Divert More Attention to Vision-Language Tracking

Anonymous Author(s)
Affiliation
Address
email

Abstract

Relying on Transformer for complex visual feature learning, object tracking has witnessed the new standard for state-of-the-arts (SOTAs). However, this advancement accompanies by larger training data and longer training period, making tracking increasingly expensive. In this paper, we demonstrate that the Transformer-reliance is not necessary and the pure ConvNets are still competitive and even better yet more economical and friendly in achieving SOTA tracking. Our solution is to unleash the power of multimodal vision-language (VL) tracking, simply using ConvNets. The essence lies in learning novel unified-adaptive VL representations with our modality mixer (ModaMixer) and asymmetrical ConvNet search. We show that our unified-adaptive VL representation, learned purely with the ConvNets, is a simple yet strong alternative to Transformer visual features, by unbelievably improving a CNN-based Siamese tracker by 14.5% in SUC on challenging LaSOT (**50.7%**$\rightarrow$**65.2%**), even outperforming several Transformer-based SOTA trackers. Besides empirical results, we theoretically analyze our approach to evidence its effectiveness. By revealing the potential of VL representation, we expect the community to divert more attention to VL tracking and hope to open more possibilities for future tracking beyond Transformer. Code and models will be released.

1 Introduction

Transformer tracking recently receives a surge of research interests and becomes almost a necessity to achieve state-of-the-art (SOTA) performance [6, 49, 8]. The success of Transformer trackers mainly attributes to *attention* that enables complex feature interactions. But, is this complex attention the only way realizing SOTA tracking? Or in other words, ***is Transformer the only path to SOTA?***

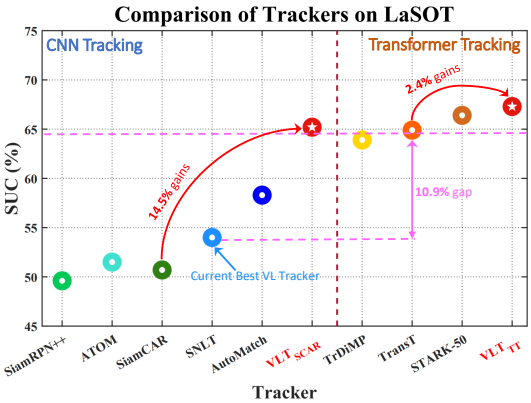

Figure 1: Comparison between CNN-based and Transformer-based trackers on LaSOT [15].

We answer **no**, and display a *Transformer-free* path using **pure** convolutional neural network (CNN). Different than complex interactions in visual feature by attention requiring more training data and longer training time, our alternative is to explore simple interactions of multimodal, *i.e.*, vision and language, through CNN. In fact, language, an equally important cue as vision, has been largely explored in vision-related tasks, and is not new to tracking. Prior works [18, 31, 17]

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

have exploited vision-language (VL) multimodal learning for improving tracking. However, the performance falls far behind current SOTAs. For instance on LaSOT [15], the gap between current best VL tracker [18] and recent Transformer tracker [6] is absolute 10.9% in SUC (see Fig. 1). ***So, what is the bottleneck of VL tracking in achieving SOTA?***

The devil is in VL representation. Feature representation has been shown to be crucial in improving tracking [48, 58, 29, 22, 8]. Given two modalities of vision and language, the VL feature is desired to be *unified* and *adaptive* [37, 55]. The former property requires deep interaction of vision and language, while the latter needs VL feature to accommodate different scenarios of visual and linguistic information. However, in existing VL trackers, vision and language are treated *independently* and processed *distantly* until the final result fusion. Although this fusion may easily get two modalities connected, it does not accords with human learning procedure that integrates multisensory by various neurons before causal inference [3], resulting in a lower upper-bound for VL tracking. Besides, current VL trackers treat template and search branches as *homoplasmic* inputs, and adopt symmetrical feature learning structures for these two branches, inherited from typical vision-only Siamese tracking [18]. We argue the mixed modality may have different intrinsic nature than the pure vision modality, and thus requires a more flexible and general design for different signals.

Our solution. Having observed the above, we introduce a novel unified-adaptive vision-language representation, aiming for SOTA VL tracking ***without using Transformer***[1]. Specifically, we first present modality mixer, or ModaMixer, a conceptually simple but effective module for VL interaction. Language is a high-level representation and its class embedding can help distinguish targets of different categories (*e.g.*, cat and dog) and meanwhile the attribute embedding (*e.g.*, color, shape) provides strong prior to separate targets of same class (*e.g.*, cars with different colors). The intuition is, channel features in vision representation also reveal semantics of objects [20, 53]. Inspired by this, ModaMixer regards language representation as a selector to reweight different channels of visual features, enhancing target-specific channels as well as suppressing irrelevant both intra- and inter-class channels. The selected feature is then fused with the original feature, using a special asymmetrical design (analyzed later in experiments), to generate the final unified VL representation. A set of ModaMixers are installed in a typical CNN from shallow to deep, boosting robustness and discriminability of the unified VL representation at different semantic levels. Despite simplicity, ModaMixer brings 6.9% gains over a pure CNN baseline [21] (*i.e.*, 50.7%→57.6%).

Despite huge improvement, the gap to SOTA Transformer tracker [6] remains (57.6% *v.s.* 64.9%). To mitigate the gap, we propose an asymmetrical searching strategy (ASS) to adapt the unified VL representation for improvements. Different from current VL tracking [18] adopting symmetrical and fixed template and search branches as in vision-only Siamese tracking [29], we argue that the learning framework of mixed modality should be adaptive and not fixed. To this end, ASS borrows the idea from neural architecture search (NAS) [60, 39] to separately learn *distinctive* and *asymmetrical* networks for mixed modality in different branches and ModaMixers. The asymmetrical architecture, to our best knowledge, is the first of its kind in matching-based tracking. Note, although NAS has been adopted in matching-based tracking [52], this method finds *symmetrical* networks for single modality. Differently, ASS is applied on mixed modality and the resulted architecture is *symmetrical*. Moreover, the network searched in ASS avoids burdensome re-training on ImageNet [12], enabling quick reproducibility of our work (only 0.625 GPU days with a single RTX-2080Ti). Our ASS is general and flexible, and together with ModaMixer, it surprisingly shows additional 7.6% gains (*i.e.*, 57.6%→65.2%), evidencing our argument and effectiveness of ASS.

Eventually, with the unified-adaptive representation, we implement the first pure CNN-based VL tracker that shows SOTA results comparable and even better than Transformer-based solutions, without bells and whistles. Specifically, we apply our method to a CNN baseline SiamCAR [21], and the resulted VL tracker VLT_{SCAR} shows 65.2% SUC on LaSOT [15] while running at 43FPS, unbelievably improving the baseline by 14.5% and outperforming SOTA Transformer trackers [6, 49] (see again Fig. 1). We observe similar improvements by our approach on other four benchmarks.

[1]Here we stress that we do not use Transformer for visual feature learning as in current Transformer trackers or for multimodal learning. We only use it in language embedding extraction (*i.e.*, BERT [13])

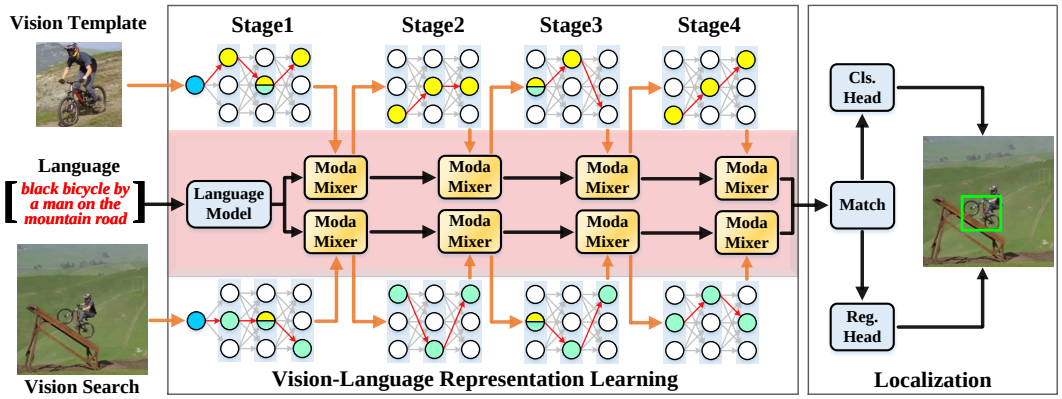

Figure 2: The proposed vision-language tracking framework. The semantic information of language description is injected to vision from shallow to deep layers of the asymmetrical modeling architecture to learn unified-adaptive vision-language representation.

Besides empirical results, we provide theoretical analysis to evidence the effectiveness of our method. Note that, our approach is general in improving vision-only trackers including Transformer-based ones. We show this by applying it to TransT [6] and the resulted tracker VLT_{TT} shows 2.4% SUC gains (*i.e.*, 64.9%→67.3%), evidencing its effectiveness and generality.

We are aware that one can certainly leverage the Transformer [36] to learn a good (maybe better) VL representation for tracking, with larger data and longer training period. Different than this, our goal is to explore a cheaper way with simple architectures such as pure CNN for SOTA tracking performance and open more possibilities for future tracking beyond Transformer. In summary, our **contributions** are four-fold: **(i)** we introduce a novel unified-adaptive vision-language representation for SOTA VL tracking; **(ii)** we propose the embarrassingly simple yet effective ModaMixer for unified VL representation learning; **(iii)** we present ASS to adapt mixed VL representation for better tracking and **(iv)** using pure CNN architecture, we achieve SOTA results on multiple benchmarks .

2 Related Work

Visual Tracking. Tracking has witnessed great progress in the past decades. Particularly, Siamese tracking [4, 44], that aims to learn a generic matching function, is a representative branch and has revolutionized with numerous extensions [30, 58, 29, 54, 16, 23, 57, 59, 7]. Recently, Transformer [46] has been introduced to Siamese tracking for better interactions of visual features and greatly pushed the standard of state-of-the-art performance [6, 49, 43, 8, 33]. From a different perspective than using complex Transformer, we explore multimodal with simple CNN to achieve SOTA tracking.

Vision-Language Tracking. Natural language contains high-level semantics and has been leveraged to foster vision-related tasks [19, 28, 2] including tracking [31, 17, 18]. The work [31] first introduces linguistic description to tracking and shows that language enhances the robustness of vision-based method. Most recently, SNLT [18] integrates linguistic information into Siamese tracking by fusing results respectively obtained by vision and language. Different from these VL trackers that regard vision and language as independent cues with weak connections only at result fusion, we propose ModaMixer to unleash the power of VL tracking by learning unified VL representation.

NAS for Tracking. Neural architecture search (NAS) aims at finding the optimal design of deep network architectures [60, 39, 35, 24] and has been introduced to tracking [52, 57]. LightTrack [52] tends to search a lightweight backbone but is computationally demanding (about 40 V100 GPU days). AutoMatch uses DARTS [35] to find better matching networks for Siamese tracking. All these methods leverage NAS for vision-only tracking and search a *symmetrical* Siamese architecture. Differently, our work searches the network for multimodal tracking and tries to find a more general and flexible *asymmetrical* two-stream counterpart. In addition, our search pipeline only takes 0.625 RTX-2080Ti GPU days, which is much more resource-friendly.

3 Unified-Adaptive Vision-Language Tracking

This section details our unified-adaptive vision-language (VL) tracking as shown in Fig. 2. In specific, we first describe the proposed modality mixer for generating unified multimodal representation and then asymmetrical network which searches for learning adaptive VL representation. Afterwards, we illustrate the proposed tracking framework, followed by theoretical analysis of our method.

3.1 Modality Mixer for Unified Representation

The essence of multimodal learning is a simple and effective modality fusion module. As discussed before, existing VL trackers simply use a *later fusion* way, in which different modalities are treated independently and processed distantly until merging their final results [18, 31]. Despite the effectiveness to some extent, the complementarity of different modalities in representation learning is largely unexplored, which may impede the multimodal learning to unleash its power for VL tracking. In this work, we propose the modality mixer (dubbed **ModaMixer**) to demonstrate a compact way to learn a unified vision-language representation for tracking.

ModaMixer considers language representation as selector to reweight channels of vision features. In specific, given the language description with N words of a video[2], a language model [13] is adopted to abstract the sentence to semantic features with size of $(N+2) \times d$. The extra "2" denotes the "[CLS][SEP]" characters in language model processing (see [13] for more details). Notably, descriptions for different videos may contain various length N. To ensure the ModaMixer applicable for all videos, we first average the features for all words along sequence length dimension "(N+2)" to generate a unique language representation $\mathbf{f}_l \in \mathbb{R}^{1\times d}$ for each description. Then a linear layer is followed to align the channel number of $\mathbf{f}_l$ with the corresponding vision feature $\mathbf{f}_v \in \mathbb{R}^{H\times W\times C}$. Channel selector is expressed as Hadamard product operator, which point-wisely multiplies language representation $(1 \times C)$ to embedding of each spatial position in the vision feature $\mathbf{f}_v$. Finally, a residual connection between the mixed feature $\mathbf{f}_m$ and vision feature $\mathbf{f}_v$ is conducted to avoid losing informative vision details. In a nutshell, the ModaMixer can be formulated as,

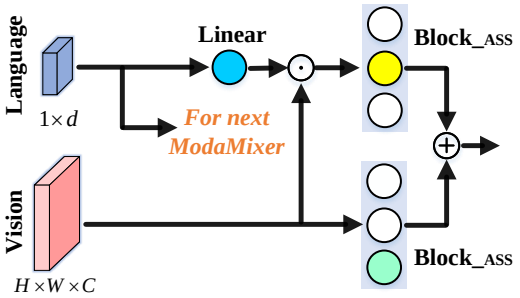

Figure 3: Illustration of the ModaMixer.

$$\mathbf{f}_m = \mathrm{Block}_{\mathrm{ASS}}\left(\mathrm{Linear}(\mathbf{f}_l) \odot \mathbf{f}_v\right) + \mathrm{Block}_{\mathrm{ASS}}\left(\mathbf{f}_v\right), \tag{1}$$

where $\odot$ denotes Hadamard product, Linear is a linear projection layer with weight matrix size of $d \times C$ for channel number alignment, and $\mathrm{Block}_{\mathrm{ASS}}$ indicates post-processing block before residual connection. Please note that, to enable adaptive feature modeling for different modalities, we search different $\mathrm{Block}_{\mathrm{ASS}}$ to process features before and after fusion (see Sec. 3.2 for more details). The proposed ModaMixer is illustrated in Fig. 3. Akin to channel attention [41, 20], the high-level semantics in language representation dynamically enhance target-specific channels in vision features, and meanwhile suppress the responses of distractors belonging to both inter- and intra-classes.

3.2 Asymmetrical Search for Adaptive Vision-Language Representation

Besides the fusion module, the other crucial key for vision-language tracking is how to construct the basic modeling structure. The simplest strategy is to inherit a symmetrical Siamese network from vision-based tracking (*e.g.*, [4, 29]), as in current VL trackers [18]. But the performance gap still remains if using this manner, which is mostly blamed on the neglect of the different intrinsic nature between VL-based multimodal and vision-only single modality. To remedy this, we propose

[2]The language description in tracking is generated *only* by the initial target object in the first frame.

Table 1: The asymmetrical architecture learned by ASS. 🔵 is the stem convolution layer. (🔴🟠🟢🔵) represents the basic ASS unit, where the first three ones indicate *Shuffle block* [56] with kernel sizes of (3,5,7), respectively, and the last one denotes a *Shuffle Xception block* [56] with kernel size of 3.

	Stem	**Stage1**	**Moda Mixer**	**Stage2**	**Moda Mixer**	**Stage3**	**Moda Mixer**	**Stage4**	**Moda Mixer**
Template	🔵	🟢🟢🟢	🔴🔵	🔵🟠🔵	🔵🔴	🔵🔴🟢 🔵🔴🟢🟢	🟢🔵	🟢🔵🟢	🟠🔴
Search	🔵	🟠🔵🔴	🟢🔵	🟠🟢🟢	🔵🔴	🟠🟢🔵 🔴🔵🟠🟢	🔵🔵	🔵🔵🟢	🔴🔴

an asymmetrical searching strategy (dubbed **ASS**) to learn an adaptive modeling structure for pairing with ModaMixer.

The spirits of network search are originated from the field of Neural Architecture Search (NAS). We adopt a popular NAS model, in particular the single-path one-shot method SPOS [24], for searching the optimal structure of our purpose. Although SPOS has been utilized for tracking [52], our work significantly differs from it from two aspects: **1)** Our ASS is tailored for constructing an ***asymmetrical*** two-stream network for ***multimodal*** tracking, while [52] is designed to find a ***symmetrical*** Siamese network for vision-only ***single-modality*** tracking. Besides, we search layers both in the backbone network and the post-processing $\mathrm{Block}_{\mathrm{ASS}}$ in the ModaMixer (see Eq. 1); **2)** Our ASS reuses the pre-trained supernet from SPOS, which avoids the burdensome re-training on ImageNet [12] (both for the supernet and found subnet) and thus reduces the time complexity of our search pipeline to $1/64$ of that in LightTrack [52] (*i.e.*, **0.625 RTX-2080Ti GPU days *v.s.* 40 V100 GPU days**). Due to limited space, please refer to **appendix** for more details and comparison of our ASS and [52].

The search space and search strategy of ASS are kept consistent with the original SPOS [24]. In particular, the search pipeline is formulated as,

$$W_{\mathcal{A}} = \underset{W}{\operatorname{argmin}} \; \mathbb{E}_{a \sim \Gamma(\mathcal{A})} \left[\mathcal{L}_{\text{train}}(\mathcal{N}(a, W(a))) \right], \tag{2}$$

$$a^{*} = \underset{a \in \mathcal{A}}{\operatorname{argmax}} \; \mathrm{SUC}_{\text{val}} \left(\mathcal{N}(a, W_{\mathcal{A}}(a)) \right), \tag{3}$$

where $\mathcal{A}$ represents the architecture search space of the network $\mathcal{N}$, a is a sample from $\mathcal{A}$ and W denotes the corresponding network weights. Notably, the network $\mathcal{N}$ includes three components $\mathcal{N} = \{\varphi_t, \varphi_s, \varphi_m\}$, where each indicates backbone for the template branch φ_t, backbone for the search branch φ_s and layers in the ModaMixer φ_m. The whole pipeline consists of training supernet on tracking datasets via random sampling Γ from search space $\mathcal{A}$ (Eq. 2) and finding the optimal subnet via evolutionary algorithms (Eq. 3). The SUC (success score) on validation data is used as rewards of evolutionary algorithms. Tab. 1 demonstrates the searched asymmetrical networks in our VL tracking. For more details of ASS, we kindly refer readers to **appendix** or [24].

3.3 Tracking Framework

With the proposed ModaMixer and the searched asymmetrical networks, we construct a new vision-language tracking framework, as shown in Fig. 2 and Tab. 1. Our framework is matching-based tracking. Both template and search backbone networks contain 4 stages with the maximum stride of 8, the chosen blocks of each stage are denoted with different colors in Tab. 1. ModaMixer is integrated into each stage of the template and search networks to learn informative mixed representation. It is worth noting that, the asymmetry is revealed in not only the design of backbone networks, but also the ModaMixer. Each ModaMixer shares the same meta-structure as in Fig. 3, but comprises different post-processing layers $\mathrm{Block}_{\mathrm{ASS}}$ to allow adaption to different semantic levels (*i.e.*, network depth) and input signals (*i.e.*, template and search, pure-vision and mixed feature in each ModaMixer). With the learned unified-adaptive VL representations from the template and search branches, we perform feature matching and target localization, the same as in our baseline.

3.4 A Theoretical Explanation

This section presents a theoretical explanation of our method. Based on the Empirical Risk Minimization (ERM) principle [42], the objective of representation learning is to find better network parameters θ by minimizing the empirical risk, *i.e.*,

$$\min \quad \hat{r}\left(\theta_{\mathcal{M}}\right) \triangleq \frac{1}{n} \sum_{i=1}^{n} \mathcal{L}\left(\mathcal{X}_i, y_i; \theta_{\mathcal{M}}\right) \qquad \text{s.t.} \;\; \theta_{\mathcal{M}} \in \mathcal{F}. \tag{4}$$

where $\mathcal{L}$ denotes loss function, $\mathcal{M}$ represents the modality set, n indicates sample number, $\mathcal{X}_i = \{x_i^1, x_i^2 ... x_i^{|\mathcal{M}|}\}$ is the input mutimodal signal, y_i is training label, and $\mathcal{F}$ demotes optimization space of θ. Given the empirical risk $\hat{r}\left(\theta_{\mathcal{M}}\right)$, its corresponding population risk is defined as,

$$r\left(\theta_{\mathcal{M}}\right) = \mathbb{E}_{(\mathcal{X}_i, y_i) \sim \mathcal{D}_{train}}\left[\hat{r}\left(\theta_{\mathcal{M}}\right)\right] \tag{5}$$

Similar to [1, 45], we use population risk to measure the learning quality. Then the **latent representation quality** [27] is defined as,

$$\eta(\theta) = \inf_{\theta \in \mathcal{F}}\left[r\left(\theta\right) - r(\theta^*)\right] \tag{6}$$

where $*$ represents the optimal case, $\inf$ indicates the best achievable population risk. With the empirical Rademacher complexity $\mathfrak{R}$ [27], we present the main result regarding multi-modal learning.

Theorem 1. *Assuming we have produced the empirical risk minimizers $\hat{\theta}_{\mathcal{M}}$ and $\hat{\theta}_{\mathcal{S}}$, training with the $|\mathcal{M}|$ and $|\mathcal{S}|$ modalities separately ($|\mathcal{M}| > |\mathcal{S}|$). Then, for all $1 > \delta > 0$, with probability at least $1 - \frac{\delta}{2}$:*

$$r\left(\hat{\theta}_{\mathcal{M}}\right) - r\left(\hat{\theta}_{\mathcal{S}}\right) \leq \gamma_{\mathcal{D}}(\mathcal{M}, \mathcal{S}) + 8L\mathfrak{R}_m(\mathcal{F}_{\mathcal{M}}) + \frac{4C}{\sqrt{n}} + 2C\sqrt{\frac{2\ln(2/\delta)}{n}} \tag{7}$$

where

$$\gamma_{\mathcal{D}}(\mathcal{M}, \mathcal{S}) \triangleq \eta(\hat{\theta}_{\mathcal{M}}) - \eta(\hat{\theta}_{\mathcal{S}}) \qquad \mathfrak{R}_m(\mathcal{F}_{\mathcal{M}}) \sim \sqrt{Complexity(\mathcal{F}_{\mathcal{M}})/n} \tag{8}$$

$\gamma_{\mathcal{D}}(\mathcal{M}, \mathcal{S})$ defined in (8) compares the quality between latent representations learned from $\mathcal{M}$ and $\mathcal{S}$ modalities with respect to the given dataset $\mathcal{D}$. Theorem 1 bounds the population risk of training with two different number of modalities, **which proves that more modalities could enhance the representation quality.** Furthermore, for the commonly used function classes in machine learning, Rademacher complexity $\mathfrak{R}_m(\mathcal{F})$ for a sample size n is bounded by $\sqrt{Complexity(\mathcal{F})/n}$, where $Complexity(\mathcal{F})$ represents the complexity of a network. This demonstrates that heterogeneous network would theoretically rise the upper bound of $r\left(\hat{\theta}_{\mathcal{M}}\right) - r\left(\hat{\theta}_{\mathcal{S}}\right)$, **which exhibits that our asymmetrical design has larger optimization space when learning with $|\mathcal{M}|$ modalities compared to $|\mathcal{S}|$ modalities ($|\mathcal{M}| > |\mathcal{S}|$).** The proof is beyond our scoop, and please refer to [27] for details.

4 Experiment

4.1 Implementation Details

We apply our method to both CNN-based SiamCAR [21] (dubbed $\mathbf{VLT}_{\mathrm{SCAR}}$) and Transformer-based TransT [6] (dubbed $\mathbf{VLT}_{\mathrm{TT}}$). The matching module and localization head are inherited from the baseline tracker without any modifications.

Searching for VLT. The proposed ASS aims to find a more flexible modeling structure for vision-language tracking (VLT). Taking $\mathrm{VLT}_{\mathrm{SCAR}}$ as example, the supernet from SPOS [24] is used as feature extractor to replace the ResNet [25] in SiamCAR. We train the trackers with supernet using training splits of COCO [34], Imagenet-VID [12], Imagenet-DET [12], Youtube-BB [40], GOT-10k [26], LaSOT [15] and TNL2K [50] for 5 epochs, where each epoch contains 1.2×10^6

Table 2: State-of-the-art comparisons on LaSOT [15], $\text{LaSOT}_{\text{Ext}}$ [14], TNL2K [50], GOT-10k [26] and OTB99-LANG (OTB99-L) [32]. TransT and SiamCAR are baselines of the proposed $\textbf{VLT}_{\text{TT}}$ and $\textbf{VLT}_{\text{SCAR}}$, respectively. All metrics of performance are in % in tables unless otherwise specified.

Type	Method	LaSOT		$\text{LaSOT}_{\text{Ext}}$		TNL2K		GOT-10k			OTB99-L	
		SUC	P	SUC	P	SUC	P	AO	$\text{SR}_{0.5}$	$\text{SR}_{0.75}$	SUC	P
CNN-based	SiamRCNN [47]	64.8	68.4	-	-	52.3	52.8	64.9	72.8	59.7	70.0	89.4
	PrDiMP [11]	59.8	60.8	-	-	47.0	45.9	63.4	73.8	54.3	69.5	89.5
	AutoMatch [57]	58.3	59.9	37,6	43.0	47.2	43.5	65.2	76.6	54.3	71.6	93.2
	Ocean [59]	56.0	56.6	-	-	38.4	37.7	61.1	72.1	47.3	68.0	92.1
	KYS [5]	55.4	-	-	-	44.9	43.5	63.6	75.1	51.5	-	-
	ATOM [10]	51.5	50.5	37.6	43.0	40.1	39.2	55.6	63.4	40.2	67.6	82.4
	SiamRPN++ [29]	49.6	49.1	34.0	39.6	41.3	41.2	51.7	61.6	32.5	63.8	82.6
	C-RPN [16]	45.5	42.5	27.5	32.0	-	-	-	-	-	-	-
	SiamFC [4]	33.6	33.9	23.0	26.9	29.5	28.6	34.8	35.3	9.8	58.7	79.2
	ECO [9]	32.4	30.1	22.0	24.0	-	-	31.6	30.9	11.1	-	-
	SiamCAR [21]	50.7	51.0	33.9	41.0	35.3	38.4	56.9	67.0	41.5	68.8	89.1
CNN-VL	SNLT [18]	54.0	57.6	26.2	30.0	27.6	41.9	43.3	50.6	22.1	66.6	80.4
	$\textbf{VLT}_{\text{SCAR}}$ (Ours)	**65.2**	**69.1**	**44.7**	**51.6**	**49.8**	**51.0**	**61.4**	**72.4**	**52.3**	**73.9**	**89.8**
Trans-based	STARK [51]	66.4	71.2	47.8	55.1	-	-	68.0	77.7	62.3	69.6	91.4
	TrDiMP [49]	63.9	66.3	-	-	-	-	67.1	77.7	58.3	70.5	92.5
	TransT [6]	64.9	69.0	44.8	52.5	50.7	51.7	67.1	76.8	60.9	70.8	91.2
Trans-VL	$\textbf{VLT}_{\text{TT}}$ (Ours)	**67.3**	**72.1**	**48.4**	**55.9**	**53.1**	**53.3**	**69.4**	**81.1**	**64.5**	**76.4**	**93.1**

template-search pairs. Once finishing supernet training, evolutionary algorithms as in SPOS [24] is applied to search for optimal subnet and finally obtains VLT_{SCAR}. The whole search pipeline consumes 15 hours on a single RTX-2080Ti GPU. The search process of VLT_{TT} is similar to VLT_{SCAR}. We present more details in the **appendix** due to space limitation.

Optimizing $\textbf{VLT}_{\text{SCAR}}$ and $\textbf{VLT}_{\text{TT}}$. The training protocol of VLT_{SCAR} and VLT_{TT} follows the corresponding baselines SiamCAR [21] and TransT [6]. Notably, for each epoch, half training pairs come from datasets without language annotations (*i.e.*, COCO [34], Imagenet-VID [12], Imagenet-DET [12], Youtube-BB [40]). The language representation is set as 0-tensor or visual pooling feature under this circumstances (discussed in Tab. 4)[3].

4.2 State-of-the-art Comparison

Tab. 2 presents the results and comparisons of our trackers with other SOTAs on LaSOT [15], $\text{LaSOT}_{\text{Ext}}$ [14], TNL2K [50], OTB99-LANG [32] and GOT-10K [26]. The proposed VLT_{SCAR} and VLT_{TT} run at 43/35 FPS on a single RTX-2080Ti GPU, respectively. Compared with the speeds of baseline trackers SiamCAR [21]/TransT [6] with 52/32 FPS, the computation cost of our method is small. Moreover, our VLT_{TT} outperforms TransT in terms of both accuracy and speed.

Compared with SiamCAR [21], VLT_{SCAR} achieves considerable SUC gains of 14.5%/10.8%/14.5% on LaSOT/$\text{LaSOT}_{\text{Ext}}$/TNL2K, respectively, which demonstrates the effectiveness of the proposed VL tracker. Notably, our VLT_{SCAR} outperforms the current best VL tracker SNLT [18] for 11.2%/18.5% on LaSOT/$\text{LaSOT}_{\text{Ext}}$, showing that the unified-adaptive vision-language representation is more robust for VL tracking and is superior to simply fusing tracking results of different modalities. The advancement of our method is preserved across different benchmarks. What surprises us more is that the CNN-based VLT_{SCAR} is competitive and even better than recent vision Transformer-based approaches. For example, VLT_{SCAR} outperforms TransT [6] on LaSOT and meanwhile runs faster (43 FPS *v.s.* 32 FPS) and requires less training pairs (2.4×10^7 *v.s.* 3.8×10^7). By applying our method to TransT, the new tracker VLT_{TT} improves the baseline to 67.3% in SUC with 2.4% gains on LaSOT while being faster, showing its generality.

[3] GOT-10k [26] provides simple descriptions for object/motion/major/root class, *e.g.*, "dove, walking, bird, animal", in each video. We concatenate these words to obtain a pseudo language description.

Table 3: Ablation on ModaMixer and asymmetrical searching strategy (ASS).

#	Method	ModaMixer	ASS	LaSOT			TNL2K		
				SUC	P_{Norm}	P	SUC	P_{Norm}	P
①	Baseline	-	-	50.7	60.0	51.0	35.3	43.6	38.4
②	**VLT**$_{\mathrm{SCAR}}$	√	-	57.6	65.8	61.1	41.5	49.2	43.2
③	**VLT**$_{\mathrm{SCAR}}$	-	√	52.1	59.8	50.6	40.7	47.2	40.2
④	**VLT**$_{\mathrm{SCAR}}$	√	√	**65.2**	**74.9**	**69.1**	**48.3**	**55.2**	**46.6**

Table 4: Comparing different strategies to handle training videos without language annotation.

Method	Setting	LaSOT		LaSOT$_{\mathrm{Ext}}$		TNL2K		GOT-10k			OTB99-L	
		SUC	P	SUC	P	SUC	P	AO	$SR_{0.5}$	$SR_{0.75}$	SUC	P
VLT$_{\mathrm{SCAR}}$	**0-tensor**	65.2	69.1	41.2	47.5	48.3	46.6	61.4	72.4	52.3	72.7	88.8
	template	63.9	67.9	44.7	51.6	49.8	51.1	61.0	70.8	52.2	73.9	89.8

Table 5: Evaluating different settings on LaSOT: (a) the influence of symmetrical and our asymmetrical design, (b) adopting fixed ShuffleNet block or searching the post-processing block in ModaMoxer, and (c) removing language description or using caption generated by [38] during inference.

(a)

Settings	SUC	P
symmetrical	60.0	61.7
asymmetrical	63.9	67.9

(b)

Settings	SUC	P
Shuffle-ModaMixer	59.1	62.2
NAS-ModaMixer	63.9	67.9

(c)

Settings	SUC	P
w/o. language	53.4	54.6
Pse. language	51.6	53.4

4.3 Component-wise Ablation

We analyze the influence of each component in our method to show the effectiveness and rationality of the proposed ModaMixer and ASS. The ablation experiments are conducted on VLT$_{\mathrm{SCAR}}$, and results are presented in Tab. 3. By directly applying the ModaMixer on the baseline SiamCAR [21] ("ResNet50+ModaMixer"), it obtains SUC gains of 6.9% on LaSOT (②*v.s.*①). This verifies that the unified VL representation effectively improves tracking robustness. One interesting observation is that ASS improves vision-only baseline for 1.4% percents on LaSOT (③*v.s.*①), but when equipping with ModaMixer, it surprisingly further brings 7.6% SUC gains (④*v.s.*②), which shows the complementarity of multimodal representation learning (ModaMixer) and the proposed ASS.

4.4 Further Analysis

Dealing with videos without language description during training. As mentioned above, language annotations are not provided in several training datasets (*e.g.*, YTB-BBox [40]). We design two strategies to handle that. One is to use "0-tensor" as language embedding, and the other is to replace the language embedding with visual features which are generated by pooling template feature in the bounding box. As shown in Tab. 4, the two strategies perform competitively, but the one with visual feature is slightly better in average. Therefore, in the following, we take VLT$_{\mathrm{SCAR}}$ trained by replacing language embedding with template visual embedding as the baseline for further analysis.

Symmetrical or Asymmetrical? The proposed asymmetrical searching strategy is essential for achieving an adaptive vision-language representation. As illustrated in Tab. 5a, we experiment by searching for a symmetrical network (including both backbone and $\mathrm{Block}_{\mathrm{ASS}}$ in the ModaMixer), but it is inferior to the asymmetrical counterpart for 3.9%/6.2% of success rate (SUC) and precision (P) on LaSOT [15], respectively, which empirically proves our argument.

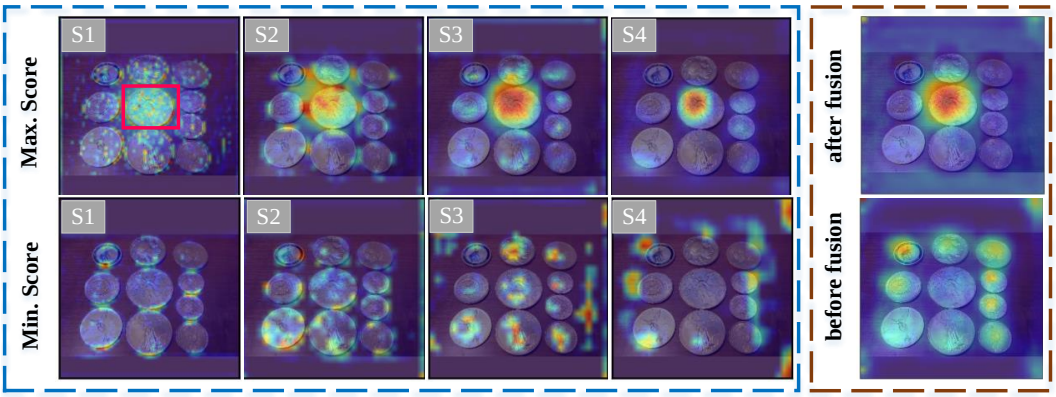

(a) Channels visualization with max/min selection scores **(b) Activation maps**

Figure 4: (a) feature channel with maximum/minimum (top/bottom) selection scores from ModaMixer in stage1-4. (b) activation map before/after (top/bottom) multimodal fusion in ModaMixer.

Asymmetry in ModaMixer. The asymmetry is used in not only the backbone network, but also the ModaMixer. In our work, the post-processing layers for different signals (visual and mixed features) are decided by ASS, which enables the adaption at both semantic levels (*i.e.*, network depth) and different input signals (*i.e.*, template and search, pure-vision and mixed feature in each ModaMixer). As in Tab. 5b, when replacing the post-processing layers with a fixed ShuffleNet block from SPOS [24] (*i.e.*, inheriting structure and weights from the last block in each backbone stage), the performance drops from 63.9% to 59.1% in SUC on LaSOT. This reveals that the proposed ASS is important for building a better VL learner.

No/Pseudo description during inference. VL trackers require the first frame of a video is annotated with a language description. One may wonder that what if there is no language description? Tab. 5c presents the results by removing the description and using that generated with an recent advanced image-caption method [38] (in ICML2021). The results show that, without language description, performance heavily degrades (65.2% $\rightarrow$ 53.4% SUC on LaSOT), verifying that the high-level semantics in language do help in improving robustness. Even though, the performance is still better than the vision-only baseline. Surprisingly, when using the generated description, it doesn't show promising results (51.6%), indicating that it is still challenging to generate accurate caption in real-world cases and noisy caption even brings negative effects to the model.

Channel selection by ModaMixer. ModaMixer translates the language description to a channel selector to reweight visual features. As shown in Fig. 4, the channel activation maps with maximum selection scores always correspond to the target, while the surrounding distractors are successfully assigned with minimum scores (Fig. 4 (a)-bottom). Besides, with multimodal fusion (or channel selection), the network can enhance the response of target and meanwhile suppress the distractors (see Fig. 4 (b)). This evidences our argument that language embedding can identify semantics in visual feature channels and effectively select useful information for localizing targets. More visualization results are presented in **appendix** due to limited space.

5 Conclusion

In this work, we explore a different path to achieve SOTA tracking without complex Transformer, *i.e.,* multimodal VL tracking. The essence is a unified-adaptive VL representation, learned by our ModaMixer and asymmetrical networks. In experiments, our approach surprisingly boosts a pure CNN-based Siamese tracker to achieve competitive or even better performances compared to recent SOTAs. Besides, we provide an theoretical explanation to evidence the effectiveness of our method. We hope that this work inspires more possibilities for future tracking beyond Transformer.

References

[1] Massih R Amini, Nicolas Usunier, and Cyril Goutte. Learning from multiple partially observed views-an application to multilingual text categorization. *Advances in Neural Information Processing Systems*, 2009.

[2] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.

[3] Stephanie Badde, Fangfang Hong, and Michael S Landy. Causal inference and the evolution of opposite neurons. *Proceedings of the National Academy of Sciences*, 2021.

[4] Luca Bertinetto, Jack Valmadre, João F Henriques, Andrea Vedaldi, and Philip H S Torr. Fully-convolutional siamese networks for object tracking. In *European Conference on Computer Vision Workshops*, 2016.

[5] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Know your surroundings: Exploiting scene information for object tracking. In *European Conference on Computer Vision*, 2020.

[6] Xin Chen, Bin Yan, Jiawen Zhu, Dong Wang, Xiaoyun Yang, and Huchuan Lu. Transformer tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.

[7] Zedu Chen, Bineng Zhong, Guorong Li, Shengping Zhang, and Rongrong Ji. Siamese box adaptive network for visual tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.

[8] Yutao Cui, Jiang Cheng, Limin Wang, and Gangshan Wu. Mixformer: End-to-end tracking with iterative mixed attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.

[9] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. ECO: Efficient convolution operators for tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017.

[10] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. ATOM: Accurate tracking by overlap maximization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.

[11] Martin Danelljan, Luc Van Gool, and Radu Timofte. Probabilistic regression for visual tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.

[12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2009.

[13] 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.

[14] Heng Fan, Hexin Bai, Liting Lin, Fan Yang, Peng Chu, Ge Deng, Sijia Yu, Mingzhen Huang, Juehuan Liu, Yong Xu, et al. Lasot: A high-quality large-scale single object tracking benchmark. *International Journal of Computer Vision*, 2021.

[15] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.

[16] Heng Fan and Haibin Ling. Siamese cascaded region proposal networks for real-time visual tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.

[17] Qi Feng, Vitaly Ablavsky, Qinxun Bai, Guorong Li, and Stan Sclaroff. Real-time visual object tracking with natural language description. In *IEEE Winter Conference on Applications of Computer Vision*, 2020.

[18] Qi Feng, Vitaly Ablavsky, Qinxun Bai, and Stan Sclaroff. Siamese natural language tracker: Tracking by natural language descriptions with siamese trackers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.

[19] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *Conference on Empirical Methods in Natural Language Processing*, 2016.

[20] Xitong Gao, Yiren Zhao, Łukasz Dudziak, Robert Mullins, and Cheng-zhong Xu. Dynamic channel pruning: Feature boosting and suppression. *arXiv preprint arXiv:1810.05331*, 2018.

[21] Dongyan Guo, Jun Wang, Ying Cui, Zhenhua Wang, and Shengyong Chen. SiamCAR: Siamese fully convolutional classification and regression for visual tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.

[22] Mingzhe Guo, Zhipeng Zhang, Heng Fan, Liping Jing, Yilin Lyu, Bing Li, and Weiming Hu. Learning target-aware representation for visual tracking via informative interactions. In *International Joint Conference on Artificial Intelligence*, 2022.

[23] Qing Guo, Wei Feng, Ce Zhou, Rui Huang, Liang Wan, and Song Wang. Learning dynamic siamese network for visual object tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2017.

[24] Zichao Guo, Xiangyu Zhang, Haoyuan Mu, Wen Heng, Zechun Liu, Yichen Wei, and Jian Sun. Single path one-shot neural architecture search with uniform sampling. In *European Conference on Computer Vision*, 2020.

[25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016.

[26] Lianghua Huang, Xin Zhao, and Kaiqi Huang. Got-10k: A large high-diversity benchmark for generic object tracking in the wild. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.

[27] Yu Huang, Chenzhuang Du, Zihui Xue, Xuanyao Chen, Hang Zhao, and Longbo Huang. What makes multi-modal learning better than single (provably). *Advances in Neural Information Processing Systems*, 2021.

[28] Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. Bilinear attention networks. In *Advances in Neural Information Processing Systems*, 2018.

[29] Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, and Junjie Yan. Siamrpn++: Evolution of siamese visual tracking with very deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.

[30] Bo Li, Junjie Yan, Wei Wu, Zheng Zhu, and Xiaolin Hu. High performance visual tracking with siamese region proposal network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.

[31] Zhenyang Li, Ran Tao, Efstratios Gavves, Cees GM Snoek, and Arnold WM Smeulders. Tracking by natural language specification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017.

[32] Zhenyang Li, Ran Tao, Efstratios Gavves, Cees GM Snoek, and Arnold WM Smeulders. Tracking by natural language specification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017.

[33] Liting Lin, Heng Fan, Yong Xu, and Haibin Ling. Swintrack: A simple and strong baseline for transformer tracking. *arXiv preprint arXiv:2112.00995*, 2021.

[34] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common objects in context. In *European Conference on Computer Vision*, 2014.

[35] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. In *International Conference on Learning Representations*, 2019.

[36] Tim Meinhardt, Alexander Kirillov, Laura Leal-Taixe, and Christoph Feichtenhofer. Trackformer: Multi-object tracking with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.

[37] Juan-Manuel Pérez-Rúa, Valentin Vielzeuf, Stéphane Pateux, Moez Baccouche, and Frédéric Jurie. Mfas: Multimodal fusion architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.

[38] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 2021.

[39] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *The Association for the Advancement of Artificial Intelligence*, 2019.

[40] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017.

[41] Zhuoran Shen, Mingyuan Zhang, Haiyu Zhao, Shuai Yi, and Hongsheng Li. Efficient attention: Attention with linear complexities. In *IEEE Winter Conference on Applications of Computer Vision*, 2021.

[42] Ajit Singh. Foundations of machine learning. *Available at SSRN 3399990*, 2019.

[43] Peize Sun, Yi Jiang, Rufeng Zhang, Enze Xie, Jinkun Cao, Xinting Hu, Tao Kong, Zehuan Yuan, Changhu Wang, and Ping Luo. Transtrack: Multiple-object tracking with transformer. *arXiv:2012.15460*, 2020.

[44] Ran Tao, Efstratios Gavves, and Arnold WM Smeulders. Siamese instance search for tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016.

[45] Nilesh Tripuraneni, Michael Jordan, and Chi Jin. On the theory of transfer learning: The importance of task diversity. *Advances in Neural Information Processing Systems*, 2020.

[46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017.

[47] Paul Voigtlaender, Jonathon Luiten, Philip HS Torr, and Bastian Leibe. Siam r-cnn: Visual tracking by re-detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.

[48] Naiyan Wang, Jianping Shi, Dit-Yan Yeung, and Jiaya Jia. Understanding and diagnosing visual tracking systems. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2015.

[49] Ning Wang, Wengang Zhou, Jie Wang, and Houqiang Li. Transformer meets tracker: Exploiting temporal context for robust visual tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.

[50] Xiao Wang, Xiujun Shu, Zhipeng Zhang, Bo Jiang, Yaowei Wang, Yonghong Tian, and Feng Wu. Towards more flexible and accurate object tracking with natural language: Algorithms and benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.

[51] Bin Yan, Houwen Peng, Jianlong Fu, Dong Wang, and Huchuan Lu. Learning spatio-temporal transformer for visual tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.

[52] Bin Yan, Houwen Peng, Kan Wu, Dong Wang, Jianlong Fu, and Huchuan Lu. Lighttrack: Finding lightweight neural networks for object tracking via one-shot architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.

[53] Bin Yang, Junjie Yan, Zhen Lei, and Stan Z Li. Convolutional channel features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2015.

[54] Yuechen Yu, Yilei Xiong, Weilin Huang, and Matthew R Scott. Deformable siamese attention networks for visual object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.

[55] Zhou Yu, Yuhao Cui, Jun Yu, Meng Wang, Dacheng Tao, and Qi Tian. Deep multimodal neural architecture search. In *ACM International Conference on Multimedia*, 2020.

[56] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.

[57] Zhipeng Zhang, Yihao Liu, Xiao Wang, Bing Li, and Weiming Hu. Learn to match: Automatic matching network design for visual tracking. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.

[58] Zhipeng Zhang and Houwen Peng. Deeper and wider siamese networks for real-time visual tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.

[59] Zhipeng Zhang, Houwen Peng, Jianlong Fu, Bing Li, and Weiming Hu. Ocean: Object-aware anchor-free tracking. In *European Conference on Computer Vision*, 2020.

[60] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. In *International Conference on Learning Representations*, 2017.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] See Sec. 1.
   (b) Did you describe the limitations of your work? [Yes] See Sec. 1.
   (c) Did you discuss any potential negative societal impacts of your work? [No]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Sec. 3.4.
   (b) Did you include complete proofs of all theoretical results? [Yes] See Sec. 3.4.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Sec. 4.1.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Sec. 4.1.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]