


Efficient Visual Tracking with Exemplar Transformers

Conference Paper**Author(s):**

Blatter, Philippe; Kanakis, Menelaos; [Danelljan, Martin](#) ; Van Gool, Luc

Publication date:

2023

Permanent link:

<https://doi.org/10.3929/ethz-b-000592890>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

<https://doi.org/10.1109/WACV56688.2023.00162>

Efficient Visual Tracking with Exemplar Transformers

Philippe Blatter^{*,1} Menelaos Kanakis^{*,1} Martin Danelljan¹ Luc Van Gool^{1,2}

¹ETH Zürich ²KU Leuven

Abstract

The design of more complex and powerful neural network models has significantly advanced the state-of-the-art in visual object tracking. These advances can be attributed to deeper networks, or the introduction of new building blocks, such as transformers. However, in the pursuit of increased tracking performance, runtime is often hindered. Furthermore, efficient tracking architectures have received surprisingly little attention. In this paper, we introduce the Exemplar Transformer, a transformer module utilizing a single instance level attention layer for realtime visual object tracking. E.T.Track, our visual tracker that incorporates Exemplar Transformer modules, runs at 47 FPS on a CPU. This is up to 8× faster than other transformer-based models. When compared to lightweight trackers that can operate in realtime on standard CPUs, E.T.Track consistently outperforms all other methods on the LaSOT [16], OTB-100 [52], NFS [27], TrackingNet [36], and VOT-ST2020 [29] datasets. Code and models are available at <https://github.com/pblatter/ettrack>.

1. Introduction

Estimating the trajectory of an object in a video sequence, referred to as visual tracking, is one of the fundamental problems in computer vision. Deep neural networks have significantly advanced the performance of visual tracking methods with deeper networks [4], more accurate bounding boxes [31], or with the introduction of new modules, such as transformers [53, 8, 48]. However, these advances often come at the cost of more expensive models. While the demand for realtime visual tracking on applications such as autonomous driving, robotics, and human-computer-interfaces is increasing, efficient deep tracking architectures have received surprisingly little attention. This calls for visual trackers that, while accurate and robust, are capable of operating in realtime under the hard computational constraints of limited hardware.

*P. Blatter and M. Kanakis contributed equally to this work.

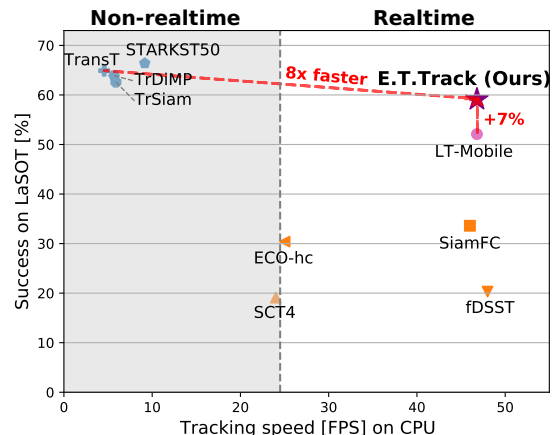


Figure 1: Comparison of tracker performance in terms of AUC score (Success in %) on LaSOT vs. tracking speed in FPS on a standard CPU. Our Exemplar Transformer Tracker (E.T.Track) outperforms all other realtime trackers. It achieves a 7% higher AUC score than LT-Mobile [54]. Furthermore, our approach achieves up to 8× faster runtime on a CPU compared to previous Transformer-based trackers.

Transformers [45], proposed for machine translation, have also demonstrated superior performance in a number of vision based tasks, including image [3] and video [51] classification, object detection [7], and even multi-task learning [6]. The field of visual tracking has also observed similar performance benefits [53, 42, 48]. While transformers have enabled the trackers to improve accuracy and robustness, they severely suffer from high computational cost, leading to decreased runtime operation, as depicted in Fig. 1. In this work, we set out to find a transformer module, capable increasing tracking accuracy and robustness while not compromising runtime.

In this work we propose Exemplar Attention, a single instance-level attention layer for visual tracking. Our attention module exploits domain specific knowledge to improve the tracker’s performance, while maintaining a comparable runtime. Specifically, we build upon two hypotheses. Firstly, one global query value is sufficiently descrip-

tive when tracking a single object. Secondly, a small set of exemplar values can act as a shared memory between the samples of the dataset. Thus, our constrained instance-level Exemplar Attention captures more explicit information about the target object, compared to conventional attention modules.

We develop the Exemplar Transformer Tracker (E.T.Track) by integrating our Exemplar Transformer layer into a Siamese tracking architecture. Specifically, we replace the convolutional layers in the tracker heads with the Exemplar Transformer layer. The additional expressivity from the Exemplar Transformer layer significantly improves the performance of the models based on regular convolutional layers. The added performance gain comes at an insignificant cost in runtime, as seen in Fig. 1 when comparing it to the mobile LightTrack [54] (LT-Mobile). We further compare our transformer layer for single object tracking to other generic transformer layers. We find that Exemplar Transformer consistently outperforms competing methods, attesting to the benefits of explicitly designing attention layers for the task of visual tracking.

We validate our approach on six benchmark datasets: LaSOT [16], OTB-100 [52], UAV-123 [35], NFS [27], TrackingNet [36] and VOT-ST2020 [29]. Our proposed tracker runs at 46.8 Frames Per Second (FPS) on a CPU, while setting a new state-of-the-art among realtime CPU trackers by achieving 59.1% AUC on the challenging LaSOT dataset.

In summary, our contributions are:

- We introduce Exemplar Transformer, a transformer layer based on a single instance-level attention layer referred to as Exemplar Attention.
- We develop a transformer-based tracking architecture based upon our Exemplar Transformer layer.
- Our tracker runs in realtime on a CPU, while outperforming previous realtime trackers on 5 benchmarks.

2. Related Work

Siamese Trackers In recent years, Siamese trackers have gained significant popularity due to their performance capabilities and simplicity. The Siamese-based tracking framework formulates visual object tracking as a template matching problem, utilizing cross-correlation between a search and an image patch. The original work of Bertinetto *et al.* introduced SiamFC [4], the first model incorporating feature correlation into a Siamese framework. Li *et al.* [31] introduced region proposal networks to increase efficiency and obtain more accurate bounding boxes. More recent advances on the Siamese tracker front include the use of additional branches [49], refinement modules for more precise bounding box regression [55], and various model update mechanisms [17, 18, 56, 58]. Unlike previous Siamese

trackers, we propose the Exemplar Transformer module that is incorporated into the prediction heads, and improves the tracker’s performance at an insignificant runtime increase.

Transformers in Tracking The Transformer [45] was introduced as a module to improve the learning of long-range dependencies in neural machine translation, by enabling every element to attend to all others. In computer vision, transformers have been used in image [3] and video [51] classification, object detection [7], and even multi-task learning of dense prediction tasks [6]. More related to our work, transformers have also been utilized to advance the performance of visual trackers. STARK [53] utilizes transformers to model the global spatio-temporal feature dependencies between target object and search regions. This is achieved by integrating a dynamically updated template into the encoder, in addition to the regular search and template patch. [48] introduced a transformer architecture that improves the standard Siamese-like pipeline by additionally exploiting temporal context. The encoder model mutually reinforces multiple template features by leveraging self-attention blocks. In the decoder, the template and search branch are bridged by cross-attention blocks in order to propagate temporal contexts. [8] also improve Siamese-based trackers by replacing the regular correlation operation by a Transformer-based feature fusion network. The Transformer-based fusion model aggregates global information, providing a superior alternative to the standard linear correlation operation. ToMP [34], on the other hand, utilizes a transformer to predict the weights of a convolutional kernel in order to localize the target in the search region and template patches. In this work, we also design transformer architecture for tracking. Unlike the previous transformers for tracking, Exemplar Transformer is lightweight and can be utilized in computationally limited hardware running at realtime.

Efficient Tracking Architectures With an increase in demand for realtime visual tracking in applications such as autonomous driving, and human-computer-interfaces, efficient deep tracking architectures are essential. Surprisingly, however, little attention has been provided on efficient trackers that can operate on computationally limited hardware. KCF [21] and fDSST [13] employ hand-crafted features to enable realtime operation on CPUs. While fast, their reliance on hand crafted features significantly hinders their performance compared to newer and more complex methods. In contrast, we present an efficient deep tracker that operates at a comparable runtime but performs on par with the more expensive deep trackers. More related to our work, LightTrack [54] employs neural architecture search (NAS) to find a lightweight and efficient Siamese tracking architecture. We instead propose an efficient transformer layer that can complement existing architecture advances




Figure 2: Comparison of our Exemplar Attention module (right) and the standard scaled dot-product attention module [45]. The matching blocks are indicated by identical colours. The line thickness is indicative of the tensor size.

such as LightTrack. Specifically, our transformer layer can act as a drop in replacement for convolutional layers, increasing performance with negligible effects on runtime.

Efficient Transformers The immense interest in transformer architectures [15, 37, 7, 51] resulted in the development of various efficient model variants that can be grouped in 4 major categories [44]. *Low Rank/ Kernel* methods assume and leverage low-rank approximations of the self-attention matrix [26, 9]. *Memory/ Downsampling* methods learn a side memory module to access multiple tokens simultaneously, or simply reduce the sequence length [43, 59, 33]. *Fixed/ Factorized/ Random Patterns* limit the field of view of the self-attention, such as using block patterns [40, 59, 33, 38]. *Learnable Patterns* replace the fixed pattern, as in standard transformers, with dynamic patterns [47, 50, 28]. Our work falls in the intersection of *Memory/ Downsampling* and *Fixed/ Factorized/ Random Patterns*. Unlike the aforementioned works that aim to design generic attention layers, Exemplar Attention is instead designed for the task of single object visual tracking by exploiting domain specific knowledge.

3. Efficient Tracking with Transformers

Striking a balance between well performing object trackers and runtime speeds that fall in the realtime envelope, is a challenging problem when deploying on computationally limited devices. In this section, we introduce the Exemplar Transformer, a transformer architecture based on single instance level attention layers for single object tracking. While lightweight, our Exemplar Transformer significantly closes the performance gap with the computationally expensive transformer-based trackers [53, 48, 8]. Sec. 3.1 first

presents the original Transformer of Vaswani *et al.* [45], followed by our Exemplar Transformer formulation. Sec. 3.2 introduces our E.T.Track. Specifically, it first outlines the overall architecture, and presents how Exemplar Transformers are utilized within the tracker.

3.1. Exemplar Transformers

Standard Transformer The Transformer [45], introduced for machine translation, receives a one dimensional input sequence $x \in \mathbb{R}^{N \times D}$ with N feature vectors of dimensions D . The input sequence is processed by a series of transformer layers defined as

$$T(x) = f(A(x) + x). \quad (1)$$

The function $f(\cdot)$ is a lightweight Feed-Forward Network (FFN) that projects independently each feature vector. The function $A(\cdot)$ represents a self-attention layer that acts across the entire sequence. Specifically, the authors used the ‘‘Scaled Dot-Product Attention’’, defined as

$$\begin{aligned} A(x) &= \text{softmax} \left(\underbrace{\frac{\overbrace{f_Q(x)}^Q \overbrace{f_K(x)}^{K^T}}{\sqrt{d_k}}}_{\text{constant}} \right) \overbrace{f_V(x)}^V \\ &= \text{softmax} \left(\underbrace{\frac{\overbrace{(xW_Q)}^{f_Q(x)} \overbrace{(W_K^T x^T)}^{f_K(x)}}{\sqrt{d_k}}}_{\text{constant}} \right) \overbrace{(xW_V)}^{f_V(x)}. \end{aligned} \quad (2)$$

The queries $Q \in \mathbb{R}^{N \times D_{QK}}$, keys $K \in \mathbb{R}^{N \times D_{QK}}$, and values $V \in \mathbb{R}^{N \times D_V}$ represent projections of the input sequence, while $\sqrt{d_k}$ is a normalization constant. The self attention, therefore, computes a similarity score between all representations, linearly combines the feature representations, and accordingly adapts the input representation x in Eq. 1. The computational complexity of Eq. 2 is $\mathcal{O}(N^2 D)$, *i.e.* it scales quadratically with the length of the input sequence.

Exemplar Attention We now introduce Exemplar Attention, the key building block of the Exemplar Transformer module. We hypothesize that, while the direct connection between all features is essential in machine translation and some vision tasks, this design choice can be sub-optimal when attending to the single object being tracked. We describe the required modifications of the individual components below.

The standard Query function f_Q projects every spatial location of the feature map independently to a query space. Unlike machine translation where every feature represents a specific word or token, adjacent spatial representation in

vision tasks often correspond to the same object. Consequently, we aggregate the information of the feature map $X \in \mathbb{R}^{H \times W \times D}$, where $H \times W$ represents the spatial dimensions. Specifically, we use a 2D adaptive average pooling layer with an output spatial dimension $S \times S$, followed by a flattening operation. The operation is denoted $\Psi_S(X)$, decreasing the output spatial dimension to S^2 . The compressed representation of X is then projected to a query space as in the standard self-attention formulation.

$$Q = \Psi_S(X)W_Q \in \mathbb{R}^{S^2 \times D_{QK}} \quad (3)$$

We hypothesize that for single instance tracking, one global query value is sufficient to identify the object of interest, while also decreasing the computational complexity of the module. To this extent, we set $S = 1$. This design choice is further supported by the success of global pooling in classification architectures [20], as well as transformer based object detection [7].

The keys and values, as presented in Eq. 2, are per spatial location linear projections of the input. The self-attention layer then enables the learning of spatial correlations, at the cost of every feature attending to all others. This eliminates spatial biases built into convolutional layers. Rather than requiring a fine grained feature map and relying solely on intra-sample relationships, we instead learn a small set of exemplar representations. The exemplar representations encapsulates dataset information in order to dynamically adapt the attention layer given the global query token and the captured information. To this end, we optimize a small set of exemplar keys $K = \hat{W}_K \in \mathbb{R}^{E \times D_{QK}}$ that, unlike the formulation in Eq. 2, are independent of the input. The similarity matrix therefore associates the global query, Eq. 3, to exemplars. Our attention layer then refines the input representation on the local level by replacing the projection $f_V(\cdot)$ with a convolutional operation

$$V = W_V \otimes X \in \mathbb{R}^{E \times H \times W \times D_V}, \quad (4)$$

where $W_V \in \mathbb{R}^{E \times Z \times Z}$ can be of any spatial dimension Z , while the number of exemplars E can be chosen arbitrarily. We use $E = 4$ in our experiments, which is significantly smaller than the dimensions $H \times W$, maintaining comparable runtime.

Our efficient Exemplar Attention is therefore defined as,

$$A(x) = \text{softmax} \left(\underbrace{\frac{\overbrace{f_Q(x)}{(\Psi_S(X)W_Q)} \overbrace{f_K(\cdot)}{(\hat{W}_K^T)}}{\underbrace{\sqrt{d_k}}_{\text{constant}}}}_{\text{constant}} \right) \overbrace{f_V(x)}{(W_V \otimes X)}, \quad (5)$$

but can also be written as,

$$A(x) = \left[\text{softmax} \left(\frac{(\Psi_S(X)W_Q)(\hat{W}_K^T)}{\sqrt{d_k}} \right) W_V \right] \otimes X. \quad (6)$$




Figure 3: E.T.Track - a Siamese tracking pipeline that incorporates Exemplar Transformers in the tracker head.

Exemplar Attention, while inspired by the scaled dot-product attention, is conceptually very different. In self-attention (2), $f_{\{Q,K,V\}}$ act as projections to their corresponding feature spaces, with the similarity function learning the relationships between all spatial locations. In other words, self-attention relies solely on intra-sample relationships, and therefore requires fine-grained representations. Instead, the Exemplar Attention layer enforces the attending over a single instance through the use of a global query token. The global query encapsulated the representation of the object, is dynamically generated from the input image, and applied locally on the feature map using a convolutional operation. To enable the use of a single query token, we exploit dataset information to form the exemplar representations through end-to-end optimization, eliminating the need of the intra-sample similarity function. A comparison between the two attention mechanisms is depicted in Fig. 2.

3.2. E.T.Track Architecture

In this section we introduce the base tracking architecture used throughout our work. While Exemplar Transformers can be incorporated into any tracking architecture, we evaluate its efficacy on lightweight Siamese trackers. An overview of the E.T.Track architecture can be seen in Fig. 3.

Our model employs the lightweight backbone model LT-Mobile [54]. The model was identified by NAS on a search space consisting of efficient and lightweight building blocks. The feature extracting backbone consists of 3×3 convolutional layers, depthwise separable convolutional layers and mobile inverted bottleneck layers with squeeze and excitation modules.

The Exemplar Transformer layer can act as a drop in replacement for any convolution operation of the architecture. We replace all the convolutions in the classification and bounding box regression branches, while keeping the lightweight backbone architecture untouched. This eliminates the need for retraining the backbone on ImageNet [14].

Search and template frames are initially processed through a backbone network. The similarity between the representations is computed by a pointwise cross-correlation. The resulting correlation map is then fed into the tracker head, where it is processed by a classification branch and a bounding box regression branch in parallel. The bounding box regression branch predicts the distance to all four sides of the bounding box. The classification branch predicts whether each region is part of the foreground or the background. During training, the bounding box regression branch considers all the pixels within the ground truth bounding box as training samples, therefore, the model is able to determine the exact location of the object even when only small parts of the input image are classified as foreground. The model is trained by optimizing a weighted combination of the binary cross-entropy (BCE) loss and the IoU loss [57] between the predicted and ground-truth bounding boxes. For more details, as well as more information on the data preprocessing, we refer the reader to [61].

4. Experiments

We first present implementation details of our tracker in section 4.1. The comparison to state-of-the-art is presented in Sec. 4.2, followed by an ablation study in Sec. 4.3. Code and trained models will be released on publication.

4.1. Implementation Details

Architecture We adopt the LT-Mobile architecture of LightTrack [54] as our baseline due to its performance versus efficiency trade-off. LT-Mobile is comprised of a small

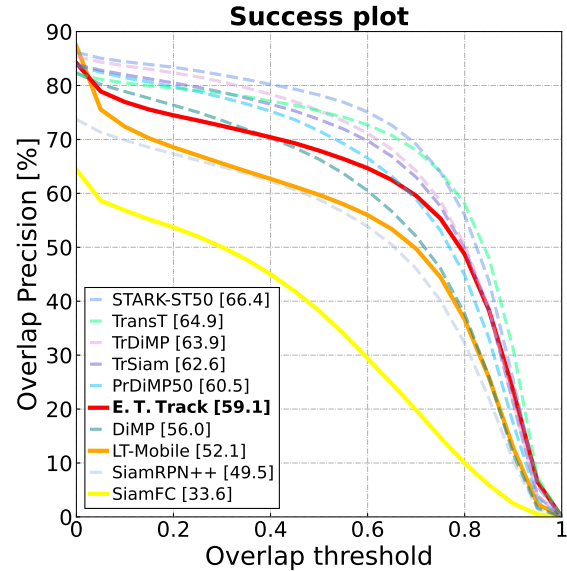


Figure 4: Success plot on the LaSOT dataset. The CPU real-time trackers are indicated by continuous lines in warmer colours, while the non-realtime trackers are indicated by dashed lines in colder colours. E.T.Track significantly outperforms the other real-time trackers, and even outperforms some of the more established trackers such as DiMP [5]. Furthermore, it significantly closes the performance gap with the more expensive transformer trackers.

encoder, and later branches to the classification and regression heads. The classification head is comprised of 6 convolutional modules, while the regression head is comprised of 8. Each convolutional module consists of a Depthwise Separable Convolution [23], a Batch Normalization layer [25] and a Rectified Linear unit. E.T.Track replaces each convolutional module with an Exemplar Transformer Layer, introduced in Sec. 3.1. The learnable “value” parameters of the attention module are initialized using kaiming initialization [19], while the learnable “keys” parameters are initialized using a normal distribution. The FFN consists of 2 linear layers with a ReLU activation, dropout [41] with a ratio of 0.1, and LayerNorm [2].

Training All models have been trained using an Nvidia GTX TITAN X, and evaluated on an Intel(R) Core(TM) i7-8700 CPU @ 3.20GHz. The training of our E.T.Track architecture is based on the training framework used in LightTrack [54] which, in turn, is based on OCEAN [61]. As is common practice, we initialize the backbone with ImageNet pre-trained weights. The models are optimized using stochastic gradient descent [39] with a momentum of 0.9, and a weight decay of $1e-4$ for 50 epochs. During the first 10 epochs, the backbone parameters remain frozen. We use a step learning rate scheduler during a warmup period of 5 epochs, increasing the learning rate from $2e-2$

	non-realtime									realtime		
	ATOM [10]	SiamRPN++ [30]	DiMP-50 [5]	PrDiMP-50 [12]	SiamR-CNN [46]	TransT [8]	TrDiMP [48]	TrSiam [48]	STARK-ST50 [53]	ECO [11]	LT-Mobile [54]	E.T.Track (Ours)
NFS	58.4	50.2	62	63.5	63.9	65.7	66.5	65.8	66.4	46.6	55.3	59.0
UAV-123	64.2	61.3	65.3	68	64.9	69.4	67.5	67.4	68.8	51.3	62.5	62.3
OTB-100	66.9	69.6	68.4	69.6	70.1	69.1	71.1	70.8	67.3	64.3	66.2	67.8
CPU Speed	20	15	15	15	15	5	6	6	9	25	47	47

Table 1: State-of-the-art comparison on the NFS, OTB-100 and UAV-123 datasets in terms of Area Under the Curve (AUC). The best score is highlighted in blue while the best realtime score is highlighted in red. We additionally report CPU runtime speeds in FPS.

	non-realtime									realtime		
	ATOM [10]	SiamRPN++ [30]	DiMP-50 [5]	PrDiMP-50 [12]	SiamR-CNN [46]	TransT [8]	TrDiMP [48]	TrSiam [48]	STARK-ST50 [53]	ECO [11]	LT-Mobile [54]	E.T.Track (Ours)
Prec. (%)	64.84	69.38	68.7	70.4	80	80.3	73.1	72.7	-	48.86	69.5	70.6
N. Prec. (%)	77.11	79.98	80.1	81.6	85.4	86.7	83.3	82.9	86.1	62.14	77.9	80.3
Success (%)	70.34	73.3	74	75.8	81.2	81.4	78.4	78.1	81.3	56.13	72.5	75.0
CPU Speed	20	15	15	15	15	5	6	6	9	25	47	47

Table 2: State-of-the-art comparison on the TrackingNet test set, comprised of 511 sequences. The trackers are compared in terms of Precision (Prec.), Normalized Precision (N. Prec.), and Success. The best score is highlighted in blue while the best realtime score is highlighted in red. We additionally report CPU runtime speeds in FPS.

to $1e-1$, followed by a logarithmically decreasing learning rate from $1e-1$ to $2e-4$ for the remainder. We utilize 3 GPUs and sample 32 image pairs per GPU for each batch. The sampled image pairs consist of a 256×256 search frame and a 128×128 template frame, sampled from training splits of LaSOT [16], TrackingNet [36], GOT10k [24] and COCO [32]. Specifically, the two frames are sampled within a range of 100 frames for LaSOT [16] and GOT10k [24], 30 frames for TrackingNet [36], and 1 frame for COCO [32]. Both patches are further shifted and scaled randomly.

4.2. Comparison to State-of-the-Art

We compare our proposed E.T.Track to state-of-the-art methods on 6 benchmarks: OTB-100 [52], NFS [27], UAV-123 [35], LaSOT [16], TrackingNet [36], and VOT2020 [29]. Specifically, we evaluate transformer-based trackers [8, 48, 48, 53], realtime CPU trackers [11, 54], as well as additional seminal trackers [10, 30, 5, 12, 46]. For all methods we further report the CPU runtime in FPS.

LaSOT [16] The LaSOT dataset is highly challenging, and includes very long sequences with an average of 2500 frames per sequence. Therefore, robustness is essential to achieving a high score. The success plot in Fig. 4 depicts the CPU realtime trackers with warmer colour continuous lines, while the non-realtime trackers are indicated by dashed lines in colder colours. Unlike online-learning methods such STARK that utilize a dynamically updated template, our model only uses the features of the template patch extracted in the first sequence of the frames. Even so, our model is very robust and reaches an AUC score

of 59.1%, outperforming the popular DiMP tracker [5] by 2.2%. Compared to the lightweight mobile architecture of LT-Mobile [54], our model improves the success score by an astonishing 7% while achieving a comparable speed.

NFS [27] We additionally evaluate our approach on the NFS dataset that contains fast moving objects. The results are presented in Table 1. E.T.Track reaches an AUC score of 59%, outperforms all realtime trackers by at least 3.7%.

OTB-100 [52] OTB-100 contains 100 sequences. As shown in Table 1, the current state-of-the-art is achieved by the recently introduced TrDiMP [48] with an AUC score of 71.1%. Our model achieves an AUC score of 67.8%, marking it as the best performing realtime tracker.

UAV-123 [35] UAV-123 contains a total of 123 sequences from aerial viewpoints. The AUC results are shown in Table 1. Unlike the other datasets, E.T.Track performs comparably to LT-Mobile, with a performance of 62.3%.

TrackingNet [36] We further evaluate the trackers on the 511 sequences of the TrackingNet test set, and report the results in Table 2. Similarly to the other datasets, E.T.Track outperforms all other realtime trackers. Specifically, E.T.Track improves LT-Mobile precision by 1.05%, normalized precision by 2.42%, and AUC by 2.48%. Comparing E.T.Track to more complex transformer-based trackers such as TrSiam [48], our model is only 2.2% worse in terms of precision, 2.32% in terms of normalized precision, and 3.12% in terms of AUC while running almost $8\times$ faster on a CPU. This further demonstrates that, while transformers have the capability to significantly improve performance, transformer modules do not need to be pro-

	non-realtime					realtime		
	SiamFC [4]	ATOM [10]	DiMP [5]	SuperDiMP [1]	STARK-ST50 [53]	KCF [21]	LT-Mobile [54]	E.T.Track (Ours)
EAO	0.179	0.271	0.274	0.305	0.308	0.154	0.242	0.267
Accuracy	0.418	0.462	0.457	0.477	0.478	0.407	0.422	0.432
Robustness	0.502	0.734	0.740	0.786	0.799	0.432	0.689	0.741
CPU Speed	6	20	15	15	9	95	47	47

Table 3: Comparison of bounding box predicting trackers on the VOT-ST2020 dataset. We report the Expected Average Overlap (EAO), the Accuracy and Robustness. The best score is highlighted in blue while the best realtime score is highlighted in red. We additionally report CPU runtime speeds in FPS.

hibitively expensive for computationally constrained devices to achieve most of the performance gains.

VOT-ST2020 [29] Finally, we also evaluate bounding box predicting trackers on the anchor-based short term tracking dataset of VOT-ST2020. Unlike other tracking datasets, VOT2020 contains various anchors that are placed Δ_{anc} frames apart. The trackers are evaluated in terms of Accuracy, Robustness and Expected Average Overlap (EAO). Accuracy represents a weighted combination of the average overlap between the ground truth and the predicted target predictions on subsequences defined by anchors. Robustness indicates the percentage of frames before the trackers fails on average. Finally, EAO is a measure for the overall tracking performance and combines the accuracy and the robustness. The results are shown in Table 3. While our model outperforms the lightweight convolutional baseline model introduced in [54] by 1 – 2% in terms of accuracy and robustness, the largest performance increase can be noted in terms of robustness, where the performance is increased by 5.2%. We find that learning exemplar representations from the dataset coupled with an image-level query representation significantly increases the tracker’s robustness compared to its convolutional counterpart.

4.3. Ablation Study

To further understand the contributions of the different components, we conduct a number of controlled experiments on three datasets. Specifically, we report AUC on OTB-100 [52], NFS [27], and LaSOT [16].

Baseline We commence our ablation study from the mobile architecture of [54], LT-Mobile, due to its performance versus efficiency trade-off. We refer to LT-Mobile, our baseline Siamese tracker, as the Convolutional(Conv) baseline. In LT-Mobile, the similarity of the search and template patch features is computed by a pointwise cross-correlation. The feature map is then passed to the tracker head consisting of two branches, the classification and the bounding box regression branches, as explained in Sec. 4.1. The performance of the baseline model, Conv, is reported in Table 5a.

Exemplar Attention We first evaluate the efficacy of the Exemplar Attention as a drop in replacement for the convolutional layer. We replace the convolutional layers with the Exemplar Attention, followed by a residual connection and a normalization layer. In other words, setting the FFN ($f(\cdot)$) in Eq. 1 to identify. We report the performance of the Attention (Att) module in Table 5a. The performance on NFS increases by 1.3%, and on LaSOT by 1.5%, demonstrating the effectiveness of our Exemplar Attention module. We note that this performance increase is without the use of the FFN, a key design choice in the transformer architectures [45].

FFN Similar to the original Transformer architecture [45], we evaluate the effect of additionally using a lightweight FFN followed by a LayerNorm layer. We find that the additional expressivity introduced by the FFN improves the performance on all three datasets, as seen in Table 5a. The highest performance increase is achieved on LaSOT, where the AUC score increases by 5.5%. This yields our final E.T.Track model, depicted in Fig. 3.

Template conditioning The queries used in the Exemplar Transformer so far are solely based on a transformed version of the initial correlation map. We further explore the impact of incorporating template information into our Exemplar Attention module. Specifically, we average pool the feature map corresponding to the template patch, and sum the representation to each layer’s input. As seen from the Template Conditioning (T-Cond) experiments in Table 5a, the richer queries lead to an improvement on NFS. However, on OTB-100 and LaSOT, the model did not benefit from the additional information. To this extend, we decide to not use the T-Cond module in our final module, keeping our final model simpler.

Number of Exemplars Table 5b reports the performance given the number of Exemplars. While more Exemplars increases the overall capacity of the model, and as such, one would expect further performance gains, our experiments yield different results. Specifically, 4 Exemplars yield consistently better results across all the datasets. We hypothesize that, while training a model with a larger number of

	Conv [54]	Standard [45]	Clustered [47]	Linear [26]	Local [38]	Swin [33]	E.T.Track (Ours)
NFS	55.3	55.3	57.5	55.8	55.8	55.4	59.0
OTB-100	66.2	65.3	67.5	65.4	64.8	64.2	67.8
LaSOT	52.1	54.2	56.5	53.5	53.4	56.9	59.1

Table 4: Comparison of the convolutional baseline (Conv) and the different attention modules in terms of AUC on NFS, OTB, and LaSOT datasets. The best score is highlighted in blue. E.T.Track consistently outperforms all other transformer variants.

Conv	Att	FFN	T-Cond.	NFS	OTB-100	LaSOT
✓				55.3	66.2	52.1
	✓			56.6	65.8	53.6
	✓	✓		58	67.3	59.1
	✓	✓	✓	59.0	66.9	57.9

(a) Ablating the different component of the Exemplar Transformer module. We evaluate the Exemplar Attention (Att) module, Feed-Forward Network (FFN), and Template Conditioning (T-Cond). The final model, depicted in Fig. 3 includes the Att and FFN modules.

	Conv	1-Ex	4-Ex	16-Ex
NFS	55.3	57.6	58.0	58.0
OTB-100	66.2	66.5	67.3	66.1
LaSOT	52.1	57.2	59.1	57.4

(b) Effect of the number of exemplars (-Ex).

	ShuffleNet [60]	MobileNetV3 [22]	ResNet-18 [20]	LT-Mobile [54]
Conv	✓	✓	✓	✓
E.T. (Ours)	✓	✓	✓	✓
NFS	54.9	56.2	56.8	56.8
OTB-100	61.3	61.8	64.5	65.3
LaSOT	48.6	49.8	52.1	52.7
			55.8	57.3
			65.3	65.7
			55.9	56.5
			52.1	59.1

(c) Comparison of Exemplar Transformer (E.T.) and Convolutional (Conv) modules on different backbone models. The E.T. consistently outperforms the Conv models, independently of the backbone used.

Table 5: Ablation experiments reported in terms of AUC on NFS, OTB, and LaSOT datasets. Conv refers to LT-Mobile [54] that acts as our convolutional baseline. The best score is highlighted in blue.

experts can increase performance, modifications in the optimization process are required to ensure the selection of the appropriate exemplar. The efficient implementation of the Exemplar Attention module, Eq. 6, ensures a comparable runtime even with a larger number of exemplars.

Interestingly, while single Exemplar Attention is mathematical equivalent to a regular convolution with a residual operation, the additional FFN layer following the Exemplar Attention increases considerably the performance. Specifically, we observe a performance increase of 2.3% on NFS, 0.3% on OTB-100, and 5.1% on LaSOT.

Backbone Alternatives All experiments reported so far utilize the LT-Mobile encoder. To demonstrate the flexibility of Exemplar Transformers, as well as their indepen-

dence to the encoder architecture, we evaluate the use of different encoder architectures. Specifically, we compare the performance of the two tracker head module variants (Convolution, Exemplar Transformer) in combination with ShuffleNet [60], MobileNetV3 [22], ResNet-18 [20], and LT-Mobile [54]. The results presented in Table 5c demonstrate consistent performance gains independent of the encoder architecture, highlighting the superiority of our Exemplar Transformer to its convolutional counterpart.

Comparison of Alternative Transformer Layers To validate the design choices and hypothesis that lead to the Exemplar Transformer module, we additionally compare to other Transformer Layer variants. All Transformer layers evaluated can also act as drop-in replacements to standard convolutions. Specifically, we evaluate the Standard [45], Clustered [47], Linear [26], Local [38], and Swin [33] Transformers. The selection ensures at least one method from every transformer category defined in Sec. 2, while using their official public implementations ensures a fair comparison. The results in Table 4 demonstrate that our Exemplar Transformer (E.T.) consistently outperforms all other attention variants across all the datasets. These findings further validate our hypothesis that one global query and a small set of exemplar representations is sufficiently descriptive when tracking a single object.

5. Conclusion

We propose a novel Transformer layer for single object visual tracking, based on Exemplar Attention. Exemplar Attention utilizes a single query token of the input sequence, and jointly learns a small set of exemplar representations. The proposed transformer layer can be used throughout the architecture, *e.g.* as a substitute for a convolutional layer. Having a comparable computational complexity to standard convolutional layers while being more expressive, the proposed Exemplar Transformer layers can significantly improve the accuracy and robustness of tracking models with minimal impact on the model’s overall runtime. E.T.Track, our Siamese tracker with Exemplar Transformer, significantly improve the performance compared to the convolutional baseline and other transformer variants. E.T.Track is capable of running in realtime on computationally limited devices such as standard CPUs.

References

- [1] Pytracking. <https://github.com/visionml/pytracking>. Accessed: 2021-11-16. 7
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 5
- [3] Irwan Bello, Barret Zoph, Ashish Vaswani, Jonathon Shlens, and Quoc V Le. Attention augmented convolutional networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3286–3295, 2019. 1, 2
- [4] Luca Bertinetto, Jack Valmadre, Joao F Henriques, Andrea Vedaldi, and Philip HS Torr. Fully-convolutional siamese networks for object tracking. In *European conference on computer vision*, pages 850–865. Springer, 2016. 1, 2, 7
- [5] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning discriminative model prediction for tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6182–6191, 2019. 5, 6, 7
- [6] David Bruggemann, Menelaos Kanakis, Anton Obukhov, Stamatios Georgoulis, and Luc Van Gool. Exploring relational context for multi-task dense prediction. *arXiv preprint arXiv:2104.13874*, 2021. 1, 2
- [7] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229. Springer, 2020. 1, 2, 3, 4
- [8] 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 (CVPR)*, pages 8126–8135, June 2021. 1, 2, 3, 6
- [9] Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, David Belanger, Lucy Colwell, et al. Masked language modeling for proteins via linearly scalable long-context transformers. *arXiv preprint arXiv:2006.03555*, 2020. 3
- [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*, pages 4660–4669, 2019. 6, 7
- [11] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. Eco: Efficient convolution operators for tracking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6638–6646, 2017. 6
- [12] 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*, pages 7183–7192, 2020. 6
- [13] Martin Danelljan, Gustav Häger, Fahad Shahbaz Khan, and Michael Felsberg. Discriminative scale space tracking. *IEEE transactions on pattern analysis and machine intelligence*, 39(8):1561–1575, 2016. 2
- [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 5
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3
- [16] 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*, pages 5374–5383, 2019. 1, 2, 6, 7
- [17] Junyu Gao, Tianzhu Zhang, and Changsheng Xu. Graph convolutional tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4649–4659, 2019. 2
- [18] 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 international conference on computer vision*, pages 1763–1771, 2017. 2
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015. 5
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4, 8
- [21] João F Henriques, Rui Caseiro, Pedro Martins, and Jorge Batista. High-speed tracking with kernelized correlation filters. *IEEE transactions on pattern analysis and machine intelligence*, 37(3):583–596, 2014. 2, 7
- [22] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1314–1324, 2019. 8
- [23] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 5
- [24] 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. 6
- [25] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015. 5
- [26] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive

- transformers with linear attention. In *International Conference on Machine Learning*, pages 5156–5165. PMLR, 2020. [3](#), [8](#)
- [27] Hamed Kiani Galoogahi, Ashton Fagg, Chen Huang, Deva Ramanan, and Simon Lucey. Need for speed: A benchmark for higher frame rate object tracking. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1125–1134, 2017. [1](#), [2](#), [6](#), [7](#)
- [28] Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. *arXiv preprint arXiv:2001.04451*, 2020. [3](#)
- [29] Matej Kristan, Aleš Leonardis, Jiří Matas, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kämäräinen, Martin Danelljan, Luka Čehovin Zajc, Alan Lukežič, Ondrej Drobní, et al. The eighth visual object tracking vot2020 challenge results. In *European Conference on Computer Vision*, pages 547–601. Springer, 2020. [1](#), [2](#), [6](#), [7](#)
- [30] 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*, pages 4282–4291, 2019. [6](#)
- [31] 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 conference on computer vision and pattern recognition*, pages 8971–8980, 2018. [1](#), [2](#)
- [32] 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 *European conference on computer vision*, pages 740–755. Springer, 2014. [6](#)
- [33] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. [3](#), [8](#)
- [34] Christoph Mayer, Martin Danelljan, Goutam Bhat, Matthieu Paul, Danda Pani Paudel, Fisher Yu, and Luc Van Gool. Transforming model prediction for tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8731–8740, 2022. [2](#)
- [35] Matthias Mueller, Neil Smith, and Bernard Ghanem. A benchmark and simulator for uav tracking. In *European conference on computer vision*, pages 445–461. Springer, 2016. [2](#), [6](#)
- [36] Matthias Muller, Adel Bibi, Silvio Giancola, Salman Alsubaihi, and Bernard Ghanem. Trackingnet: A large-scale dataset and benchmark for object tracking in the wild. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 300–317, 2018. [1](#), [2](#), [6](#)
- [37] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Łukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In *International Conference on Machine Learning*, pages 4055–4064. PMLR, 2018. [3](#)
- [38] Prajit Ramachandran, Niki Parmar, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jon Shlens. Stand-alone self-attention in vision models. *Advances in Neural Information Processing Systems*, 32, 2019. [3](#), [8](#)
- [39] Sebastian Ruder. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 2016. [5](#)
- [40] Michael S Ryoo, AJ Piergiovanni, Anurag Arnab, Mostafa Dehghani, and Anelia Angelova. Tokenlearner: What can 8 learned tokens do for images and videos? *arXiv preprint arXiv:2106.11297*, 2021. [3](#)
- [41] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014. [5](#)
- [42] 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 preprint arXiv:2012.15460*, 2020. [1](#)
- [43] Yehui Tang, Kai Han, Yunhe Wang, Chang Xu, Jianyuan Guo, Chao Xu, and Dacheng Tao. Patch slimming for efficient vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12165–12174, 2022. [3](#)
- [44] Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. Efficient transformers: A survey. *arXiv preprint arXiv:2009.06732*, 2020. [3](#)
- [45] 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*, pages 5998–6008, 2017. [1](#), [2](#), [3](#), [7](#), [8](#)
- [46] 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*, pages 6578–6588, 2020. [6](#)
- [47] Apoorv Vyas, Angelos Katharopoulos, and François Fleuret. Fast transformers with clustered attention. *Advances in Neural Information Processing Systems*, 33:21665–21674, 2020. [3](#), [8](#)
- [48] 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 (CVPR)*, pages 1571–1580, June 2021. [1](#), [2](#), [3](#), [6](#)
- [49] Qiang Wang, Li Zhang, Luca Bertinetto, Weiming Hu, and Philip HS Torr. Fast online object tracking and segmentation: A unifying approach. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1328–1338, 2019. [2](#)
- [50] Shuohang Wang, Luwei Zhou, Zhe Gan, Yen-Chun Chen, Yuwei Fang, Siqi Sun, Yu Cheng, and Jingjing Liu. Clusterformer: Clustering-based sparse transformer for long-range dependency encoding. *arXiv preprint arXiv:2009.06097*, 2020. [3](#)
- [51] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018. [1](#), [2](#), [3](#)

- [52] Yi Wu, Jongwoo Lim, and Ming-Hsuan Yang. Online object tracking: A benchmark. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2411–2418, 2013. [1](#), [2](#), [6](#), [7](#)
- [53] Bin Yan, Houwen Peng, Jianlong Fu, Dong Wang, and Huchuan Lu. Learning spatio-temporal transformer for visual tracking. *arXiv preprint arXiv:2103.17154*, 2021. [1](#), [2](#), [3](#), [6](#), [7](#)
- [54] 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*, pages 15180–15189, 2021. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [55] Bin Yan, Xinyu Zhang, Dong Wang, Huchuan Lu, and Xiaoyun Yang. Alpha-refine: Boosting tracking performance by precise bounding box estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5289–5298, 2021. [2](#)
- [56] Tianyu Yang and Antoni B Chan. Learning dynamic memory networks for object tracking. In *Proceedings of the European conference on computer vision (ECCV)*, pages 152–167, 2018. [2](#)
- [57] Jiahui Yu, Yuning Jiang, Zhangyang Wang, Zhimin Cao, and Thomas Huang. Unitbox: An advanced object detection network. In *Proceedings of the 24th ACM international conference on Multimedia*, pages 516–520, 2016. [5](#)
- [58] Lichao Zhang, Abel Gonzalez-Garcia, Joost van de Weijer, Martin Danelljan, and Fahad Shahbaz Khan. Learning the model update for siamese trackers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4010–4019, 2019. [2](#)
- [59] Qinglong Zhang and Yu-Bin Yang. Rest: An efficient transformer for visual recognition. *Advances in Neural Information Processing Systems*, 34:15475–15485, 2021. [3](#)
- [60] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6848–6856, 2018. [8](#)
- [61] Zhipeng Zhang, Houwen Peng, Jianlong Fu, Bing Li, and Weiming Hu. Ocean: Object-aware anchor-free tracking. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pages 771–787. Springer, 2020. [5](#)