $Z$  and  $X$  respectively. We use a Language Model to compute a sentence embedding of the NL description  $Q$ . We use  $Z_Q$  to denote this embedding. The Language Model for the SNLT can be any sentence embedding model, and in our experiments we use GloVe [30], HGLMM [5] and BERT [11] based models.

The triplet  $(Z, Z_Q, X)$  is then passed on to the proposed Siamese Natural Language Region Proposal Network (SNL-RPN), which will predict bounding box classification scores and regressions on a set of predefined anchors for both vision and language modalities. We use  $\mathcal{S}$  and  $\mathcal{B}$  to denote the classification scores and regressions in our detailed derivation of the SNL-

RPN in Sec. 3.2.

After predictions from the SNL-RPN are obtained, a Dynamic Aggregation Module combines the predictions from the vision and language modalities. The Dynamic Aggregation Module is explained in detail in Sec. 3.3.

Our SNLT enhances existing Siamese trackers by exploiting an NL description of the target and reducing the chance of model drift that is common in Siamese trackers. The SNLT significantly improves the performance of Siamese trackers and outperforms previous tracking by NL description approaches by a large margin.

### 3.2. Architecture of the SNL-RPN

The proposed SNLT is built upon a Siamese Natural Language Region Proposal Network (SNL-RPN), as shown in Fig. 2b. The novel components are presented in red blocks and arrows.

For each input triplet  $(Z, Z_Q, X)$ , the SNL-RPN outputs two sets of different predictions, one from the Visual Head and the other from the NL Head. Same as an ordinary region proposal network (RPN), the SNL-RPN has two branches for each of the Visual Head and the NL Head: the Classification branch and the Regression branch for its anchors.

For both the Classification branch and the Regression branch, a depth-wise cross correlation between  $X$  and  $Z$ , shown as  $\star$  in Fig. 2b, is used to compute visual feature maps for the Visual Head. Another depth-wise cross correlation between  $X$  and  $Z_Q$  is used to compute the RPN feature maps for the NL Head. We use  $\mathcal{S}_{\text{VIS}}$  and  $\mathcal{S}_{\text{NL}}$  to denote the scores predicted by the Classification branch, and  $\mathcal{B}_{\text{VIS}}$  and  $\mathcal{B}_{\text{NL}}$  to denote the regressions predicted by the Regression branch.

The layer-wise aggregation of the SiamRPN++ is a weighted sum between three predictions from ResNet group 3, 4, and 5 respectively. The weights for this group-wise aggregation are trained offline and remain fixed during inference. For the SNL-RPN, similarly, we train a set of weights offline that aggregates the predictions from ResNet group 3, 4, and 5 for both the Visual Head and NL Head *independently*. *i.e.*

$$\begin{aligned}
 \mathcal{S}_{\text{VIS}} &= \sum_{i=3,4,5} \mathcal{S}_{\text{VIS}}^{\text{Group } i} \\
 \mathcal{S}_{\text{NL}} &= \sum_{i=3,4,5} \mathcal{S}_{\text{NL}}^{\text{Group } i} \\
 \mathcal{B}_{\text{VIS}} &= \sum_{i=3,4,5} \mathcal{B}_{\text{VIS}}^{\text{Group } i} \\
 \mathcal{B}_{\text{NL}} &= \sum_{i=3,4,5} \mathcal{B}_{\text{NL}}^{\text{Group } i}
 \end{aligned} \tag{1}$$

Note that for a simpler backbone tracker, *e.g.* SiamRPN and SiamFC, no such layer-wise aggregation is needed.

### 3.3. Aggregation of the SNL-RPN Predictions

In order to jointly predict the tracking update from both visual and language cues, we introduce another type of aggregation beyond the layer-wise aggregation in SiamRPN++: the aggregation between the Visual Head and the NL Head predicted by the SNL-RPN. We define the aggregation as

$$\begin{aligned}
 \mathcal{S} &= w_{\text{VIS}} \cdot \mathcal{S}_{\text{VIS}} + w_{\text{NL}} \cdot \mathcal{S}_{\text{NL}}; \\
 \mathcal{B} &= w_{\text{VIS}} \cdot \mathcal{B}_{\text{VIS}} + w_{\text{NL}} \cdot \mathcal{B}_{\text{NL}}.
 \end{aligned} \tag{2}$$

Intuitively, we can train the aggregation weights  $w_{\text{VIS}}$  and  $w_{\text{NL}}$  offline, which is essentially an estimate of the reliability of the predictions based on language cues and visual cues. However, as we are taking predictions from two networks that consume different inputs, similar to online learning setups in [32], it is not ideal to keep fixed weights between them.

Therefore, we design the aggregation between the NL Head and the Visual Head to be dynamic based on the predictions and the inputs. The entropies for a predicted score map are defined as:

$$\begin{aligned}
 H_{\text{VIS}} &= - \sum \mathcal{S}_{\text{VIS}} \cdot \log \mathcal{S}_{\text{VIS}}; \\
 H_{\text{NL}} &= - \sum \mathcal{S}_{\text{NL}} \cdot \log \mathcal{S}_{\text{NL}}.
 \end{aligned} \tag{3}$$

Ablation studies in Sec. 4.4 show a negative correlation between the entropy of  $\mathcal{S}$  and the Intersection over Union (IoU) between the prediction and the ground truth bounding box. Therefore, we give less weight to either the NL Head or the Visual Head when they have a high entropy on their classification scores:

$$w_{\text{VIS}}, w_{\text{NL}} = \sigma([\alpha \cdot H_{\text{NL}}, \alpha \cdot H_{\text{VIS}}]), \tag{4}$$

where  $\sigma$  is the softmax function and  $\alpha$  is the ‘‘temperature,’’ *i.e.* a constant scalar to scale the entropies. Note that in Eq. 4, the subscript ‘‘VIS’’ and ‘‘NL’’ for  $w$  are swapped in the left hand side compared with that for the  $H$  in the right hand side. As a result, when aggregating the Visual Head and NL Head, the one with a higher entropy will have a lower weight.

### 3.4. Training the SNL-RPN and Loss Functions

To construct training instances that resemble the test-time distribution, we randomly choose two frames at different time steps, together with the corresponding ground truth bounding boxes. We crop and resize  $Z$  for the visual exemplar, and  $X$  for the search patch.



Thus, a triplet  $(Z, Z_Q, X)$ , is constructed as the input for training our proposed tracker.

We follow the training process of the RPN in Faster RCNN [15] for the SNL-RPN to sample 16 positive anchors and 48 negative anchors. Positive anchors have an IoU with the ground truth bounding box greater than 0.7, while negative anchors have an IoU less than 0.3. We use a softmax cross entropy loss, denoted by  $L_{\text{cls}}$ , for training Classification branches and a smoothed  $L1$ -loss, denoted by  $L_{\text{reg}}$ , for training Regression branches. The overall training loss is

$$L(Z, Z_Q, X) = L_{\text{cls}}(\mathcal{S}_{\text{NL}}) + L_{\text{reg}}(\mathcal{B}_{\text{NL}}) + L_{\text{cls}}(\mathcal{S}_{\text{VIS}}) + L_{\text{reg}}(\mathcal{B}_{\text{VIS}}). \quad (5)$$

## 4. Experiments

In this section, we first describe the datasets and implementation details in our experiments. Then, we compare our tracker with state-of-the-art visual and NL trackers. Finally, we present ablation studies to demonstrate the effectiveness of our proposed SNLT, the SNL-RPN and the novel Dynamic Aggregation Module.

### 4.1. Datasets

**Training Datasets:** The backbone networks used in this work, AlexNet [23] and ResNet [16], are pretrained on ImageNet [10]. We use all images and phrases from VisualGenome [22], and frames from MSCOCO [27] and YouTube-BoundingBox [31], together with images and phrases from training splits of LaSOT and OTB-99-LANG<sup>2</sup> [26] for training the SNL-RPN. We follow the same size of  $127 \times 127$  pixels for the template patch  $Z$  and the size of  $255 \times 255$  pixels for the search patch  $X$  during training.

**Evaluation Datasets:** Given the novelty of the tracking by NL description setup, we are aware of only two publicly available tracking benchmarks that are **annotated with NL** for targets. In Li *et al.*'s early work on NL tracking [26], they annotated OTB-100 [39] with NL to produce the OTB-99-LANG dataset. In a more recent work [13], LaSOT, a large single object tracking benchmark dataset annotated with NL for targets, was introduced with 70 different categories of objects and 20 sequences for each category, totaling at 1,400 sequences. We choose to follow protocol 2 from [13], to evaluate our tracking by One Pass Evaluation (OPE) on the testing split of the dataset.

<sup>2</sup>Note that this is different from OTB-100.

### 4.2. Implementation Details

**Training Initialization:** We initialize the AlexNet [23] and the stride-reduced ResNet [24] with pretrained weights on ImageNet [10] and randomly initialize layers in the SiamRPN/SiamRPN++ [25, 24]. Layers other than the sentence encoder in the proposed SNL-RPN are initialized randomly from  $N(0, 1)$ .

We use the average of the word embeddings as the sentence embedding with HGLMM [5] and GLOVE [30] in addition to a pretrained BERT sentence encoding model [11] in our ablation studies. Additionally, we fine-tune a BERT sentence embedding model [11] to achieve the state-of-the-art tracking performance.

**Learning Rates and Convergence:** We train our proposed SNL-RPN using a PyTorch implementation on GPUs with an Adagrad [12] optimizer and an initial learning rate of 0.001. We decay the learning rate after 5 epochs to 0.0005 and continue the training for another 5 epochs. Batch size is set to 256/64 triplets of  $Z$ ,  $Q$ , and  $X$  per GPU. Gradients are averaged over each batch, while the gradients for the NL Head in SNL-RPN are omitted if  $Q$  is not present. Under these settings, the training process takes 5 hours to converge on 16 GPUs using the loss described in Sec. 3.4.

**Inference Hyper-parameter Selection:** The  $\alpha$  used in the Dynamic Aggregation Module between vision and language modalities is set to 300 throughout our experiments. In addition to the introduced  $\alpha$ , existing hyper-parameters that are standard in Siamese trackers, *e.g.* sub-window attention, gating, *etc.*, are chosen via validation experiments. We use the values released by the original authors for all hyper-parameters shared by SNLT and backbone Siamese trackers. We chose hyper-parameters that are unique to our SNLT on the training split of LaSOT. The  $\alpha$  we chose will not result in a hard switch. The  $w_{\text{VIS}}$  and  $w_{\text{NL}}$  on LaSOT test videos range from roughly 0.13 to 0.86.

### 4.3. Comparison with Visual and NL Trackers

We compare the proposed SNLT with the following state-of-the-art real-time trackers: SiamRPN++ [24], SiamRPN [25], PrDiMP [8], DiMP [2], ATOM [6], MD-Net [28], and VITAL [35]. For the fairness of comparisons, we use their released codes, model weights, and hyper-parameters in all experiments. Success, Precision, and Normalized Precision Plots on LaSOT [13] and OTB-99-LANG test splits are presented in Fig. 3 and Fig. 4 respectively. The SNLT tracker improves the SiamRPN++ baseline on both the LaSOT and the OTB-99-LANG. In OTB-99-LANG, where videos are typically less than 300 frames, the model drift problem is much less frequent than that in the LaSOT




Figure 3: Success, Precision, and Normalized Precision Plots on the LaSOT test set.




Figure 4: Success, Precision, and Normalized Precision Plots on OTB-99-LANG.

Backbone	NL	Success	Norm. Prec.
SiamRPN++	BERT	<b>54.0</b>	<b>63.6</b>
SiamRPN++	HGLMM	<i>50.4</i>	<i>60.0</i>
SiamRPN++	GLOVE	49.9	59.4
SiamRPN++	N/A	48.9	58.0
SiamRPN	BERT	<b>46.0</b>	<b>54.6</b>
SiamRPN	HGLMM	<i>44.1</i>	<i>53.7</i>
SiamRPN	GLOVE	43.6	52.8
SiamRPN	N/A	42.2	50.9

Table 1: Ablation studies of the proposed SNLT tracker on LaSOT test split. Three different sentence embedding models (BERT [11], HGLMM [5] and GLOVE [30]) are used to train our SNL-RPN on both SiamRPN [25] and SiamRPN++ [24] backbones. The best and second best performances are highlighted with bold and italic fonts.

dataset. Our SNLT improves the performance of the SiamRPN++ by 0.8 percentage points. In the LaSOT benchmark, as shown in Fig. 1, the SNLT consistently improves the Siamese trackers by 3 to 7 percentage points, and the top SNLT variant based on

Tracker	OTB		LaSOT	
	Suc.	Norm.	Suc.	Norm.
SNLT	<b>0.67</b>	<b>0.80</b>	<b>0.54</b>	<b>0.64</b>
FENG [14]	0.61	0.73	0.35	0.43
LI [26]	0.55	0.67	-	-
SNLT (SEM)	-	-	0.40	0.48

Table 2: Comparison between LI [26], FENG [14] and the SNLT using both bounding box and NL for initialization. “SEM” stands for using the semantic class as the NL description input to the SNLT tracker. The best performance is highlighted with bold font.

SiamRPN++ is very competitive among state-of-the-art real-time trackers.

Following [14], we evaluate our SNLT on the NL-Consistent LaSOT, a subset of LaSOT test split in which the NL descriptions uniquely describe the target in the video (selected by crowd workers). We show that when the NL descriptions uniquely describe targets, the proposed SNL-RPN outperforms all prior works. Success, Precision and Normalized Precision Plots on NL-consistent LaSOT are presented in Fig. 5.





Figure 5: Success, Precision, and Normalized Precision Plots on NL-Consistent LaSOT [14]. Our SNLT outperforms siamese trackers by a large margin as the NL descriptions in the NL-Consistent LaSOT uniquely describe the targets.



(a) The given NL uniquely describes the airplane and helps our tracker (red) stay on the target, while SiamFC [1] (purple) and SiamRPN++ [24] (blue) suffers from model drifts.



(b) The given NL (goldfish swimming among other fishes in the water) does not uniquely describes the target. As multiple goldfishes are present in the scene, the NL description does not help our tracker (red) to avoid model drifting.

Figure 6: The ambiguity of the NL description may affect our tracker.

We also compare our proposed tracker with best published NL trackers, Feng [14] and Li [26], on LaSOT and OTB-99-LANG for the tracking with NL task, *i.e.* tracking with both the bounding box initialization and the NL description of the target. Evaluations of [26] on LaSOT are omitted, since no training code was released for it to guarantee a fair comparison. As shown in Tbl. 2, the SNLT outperforms the best published NL trackers by a large margin.

#### 4.4. Ablation Studies

In this section, we conduct comprehensive experiments and ablation studies analyzing the performance of our proposed SNLT, SNL-RPN, and the Dynamic Aggregation Module.

OTB-99-LANG and LaSOT are the only available single object tracking benchmarks with NL annota-

tions that are publicly available. Only one NL description is provided for each sequence. Leveraging such NL descriptions to consistently improve visual tracking ends up very challenging. In Fig. 6, we show one case when the NL description would help our tracker recover from model drift, and another case when the given NL description does not uniquely describe the target and eventually makes our tracker drift away to another object that also matches the NL description. This showcase explains why our SNLT has an even further advantage on the NL-Consistent LaSOT.

We evaluate the correlation coefficient between the entropies defined in Eq. 3 and the IoU between the predicted bounding boxes and ground truth bounding boxes for each video. We plot the distribution of the correlation coefficient on the LaSOT test split in Fig. 7. The negative correlation between the entropies and the




Figure 7: Distribution of the correlation coefficient between the entropies defined in Eq. 3 and IoU of the predicted and ground truth bounding box.

predicted IoU motivates our Dynamic Aggregation, in Eq. 4, of the NL Head and the Visual Head in the SNL-RPN.

Additionally, an example frame from the *airplane-13* video is visualized together with the score maps from the NL Head and the Visual Head in Fig. 8. The proposed Dynamic Aggregation between the two modalities give higher weights to the modality with lower entropy, resulting in a more stable tracker.

We test the SNLT with different Language Models, including sentence embedding model that takes average of word embeddings from GloVe [16] and HGLMM [5], and a BERT sentence embedding model [11]. The SNLT were individually trained with the same procedure as discussed in Sec. 4.2. Comparisons of these trackers are shown in Tbl. 1. The proposed SNL-RPN and the Dynamic Aggregation between vision and language modalities, effectively boost the tracking performance of both SiamRPN and SiamRPN++. HGLMM, a more recent sentence embedding model trained for vision tasks, further pushes the SNLT to a better performance compared to the GloVe based sentence embedding model. The sentence embedding models are not fine-tuned for two reasons: 1. They are pre-trained with large corpuses of texts and the NL tracking datasets are relatively small. 2. There are several layers after  $Z_Q$  that are trained from scratch.

The tracking by NL description problem has its unique challenges compared to the tracking by semantic information [36]. An ablation study using the ground truth semantic class label, *i.e.* category label in LaSOT (*e.g.* airplane, person, *etc.*), as the NL description is reported in Tbl 2. Results show that the SNLT




Figure 8: Visualization of an example of the Visual Search Patch  $X$  (left), Score Map from Visual Head  $S_{VIS}$  (middle) and Score Map from NL Head  $S_{NL}$  (right) in the *airplane-13* video of LaSOT. The NL description of the target is “white airplane flying in the air.” The NL Head is more confident and has a more concentrated prediction of the airplane than the Visual Head, possibly due to the occlusion of the target. Such difference in prediction confidence is reflected in the entropies computed by Eq. 3, then follow Eq. 4, the aggregation weights for this example are  $(w_{VIS}, w_{NL}) = (0.39, 0.61)$ , biasing towards the prediction from the NL Head, as expected for this frame. The proposed Dynamic Aggregation that gives higher weight to the modality with lower entropy results in a more stable tracker.

learns more than the semantic class to outperform the baseline SiamRPN++.

Regarding the speed of the proposed SNLT, it only adds a small overhead to Siamese trackers (most computations from ResNet are shared between Siamese RPN and SNL-RPN). As shown in Fig. 1, the SNLTs are around 5% to 10 % slower than their corresponding backbones when using the NL-RPN for inference, still achieving over 50 frames per second.

## 5. Conclusion and Future Work

We present a novel Siamese Natural Language Tracker (SNLT) and Siamese Natural Language Region Proposal Network (SNL-RPN), which can track a target in a video given an NL description of the target. With the Dynamic Aggregation Module between vision and language modalities, our approach enjoys better robustness than other visual object trackers. Experiments on challenging datasets demonstrate that the SNLT outperforms its backbone trackers by a large margin. Our SNLT and SNL-RPN are generally applicable to all Siamese trackers. Thus, we expect they will be useful in enhancing future Siamese trackers in pursuit of state-of-the-art.

**Acknowledgment:** This work was partially supported by the US NSF under Grant No. 1928477.



## References

- [1] Luca Bertinetto, Jack Valmadre, João F Henriques, Andrea Vedaldi, and Philip HS Torr. Fully-convolutional siamese networks for object tracking. *arXiv preprint arXiv:1606.09549*, 2016.
- [2] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning discriminative model prediction for tracking. In *Proc. IEEE International Conf. on Computer Vision (ICCV)*, pages 6182–6191, 2019.
- [3] Samuel S Blackman. Multiple hypothesis tracking for multiple target tracking. *IEEE Aerospace and Electronic Systems Magazine*, 19(1):5–18, 2004.
- [4] Eli Brookner. g-h and g-h-k filters. *Tracking and Kalman Filtering Made Easy*, pages 3–63, 1998.
- [5] Andrea Burns, Reuben Tan, Kate Saenko, Stan Sclaroff, and Bryan A Plummer. Language features matter: Effective language representations for vision-language tasks. In *Proc. IEEE International Conf. on Computer Vision (ICCV)*, pages 7474–7483, 2019.
- [6] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. Atom: Accurate tracking by overlap maximization. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4660–4669, 2019.
- [7] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. Eco: Efficient convolution operators for tracking. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 6638–6646, 2017.
- [8] Martin Danelljan, Luc Van Gool, and Radu Timofte. Probabilistic regression for visual tracking. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 7183–7192, 2020.
- [9] Martin Danelljan, Andreas Robinson, Fahad Shahbaz Khan, and Michael Felsberg. Beyond correlation filters: Learning continuous convolution operators for visual tracking. In *Proc. European Conf. on Computer Vision (ECCV)*, pages 472–488. Springer, 2016.
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255. Ieee, 2009.
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [12] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.
- [13] Heng Fan, Liting Lin, Fan Yang, Peng Chu, Ge Deng, Sijia Yu, Hexin Bai, Yong Xu, Chunyuan Liao, and Haibin Ling. Lasot: A high-quality benchmark for large-scale single object tracking. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 5374–5383, 2019.
- [14] Qi Feng, Vitaly Ablavsky, Qinxun Bai, Guorong Li, and Stan Sclaroff. Real-time visual object tracking with natural language description. In *Proc. Winter Conf. on Applications of Computer Vision (WACV)*, pages 700–709, 2020.
- [15] Ross Girshick. Fast r-cnn. In *Proc. IEEE International Conf. on Computer Vision (ICCV)*, pages 1440–1448, 2015.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *Proc. European Conf. on Computer Vision (ECCV)*, pages 630–645. Springer, 2016.
- [17] Richang Hong, Daqing Liu, Xiaoyu Mo, Xiangnan He, and Hanwang Zhang. Learning to compose and reason with language tree structures for visual grounding. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 2019.
- [18] Chen Huang, Simon Lucey, and Deva Ramanan. Learning policies for adaptive tracking with deep feature cascades. In *Proc. IEEE International Conf. on Computer Vision (ICCV)*, pages 105–114, 2017.
- [19] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. Denscap: Fully convolutional localization networks for dense captioning. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4565–4574, 2016.
- [20] Zdenek Kalal, Krystian Mikolajczyk, and Jiri Matas. Tracking-learning-detection. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 34(7):1409–1422, 2012.
- [21] Rudolph Emil Kalman et al. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.
- [22] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision (IJCV)*, 123(1):32–73, 2017.
- [23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1097–1105, 2012.
- [24] Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, and Junjie Yan. Siamrpn++: Evolution of siamese visual tracking with very deep networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4282–4291, 2019.
- [25] Bo Li, Junjie Yan, Wei Wu, Zheng Zhu, and Xiaolin Hu. High performance visual tracking with siamese region proposal network. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 8971–8980, 2018.
- [26] Zhenyang Li, Ran Tao, Efstratios Gavves, Cees GM Snoek, and Arnold WM Smeulders. Tracking by natural language specification. In *Proc. IEEE Conf. on*

- Computer Vision and Pattern Recognition (CVPR)*, pages 6495–6503, 2017.
- [27] 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 *Proc. European Conf. on Computer Vision (ECCV)*, pages 740–755. Springer, 2014.
- [28] Hyeonseob Nam and Bohyung Han. Learning multi-domain convolutional neural networks for visual tracking. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4293–4302, 2016.
- [29] Guanghan Ning, Zhi Zhang, Chen Huang, Xiaobo Ren, Haohong Wang, Canhui Cai, and Zhihai He. Spatially supervised recurrent convolutional neural networks for visual object tracking. In *2017 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–4. IEEE, 2017.
- [30] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014.
- [31] Esteban Real, Jonathon Shlens, Stefano Mazzocchi, Xin Pan, and Vincent Vanhoucke. Youtube-boundingboxes: A large high-precision human-annotated data set for object detection in video. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 5296–5305, 2017.
- [32] Shai Shalev-Shwartz et al. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.
- [33] Linda B Smith. Learning to recognize objects. *Psychological Science*, 14(3):244–250, 2003.
- [34] Linda B Smith and Susan S Jones. Symbolic play connects to language through visual object recognition. *Developmental Science*, 14(5):1142–1149, 2011.
- [35] Yibing Song, Chao Ma, Xiaohe Wu, Lijun Gong, Linchao Bao, Wangmeng Zuo, Chunhua Shen, Rynson WH Lau, and Ming-Hsuan Yang. Vital: Visual tracking via adversarial learning. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 8990–8999, 2018.
- [36] Ardhendu Shekhar Tripathi, Martin Danelljan, Luc Van Gool, and Radu Timofte. Tracking the known and the unknown by leveraging semantic information. In *BMVC*, volume 2, page 6, 2019.
- [37] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 3156–3164, 2015.
- [38] Paul Voigtlaender, Jonathon Luiten, Philip HS Torr, and Bastian Leibe. Siam R-CNN: Visual tracking by re-detection. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 6578–6588, 2020.
- [39] Yi Wu, Jongwoo Lim, and Ming-Hsuan Yang. Online object tracking: A benchmark. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [40] Zhengyuan Yang, Boqing Gong, Liwei Wang, Wenbing Huang, Dong Yu, and Jiebo Luo. A fast and accurate one-stage approach to visual grounding. In *Proc. IEEE International Conf. on Computer Vision (ICCV)*, pages 4683–4693, 2019.
- [41] Zheng Zhu, Qiang Wang, Bo Li, Wei Wu, Junjie Yan, and Weiming Hu. Distractor-aware siamese networks for visual object tracking. In *Proc. European Conf. on Computer Vision (ECCV)*, pages 101–117, 2018.