

HiFT: Hierarchical Feature Transformer for Aerial Tracking

Ziang Cao[†], Changhong Fu^{†,*}, Junjie Ye[†], Bowen Li[†], and Yiming Li[‡]
[†]Tongji University [‡]New York University

caoang233@gmail.com, changhongfu@tongji.edu.cn, yimingli@nyu.edu

Abstract

Most existing Siamese-based tracking methods execute the classification and regression of the target object based on the similarity maps. However, they either employ a single map from the last convolutional layer which degrades the localization accuracy in complex scenarios or separately use multiple maps for decision making, introducing intractable computations for aerial mobile platforms. Thus, in this work, we propose an efficient and effective hierarchical feature transformer (HiFT) for aerial tracking. Hierarchical similarity maps generated by multi-level convolutional layers are fed into the feature transformer to achieve the interactive fusion of spatial (shallow layers) and semantics cues (deep layers). Consequently, not only the global contextual information can be raised, facilitating the target search, but also our end-to-end architecture with the transformer can efficiently learn the interdependencies among multi-level features, thereby discovering a tracking-tailored feature space with strong discriminability. Comprehensive evaluations on four aerial benchmarks have proven the effectiveness of HiFT. Real-world tests on the aerial platform have strongly validated its practicability with a real-time speed. Our code is available at <https://github.com/vision4robotics/HiFT>.

1. Introduction

Visual object tracking¹, aiming to estimate the location of object frame by frame given the initial state, has drawn considerable attention due to its prosperous applications especially for unmanned aerial vehicles (UAVs), e.g., aerial cinematography [5], visual localization [48], and collision warning [19]. Despite the impressive progress, efficient and effective aerial tracking remains a challenging task due to limited computational resources and various difficulties like fast motion, low-resolution, frequent occlusion, etc.

In the visual tracking community, deep learning (DL)-based trackers [44, 35, 9, 2, 31, 53, 18, 17, 6] stand out on

*Corresponding Author

¹This work targets single object tracking (SOT).




Figure 1. Qualitative comparison of the proposed HiFT with state-of-the-arts [23, 8, 31] on three challenging sequences (*BMX4*, *RaceCar1* from DTB70 [34], and *Car16* from UAV20L [39]). Owing to the effective tracking-tailored feature space produced by the hierarchical feature transformer, our HiFT tracker can achieve robust performance under various challenges with a satisfactory tracking speed while other trackers lose effectiveness.

account of using the convolutional neural network (CNN) with robust representation capability. However, lightweight CNNs like AlexNet [30] can hardly extract robust features which are vital for tracking performance in complex aerial scenarios. Using a larger kernel size or a deeper backbone [31] can alleviate the aforementioned shortcoming yet the efficiency and practicability will be sacrificed. In literature, the dilated convolution [49] proposed to expand the receptive field and avoid the loss of resolution caused by the pooling layer. Unfortunately, it still suffers from unstable performance while handling small objects.

Recently, the transformer has demonstrated huge potential in many domains with an encoder-decoder structure [1]. Inspired by the superior performance of the transformer in modeling global relationships, we try to exploit its architecture in aerial tracking to effectively fuse multi-level² features to achieve promising performance. Meanwhile, the loss of efficiency caused by the computations of multiple layers and the deficiency of the transformer in handling

²We use the hierarchical feature to denote the feature maps from multiple convolutional layers.

small objects (pointed out in [52]) can be mitigated simultaneously.

In specific, since the target object in visual tracking can be an arbitrary object, the learned object queries in the original transformer structure hardly generalize well in visual tracking. Therefore, we adopt low-resolution features from the deeper layer to replace object queries. Meantime, we also feed the shallow layers into the transformer to discover a tracking-tailored feature space with strong discriminability by end-to-end training, which implicitly models the relationship of spatial information from high-resolution layers and semantic cues from low-resolution layers. Moreover, to further handle the insufficiency faced with low-resolution objects [52], we design a novel feature modulation layer in the transformer to fully explore the interdependencies among multi-level features. The proposed hierarchical feature transformer (HiFT) tracker has efficiently achieved robust performance under complex scenarios, as shown in Fig. 1. The main contributions of this work are as follows:

- We propose a novel hierarchical feature transformer to learn relationships amongst multi-level features, thereby discovering a tracking-tailored feature space with strong discriminability for aerial tracking.
- We design a neat feature modulation layer and classification label to further exploit the hierarchical features in Siamese networks and improve the tracking accuracy in handling the small objects.
- Comprehensive evaluations on four authoritative aerial benchmarks have validated the promising performance of HiFT against other state-of-the-art (SOTA) trackers, even those equipped with deeper backbones.
- Real-world tests are conducted on a typical aerial platform, demonstrating the superior efficiency and effectiveness of HiFT in real-world scenarios.

2. Related Works

2.1. Visual Tracking Methods

After MOSSE [4], a variety of achievements have been witnessed in handcrafted discriminative correlation filter (DCF)-based trackers [21, 36, 29, 9]. By calculating in the Fourier domain, DCF-based trackers can achieve competitive performance with high efficiency [20]. Nevertheless, those trackers hardly maintain robustness under various tracking conditions due to the poor representation ability of the handcrafted feature. To improve the tracking performance, several works introducing deep learning to DCF-based methods have been released [9, 50, 35]. Despite the great progress, they are still faced with inferior robustness and efficiency for aerial tracking.

Another outstanding branch in the SOT community is the Siamese-based methods [2, 24, 32, 53, 31], which benefit from massive offline training data and end-to-end learning strategy. As one of the pioneering works, SiameseFC [2] exposed the advantage of the Siamese framework, formulating the tracking task as the similarity matching process of template and search patches. Based on SiameseFC, DSiam [24] was proposed to effectively handle the object appearance variation and background interference. Inspired by region proposal network (RPN) [40], SiamRPN [32] considered tracking as two subtasks, applying the classification and regression branches respectively. DaSiamRPN [53] introduced a novel distractor-aware module and an effective sampling strategy, further promoting its robustness. More recently, the potential of adopting very deep networks as the backbone is extensively tapped [31], while the efficiency is sacrificed largely. Obviously, RPN-based trackers [32, 53, 31] provide an effective tracking strategy. However, the hyper-parameters associated with anchors significantly decrease the generalization of trackers. In order to eliminate such a drawback, the anchor-free method is proposed [23, 8].

In Siamese-based trackers, robust features make a vital influence on tracking performance. However, the trackers [2, 32, 53, 18] with lightweight backbone like AlexNet [30] suffer from the lack of global context while the trackers [31, 8, 23] utilizing deep CNN like ResNet [25] are far from real-time requirements onboard UAV. Albeit several works proposed to explore multi-level features in visual tracking [31, 16], they introduce cumbersome computation inevitably which is unaffordable for mobile platforms. Differently, this work proposes a brand-new lightweight hierarchical feature transformer (HiFT) for effective and efficient multi-level feature fusion, achieving robust aerial tracking efficiently.

2.2. Transformer in Computer Vision

Vaswani *et al.* [1] firstly proposed the transformers for machine translation based on the attention mechanism. Benefiting from its high representation ability, the transformer structure is expanded to the domain of computer vision such as video captioning [51], image enhancement [47], and pose estimation [27]. After DETR [7] initiates the research of transformer in object detection, deformable DETR [52] proposed the deformable attention module for efficiently convergence, providing inspirations about the combination of transformer and CNN. Some studies attempted to introduce the transformer to multi-object tracking and achieved promising performance [38], while the study of transformer in SOT is still blocked so far.

Although the attention mechanism in the transformer shows good performance in extensive visual tasks, its superiority struggles to be extended to SOT, since predefined

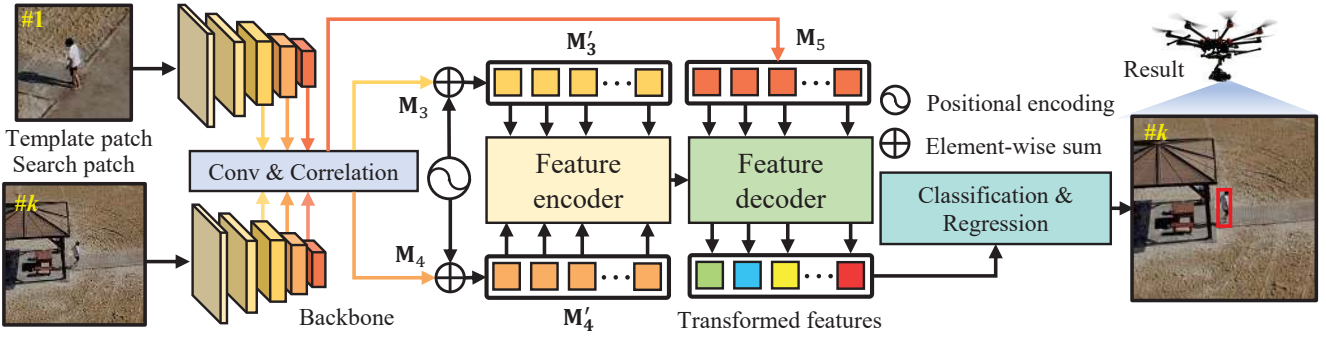


Figure 2. Overview of the HiFT tracker. The modules from the left to right are feature extraction network, hierarchical feature transformer, and classification & regression network. Three arrows with different colors represent the workflow of features from different layers respectively. Note that only the input of the encoder is combined with position encoding. Best viewed in color. (Image frames are from UAV20L [39].)

(or learned) object queries hardly maintain effectiveness when facing an arbitrary object. Moreover, the transformer hardly deals with the low-resolution object which is frequently encountered in aerial tracking. In this work, instead of redesigning object queries and related structures, we propose a hierarchical feature transformer to constructing a novel as well as robust tracking-tailored feature space. By virtue of the introduction of global context and interdependencies among multi-level features, the discriminability in the feature space is significantly raised to improve the tracking performance. Meanwhile, HiFT possesses a lightweight encoder-decoder structure which is desirable for mobile platforms.

3. Proposed Method

The workflow of HiFT is presented in Fig. 2. It can be divided into three submodules, feature extraction network, hierarchical feature transformer, and classification & regression network. Note that we utilize features from the last three layers to build the hierarchical feature transformation in this paper.

3.1. Feature Extraction Network

Deep CNNs, *e.g.*, ResNet [25], MobileNet [42], and GoogLeNet [43], have demonstrated their surprising capability, serving as popular feature extraction backbones in Siamese frameworks [31]. However, the heavy computation brought by the deep structure hardly be afforded by the aerial platform. To this concern, HiFT adopts a lightweight backbone, *i.e.*, AlexNet [30], which serves in both template and search branches. For clarity, the template/search images are respectively denoted by \mathbf{Z} and \mathbf{X} . $\phi_k(\mathbf{X})$ represents the k -th layer output of the search branch.

Remark 1: Despite the weaker feature extraction capability of AlexNet compared with those deeper networks, the proposed feature transformer can make up such a drawback significantly, at the same time saving computation resources for real-time aerial tracking.

3.2. Hierarchical Feature Transformer

The proposed hierarchical feature transformer can be mainly divided into two steps: high-resolution features encoding and low-resolution feature decoding. The former aims at learning interdependencies between different feature layers and spatial information to raise attention to objects with different scales (especially low-resolution objects). While the latter aggregates the semantic information from the low-resolution feature map. Benefiting from the abundant global context and interdependencies among hierarchical features, our method discovers a tracking-tailored feature space. Thus, the discriminability and representative capabilities of transformed features under various aerial tracking conditions are raised significantly. Specifically, features from the last three layers are utilized. The feature map from k -th layer is convoluted and reshaped to $\mathbf{M}_i \in \mathbb{R}^{W \times H \times C}$ (C , W , H represents the channel, width, and height of the feature map respectively) before being fed into the feature transformer:

$$\mathbf{M}_i = \mathcal{F}(\phi_i(\mathbf{Z}) \star \phi_i(\mathbf{X})) , i = 3, 4, 5 , \quad (1)$$

where \mathcal{F} denotes the convolution layer and \star represents the cross-correlation operator. Then, $\mathbf{M}'_3 \in \mathbb{R}^{W \times H \times C}$ and $\mathbf{M}'_4 \in \mathbb{R}^{W \times H \times C}$ can be obtained by supplementing with a learnable positional encoding.

3.2.1 Feature Encoding

To fully explore the interdependencies between hierarchical features, we use the combination of \mathbf{M}'_3 and \mathbf{M}'_4 as the input of multi-head attention module [1] as $\mathbf{M}^1_E = \text{Norm}(\mathbf{M}'_3 + \mathbf{M}'_4)$, where Norm represents the normalization layer. Generally, the scaled dot-product attention Att can be expressed by:

$$\text{Att}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{c}}\right)\mathbf{V} , \quad (2)$$

where \sqrt{c} is the scaling factor to avoid gradient vanishment in the softmax function. Then the calculation process of the multi-head attention module mAtt is expressed as:

$$\begin{aligned} \text{mAtt}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) &= \left(\text{Cat}(a^1, \dots, a^N) \right) \mathbf{W}_c, \\ a^j &= \text{Att}(\mathbf{Q}\mathbf{W}_1^j, \mathbf{K}\mathbf{W}_2^j, \mathbf{V}\mathbf{W}_3^j) \end{aligned} \quad (3)$$

where $\mathbf{W}_c \in \mathbb{R}^{C \times C}$, $\mathbf{W}_1^j \in \mathbb{R}^{C \times C_d}$, $\mathbf{W}_2^j \in \mathbb{R}^{C \times C_d}$, and $\mathbf{W}_3^j \in \mathbb{R}^{C \times C_d}$ ($C_d = C/N$, N is the number of parallel attention head) can all be regarded as fully connected layer operation. Please note that $\mathbf{Q}, \mathbf{K}, \mathbf{V}$ are only mathematical symbols to clarify the function. Therefore, they do not have practical meanings. Afterwards, the output of the first multi-head attention module, *i.e.*, $\mathbf{M}_E^2 \in \mathbb{R}^{WH \times C}$, can be obtained by:

$$\mathbf{M}_E^2 = \text{mAtt}(\mathbf{M}_E^1, \mathbf{M}_E^1, \mathbf{M}_3'). \quad (4)$$

As a result, the interdependencies between \mathbf{M}_3' and \mathbf{M}_4' are effectively learned to enrich the high-resolution feature map \mathbf{M}_E^2 . Besides, the global context in the two feature maps is also introduced in \mathbf{M}_E^2 . After that, we construct the modulation layer to fully explore the potential of interdependencies between \mathbf{M}_E^3 and \mathbf{M}_4' whose structure is shown in Fig. 3. Specifically, the input of modulation layer \mathbf{M}_E^3 is obtained by normalization of \mathbf{M}_3' and \mathbf{M}_E^2 , *i.e.*, $\mathbf{M}_E^3 = \text{Norm}(\mathbf{M}_3' + \mathbf{M}_E^2)$. After a feed-forward network (FFN) and global average pooling operation (GAP), the output of modulation layer \mathbf{M}_E^4 can be formulated as:

$$\begin{aligned} \mathbf{W}' &= \mathcal{F}(\text{Cat}(\mathbf{M}_E^3, \mathbf{M}_4')) * \text{FFN}(\text{GAP}(\mathbf{M}_4')) \\ \mathbf{M}_E^4 &= \mathbf{M}_E^3 + \gamma_1 * \mathbf{W}' * \mathbf{M}_E^3 \end{aligned} \quad (5)$$

where γ_1 represents a learning weight.

Owing to the modulation layer, the internal spatial information between \mathbf{M}_4' and \mathbf{M}_E^3 are exploited efficiently, thereby effectively distinguishing the object from the complex background. Eventually, the encoded information can be calculated through FFN and normalization.

Remark 2: Attributing to the feature encoder, the global context and interdependencies between \mathbf{M}_3' and \mathbf{M}_4' are fully explored. Additionally, to overcome the deficiency of handling small objects, the modulation layer is proposed to further explore spatial information for enriching the encoded information. Finally, based on it, the decoder can build an effective feature transformation for robust tracking.

3.2.2 Feature Decoding

Before decoding, the low-resolution feature map is firstly reshaped to $\mathbf{M}_5 \in \mathbb{R}^{WH \times C}$ in Eq. (1). The feature decoder follows the similar structure of standard transformer [1]. Differently, we build the effective feature decoder without position encoding. Since we treat the number of locations as

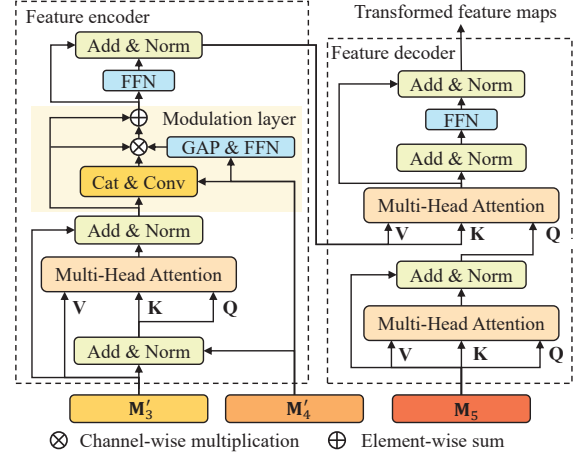


Figure 3. Detailed workflow of HiFT. The left sub-window illustrates the feature encoder. The right one shows the structure of the decoder. Best viewed in color.

the sequence length in our method, the position encoding is introduced to distinguish each location on feature maps. For avoiding the direct influence on the transformed feature, we decide to introduce the position information through the encoder implicitly. Analysis of the positional encoding strategy is conducted later in Sec. 4.3.3. The structure of the decoder is exhibited in Fig. 3.

Remark 3: By the hierarchical feature transformer, the spatial/semantic information in the high-/low-resolution features is fully utilized to improve the discriminability of the final transformed feature. Meanwhile, the modulation layer achieves the aggregation of interdependencies among different feature layers, enhancing the robustness of tracking objects with various scales.

3.3. Definition of Classification Label

The structures of classification and regression are implemented by several convolution layers. To achieve accurate classification, we apply two classification branches. One branch aims to classify via the area involved in the ground truth box. The other branch focuses on determining the positive samples measured by the distance between the center of ground truth and the corresponding point. Besides, to accelerate the convergence, we use pseudo-random number generators denoted as \mathcal{T} to constrain the number of negative labels.

Remark 4: The detailed calculation process of classification and regression can be found in the supplementary material.

Therefore, the overall loss function can be determined as:

$$L_{\text{overall}} = \lambda_1 L_{\text{cls1}} + \lambda_2 L_{\text{cls2}} + \lambda_3 L_{\text{loc}}, \quad (6)$$

where L_{cls1} , L_{cls2} , L_{loc} represent the cross-entropy, binary cross-entropy, and IoU loss. λ_1 , λ_2 , and λ_3 are the coefficients to balance the contributions of each loss.

4. Experiments

4.1. Implementation Details

During the training of 70 epochs, the last three layers of AlexNet are fine-tuned in the last 60 epochs while the first two layers are frozen. The learning rate is initialized as 5×10^{-4} and decreased in the log space from 10^{-2} to 10^{-4} . Besides, the sizes of \mathbf{Z} and \mathbf{X} are set to $3 \times 127 \times 127$ and $3 \times 287 \times 287$ respectively. The feature transformer consists of one encoder layer and two decoder layers. We use image pairs extracted from COCO [37], ImageNet VID [41], GOT-10K [28], and Youtube-BB [15] to train HiFT. In addition, the stochastic gradient descent (SGD) is adopted, and batch size, momentum, and weight decay are set to 220, 0.9, and 10^{-4} , respectively. Our tracker is trained on a PC with an Intel i9-9920X CPU, a 32GB RAM, and two NVIDIA TITAN RTX GPUs. More experimental results can be found in the supplementary.

4.2. Evaluation Metrics

The one-pass evaluation (OPE) metrics [39] including precision and success rate are applied to assess the tracking performance. Specifically, the success rate is measured by the intersection over union (IoU) of the ground truth and estimated bounding boxes. The percentage of frames whose IoU is beyond a pre-defined threshold is drawn as the success plot (SP). Besides, the center location error (CLE) between the estimated location and the ground truth is employed to evaluate the precision. The percentage of frames whose CLE is within a certain threshold is drawn as the precision plot (PP). Meanwhile, the area under the curve (AUC) of the SP and the precision at a threshold of 20 pixels is adopted to rank the trackers.

4.3. Evaluation on Aerial Benchmarks

4.3.1 Overall Performance

For overall evaluation, HiFT is tested on four challenging and authoritative aerial tracking benchmarks, and comprehensively compared with other 19 state-of-the-art (SOTA) trackers including SiamRPN++ [31], DaSiamRPN [53], UDT [44], UDT+ [44], TADT [35], CoKCF [50], ARCF [29], AutoTrack [36], ECO [9], C-COT [13], MCCT [45], DeepSTRCF [33], STRCF [33], BACF [21], SRDCF [11], fDSST [12], SiameseFC [2], DSiam [24], and KCF [26]. For fairness, all the Siamese-based trackers adopt the same backbone, *i.e.*, AlexNet [30] pre-trained on ImageNet [41].

UAV123 [39]: UAV123 is a large-scale UAV benchmark including 123 high-quality sequences with more than 112K frames which covers a variety of challenging aerial scenarios such as frequent occlusion, low resolution, out-of-view, *etc.* Therefore, UAV123 can help to exhaustively assess

tracking performance in aerial tracking. As illustrated in Table 1, HiFT outperforms other trackers in both precision and success. In terms of precision, HiFT gains first place with a precision score of 0.787, surpassing the second- and third-best SiamRPN++ (0.769) and ECO (0.752) by 2.3% and 4.7% respectively. As for the success rate, HiFT (0.589) also improves over SiamRPN++ (0.579) and ranks first place. In a word, HiFT demonstrates superior performance in all kinds of aerial tracking scenarios.

UAV20L [39]: UAV20L is composed of 20 long-term tracking sequences with 2934 frames on average and over 58K frames in total. In this paper, it is utilized to evaluate our tracker in realistic long-term aerial tracking scenes. As presented in Table 2, attributing to the global contextual information introduced by the feature transformer, our tracker achieves competitive performance compared to other SOTA trackers. Specifically, HiFT yields the best precision score (0.763), surpassing the second-best SiamRPN++ (0.696) and the third-best DaSiamRPN (0.665) by **9.6%** and **14.7%**. Similarly, in success rate, HiFT achieves the best score (0.566), followed by SiamRPN++ (0.528) and DaSiamRPN (0.465). The extraordinary performance verifies that HiFT could be a desirable choice in long-term aerial tracking scenarios.


DTB70 [34]: Compared to the aforementioned two benchmarks, DTB70 contains 70 challenging UAV sequences with a large number of severe motion scenes. The robustness of trackers in fast motion scenarios could be appropriately evaluated on this benchmark. Experimental results are shown in Fig. 4a, HiFT ranks first place in both precision (0.802) and success rate (0.594), followed by SiamRPN++

Table 1. Quantitative evaluation on UAV123 [39]. The top three performances are respectively highlighted by **red**, **green**, and **blue** color. Prec. and Succ. respectively denote precision score at 20 pixels and AUC of success plot.

Trackers	Prec.	Succ.	Trackers	Prec.	Succ.
AutoTrack [36]	0.689	0.472	C-COT [13]	0.729	0.502
ARCF [29]	0.671	0.468	UDT+ [44]	0.732	0.502
STRCF [33]	0.681	0.481	UDT [44]	0.668	0.477
fDSST [12]	0.583	0.405	TADT [35]	0.727	0.520
SRDCF [11]	0.676	0.463	DeepSTRCF [33]	0.705	0.508
CoKCF [50]	0.652	0.399	MCCT [45]	0.734	0.507
KCF [26]	0.523	0.331	DSiam [24]	0.608	0.400
BACF [21]	0.662	0.461	ECO [9]	0.752	0.528
SiamRPN++ [31]	0.769	0.579	SiameseFC [2]	0.725	0.494
DaSiamRPN [53]	0.725	0.501	HiFT (ours)	0.787	0.589

Table 2. Overall evaluation on UAV20L [39]. The top nine trackers are reported. The top three trackers are respectively marked by **red**, **green**, and **blue** color. Prec. and Succ. respectively denote precision score at 20 pixels and AUC of success plot.

	UDT+ [44]	ECO [9]	TADT [35]	DeepST-RCF [33]	Siamese-eFC [2]	DSiam [24]	DaSiam RPN [53]	SiamRP-N++ [31]	HiFT (ours)
Prec.	0.585	0.589	0.609	0.588	0.599	0.603	0.665	0.696	0.763
Succ.	0.401	0.427	0.459	0.443	0.402	0.391	0.465	0.528	0.566



(a) Results on DTB70.

(b) Results on UAV123@10fps.

Figure 4. PPs and SPs of HiFT and other SOTA trackers on (a) DTB70 and (b) UAV123@10fps. Our tracker achieves superior performance in the two benchmarks

Table 3. Average evaluation on four aerial tracking benchmarks. Our tracker outperforms all other trackers with an obvious improvement. The best three performances are respectively highlighted with red, green, and blue color.

Trackers	HiFT (ours)	SiamRPN++ [31]	DaSiamRPN [53]	AutoTrack [36]	ARCF [29]	C-COT [13]	SiameseFC [2]	UDT+ [44]	TADT [35]	DeepSTRCF [33]	MCCT [45]	ECO [9]
Avg. Prec.	0.776	0.750	0.693	0.648	0.643	0.691	0.680	0.662	0.678	0.677	0.686	0.693
Avg. Succ.	0.581	0.563	0.480	0.445	0.448	0.479	0.463	0.461	0.488	0.489	0.472	0.494

with a precision of 0.795 and a success rate of 0.589. The promising ability of HiFT in handling fast motion can be attributed to the proposed hierarchical feature transformer which is able to promote the discrimination ability of HiFT. **UAV123@10fps [39]:** UAV123@10fps is created by down-sampling from the original 30FPS recording. Consequently, the issue of strong motion in UAV123@10fps is more severe compared to UAV123. The PPs and SPs shown in Fig. 4b demonstrate that HiFT can consistently obtain satisfactory performance, achieving the best precision (0.754) and success rate (0.574). To sum up, HiFT provides a more stable performance comparing to other SOTA trackers, verifying its favorable robustness in various aerial tracking scenarios.

Remark 5: Table 3 reports the average precision and success rate of the top 11 trackers on four benchmarks. It shows that HiFT has improved the second-best tracker SiamRPN++ by 3.5% and 3.2% in precision and success rate respectively.

Table 4. Attribute-based evaluation of top 6 trackers on four benchmarks. The best two performances are respectively highlighted by red and green color. HiFT keeps achieving the best performance in different attributes. Δ denotes the improvement in comparison with the second best tracker.

Attributes	Low-resolution		Scale variation		Occlusion		Fast motion	
	Prec.	Succ.	Prec.	Succ.	Prec.	Succ.	Prec.	Succ.
SiamRPN++	0.591	0.390	0.728	0.559	0.601	0.405	0.680	0.489
DaSiamRPN	0.592	0.347	0.678	0.482	0.583	0.361	0.617	0.409
C-COT	0.586	0.331	0.643	0.451	0.571	0.359	0.644	0.411
TADT	0.604	0.366	0.632	0.466	0.598	0.387	0.628	0.412
ECO	0.581	0.343	0.644	0.471	0.583	0.375	0.620	0.407
HiFT (ours)	0.626	0.416	0.772	0.584	0.638	0.431	0.751	0.537
Δ (%)	3.63	6.81	5.98	4.40	6.20	6.43	10.42	9.79

4.3.2 Attribute-based Comparison

To exhaustively evaluate HiFT under various challenges, attribute-based comparisons are conducted, seen in Table 4. HiFT ranks first place in terms of both precision and success rate in comparison with other top 5 trackers. Specifically, HiFT significantly exceeds the second-best performance in attributes of low-resolution, scale variation, occlusion, and fast motion. HiFT improves the second-best performance by around 10% in fast motion scenarios. The satisfactory results demonstrate that our hierarchical feature transformer can help exploit the global contextual information to overcome issues of severe motion. In addition, when the objects are severely occluded, HiFT can learn more robust features to discriminate the occluded objects. Therefore, HiFT also yields prominent improvement in the scenarios of occlusion. Moreover, since the multi-scale feature maps are utilized for building the feature transformation, our tracker is endowed with the ability to track objects with various scales, as verified by its performance in the attributes of low-resolution and scale variation.

4.3.3 Ablation Study

To verify the effectiveness of each module of the proposed method, detailed studies amongst HiFT with different modules enabled are conducted on UAV20L.

Symbol introduction: For clarity, we first introduce the meaning of symbols used in Table 5. This work considers Baseline as the model with only feature extraction and regression & classification network. OT denotes original standard transformer (with object query). FT indicates the original transformer with the feature map (instead of object query) but without the proposed modulation layer. HFT

Table 5. Ablation study of different components of HiFT. For the detailed explanation of Baseline, OT, FT, HFT, PE, and RL, please kindly refer to the text in Sec. 4.3.3. Δ denotes the improvement compared with the Baseline tracker.

Trackers	Precision	Δ_{pre} (%)	Success	Δ_{suc} (%)
Baseline	0.611	–	0.463	–
Baseline+OT	0.597	-2.29	0.446	-3.67
Baseline+FT	0.675	+10.47	0.496	+7.13
Baseline+HFT+PE	0.689	+12.77	0.523	+12.96
Baseline+HFT+RL	0.629	+2.95	0.486	+4.97
Baseline+HFT (HiFT)	0.763	+24.88	0.566	+22.25

denotes the full version of the proposed hierarchical feature transformer. PE represents direct positional encoding to M_5 (HiFT leaves out position encoding in M_5 as demonstrated in Sec. 3.2.2). RL represents the rectangle label used in the traditional trackers. For fairness, each version of the tracker adopts the same training strategy except for the investigated module.

Discussion on transformer architecture: As shown in Table 5, adding the original transformer with object queries (Baseline+OT) directly lowers the performance of Baseline by about 2.29% on precision and 3.67% on success rate, which proves that object queries hardly perform well in SOT with novel target objects. Replacing object query with the feature map, Baseline+FT raises tracking precision by **10.47%**. Further adopting the modulation layer, Baseline+HFT, yields the best performance by **24.88%**. All the aforementioned results can be com-

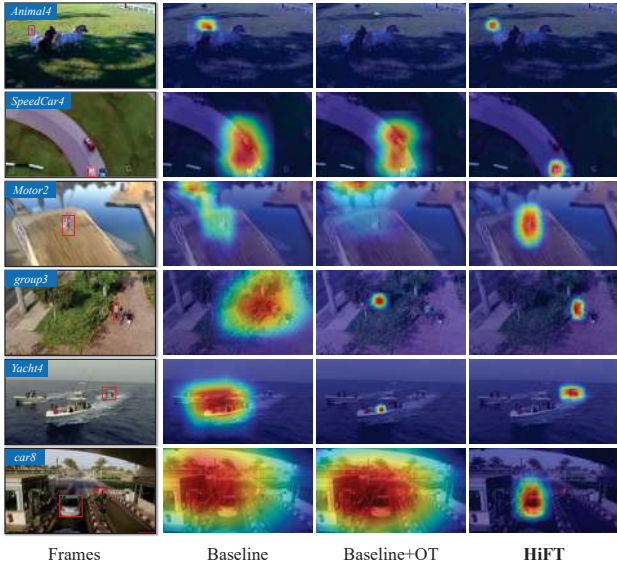


Figure 5. Visualization of the confidence map of three tracking methods on several sequences from UAV20L [39] and DTB70 [34]. The target objects are marked out by red boxes in the original frames. HiFT gets more robust performance for visual tracking in the air.

pared together to validate the efficacy of the elaborately designed hierarchical feature transformer with the modulation layer in aerial tracking.

Discussion on position encoding&classification label:

This part aims at proving the 2 strategies, position encoding in Sec. 3.2.2 and new classification label in Sec. 3.3. For position decoding, in Table 5, the tracker Baseline+HFT+PE hurts the performance of HiFT tremendously (from **24.88%** improvements to **12.77%**), proving that direct position encoding is indeed not proper for feature M_5 . Considering the distance of ground truth and sample points, the circular strategy utilized in HiFT achieves a notable improvement (**24.88%**) compared to the traditional rectangle label in Baseline+HFT+RL (**2.95%**).

Remark 6: Please note that more ablation studies are reported in supplement material.

4.3.4 Qualitative Evaluation

As shown in Fig. 5, the confidence map of our HiFT tracker consistently focuses on the object under onerous challenges in aerial tracking, e.g., fast motion in *Motor2*, low resolution in *SpeedCar4*, and occlusion in *group3* and *Yacht4*. Despite that the Baseline and Baseline+OT are trained with the same strategy as HiFT, they still fail to concentrate on the target object in those complex tracking scenarios, which proves the robustness of the proposed hierarchical feature transformer.

4.3.5 Comparison to Trackers with Deeper Backbone

The proposed hierarchical feature transformer dedicates to model effective feature mapping among multi-level features, so as to achieve SOTA performance without introducing a huge computational burden. To further evaluate its effectiveness, we employ the trackers equipped with deeper backbones for comparison. The state-of-the-art trackers, including SiamRPN++ (ResNet-50) [31], SiamRPN++ (MobileNet) [31], SiamMask (ResNet-50) [46], ATOM (ResNet-18) [10], DiMP (ResNet-50) [3], PrDiMP (ResNet-18) [14], SiamCAR (ResNet-50) [23], SiamGAT (GoogleNet) [22], and SiamBAN (ResNet-50) [8], are involved in the comparison. As illustrated in Fig. 6, HiFT achieves a satisfactory balance of tracking robustness and speed. On UAV20L, adopting AlexNet as the backbone, HiFT (0.763) surpasses the second-best tracker SiamRPN++_ResNet-50 (0.749) in precision and achieves a speed of 127 FPS, which is **1.8** times faster than the latter. Similarly, on DTB70, HiFT achieves comparable performance compared to those deeper CNN-based trackers. Eventually, the average precision and tracking speed are reported in Table 6, HiFT yields the best average precision (0.783) with a promising speed of **129.87** FPS, proving that




Figure 6. Precision-speed trade-off analysis by quantitative comparison between HiFT and trackers with deeper backbone on UAV20L [39] (left) and DTB70 [34] (right). Our method realizes an excellent trade-off on both two benchmarks.

Table 6. Average precision and tracking speed of HiFT and the trackers with deeper backbone. The proposed approach runs at a satisfactory speed of ~ 130 FPS, while achieving comparable tracking performance with those trackers equipped with a deeper backbone. The best three performances are respectively highlighted with **red**, **green**, and **blue** color.

Tracker	HiFT (ours)	SiamGAT [22]	SiamCAR [23]	SiamBAN [8]	PrDiMP [14]	SiamRPN++ [31]	SiamRPN++ [31]	SiamMask [46]	ATOM [10]
Backbone	AlexNet	GoogleNet	ResNet-50	ResNet-50	ResNet-18	ResNet-50	MobileNet	ResNet-50	ResNet-18
Avg. Prec.	0.783	0.751	0.739	0.763	0.741	0.774	0.748	0.740	0.738
Avg. FPS	129.87	90.01	71.74	73.25	25.94	71.59	115.03	77.30	34.94

HiFT achieves an excellent balance between tracking performance and efficiency.

5. Real-World Tests

In this section, HiFT is further implemented on a typical UAV platform including an embedded onboard processor, *i.e.*, NVIDIA AGX Xavier, to testify its practicability in real-world applications. Figure 7 presents three tests in the wild, including day and night scenes. The main challenges in the tests are partial occlusion, viewpoint change (the first row), low-resolution, camera motion (the second row), small object, and similar object around (the third row). Attributing to the effective feature transformer, HiFT maintains satisfying tracking robustness in various challenging scenarios. Moreover, our tracker remains at an average speed of **31.2** FPS during the tests without using TensorRT. Therefore, the real-world tests onboard the embedded system directly validate the superior performance and efficiency of HiFT under various UAV-specific challenges.

6. Conclusion

In this work, a novel hierarchical feature transformer for efficient aerial tracking is proposed for streamlining the process of exploiting the global contextual information and multi-level features. By virtue of both low-resolution semantics information and high-resolution spatial details, the transformed feature can achieve promising performance in discriminating the object from clutters via a lightweight structure. Meanwhile, attributing to the modulation layer and the new classification label, the effectiveness of the feature transformer can reach its full potential. Comprehensive experiments have validated that HiFT can achieve an excellent precision-speed trade-off and can be utilized in real-world aerial tracking scenarios. Moreover, even compared

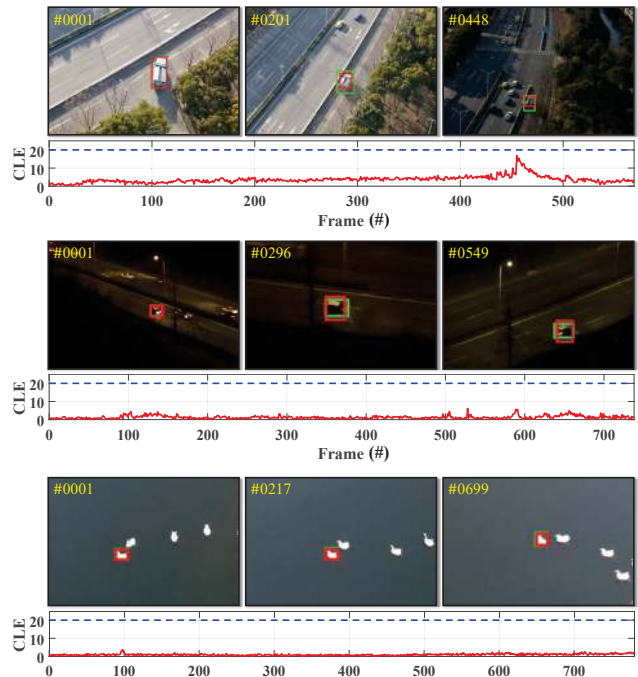


Figure 7. Visualization of real-world tests on the embedded platform. The tracking results and ground truth are marked with **red** and **green** boxes. The CLE score below the **blue** dotted line is considered as the success tracking result in the real-world tests.

to the trackers with deeper backbones, HiFT can achieve comparable performance. We are convinced that our work can advance the development of aerial tracking and promote the real-world applications of visual tracking.

Acknowledgment: This work is supported by the National Natural Science Foundation of China (No. 61806148) and the Natural Science Foundation of Shanghai (No. 20ZR1460100). We thank the anonymous reviewers for their efforts to help us improve our work.

References

- [1] V. Ashish, S. Noam, P. Niki, U. Jakob, J. Llion, N. G. Aidan, K. Lukasz, and P. Illia. Attention Is All You Need. In *Advances in neural information processing systems (NIPS)*, pages 6000–6010, 2017. 1, 2, 3, 4
- [2] L. Bertinetto, J. Valmadre, J. F. Henriques, A. Vedaldi, and P. H. Torr. Fully-Convolutional Siamese Networks for Object Tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 850–865, 2016. 1, 2, 5, 6
- [3] G. Bhat, M. Danelljan, L. Van Gool, and R. Timofte. Learning discriminative model prediction for tracking. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 6181–6190, 2019. 7
- [4] D. S. Bolme, J. R. Beveridge, B. A. Draper, and Y. M. Lui. Visual Object Tracking Using Adaptive Correlation Filters. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2544–2550, 2010. 2
- [5] R. Bonatti, C. Ho, W. Wang, S. Choudhury, and S. Scherer. Towards a Robust Aerial Cinematography Platform: Localizing and Tracking Moving Targets in Unstructured Environments. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 229–236, 2019. 1
- [6] Z. Cao, C. Fu, J. Ye, B. Li, and Y. Li. SiamAPN++: Siamese Attentional Aggregation Network for Real-Time UAV Tracking. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1–7, 2021. 1
- [7] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko. End-to-End Object Detection with Transformers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 213–229, 2020. 2
- [8] Z. Chen, B. Zhong, G. Li, S. Zhang, and R. Ji. Siamese Box Adaptive Network for Visual Tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6668–6677, 2020. 1, 2, 7, 8
- [9] M. Danelljan, G. Bhat, F. S. Khan, and M. Felsberg. ECO: Efficient Convolution Operators for Tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6931–6939, 2017. 1, 2, 5, 6
- [10] M. Danelljan, G. Bhat, F. S. Khan, and M. Felsberg. ATOM: Accurate Tracking by Overlap Maximization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4655–4664, 2019. 7, 8
- [11] M. Danelljan, G. Häger, F. S. Khan, and M. Felsberg. Learning Spatially Regularized Correlation Filters for Visual Tracking. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 4310–4318, 2015. 5
- [12] M. Danelljan, G. Häger, F. S. Khan, and M. Felsberg. Discriminative Scale Space Tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(8):1561–1575, 2017. 5
- [13] M. Danelljan, A. Robinson, F. S. Khan, and M. Felsberg. Beyond Correlation Filters: Learning Continuous Convolution Operators for Visual Tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 472–488, 2016. 5, 6
- [14] M. Danelljan, L. Van Gool, and R. Timofte. Probabilistic Regression for Visual Tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7181–7190, 2020. 7, 8
- [15] E. Real, J. Shlens, S. Mazzocchi, X. Pan, and V. Vanhoucke. YouTube-BoundingBoxes: A Large High-Precision Human-Annotated Data Set for Object Detection in Video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7464–7473, 2017. 5
- [16] H. Fan and H. Ling. Siamese Cascaded Region Proposal Networks for Real-Time Visual Tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7944–7953, 2019. 2
- [17] C. Fu, Z. Cao, Y. Li, J. Ye, and C. Feng. Onboard Real-Time Aerial Tracking With Efficient Siamese Anchor Proposal Network. *IEEE Transactions on Geoscience and Remote Sensing*, pages 1–13, 2021. 1
- [18] C. Fu, Z. Cao, Y. Li, J. Ye, and C. Feng. Siamese Anchor Proposal Network for High-Speed Aerial Tracking. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–7, 2021. 1, 2
- [19] C. Fu, A. Carrio, M. A. Olivares-Mendez, R. Suarez-Fernandez, and P. Campoy. Robust real-time vision-based aircraft tracking from Unmanned Aerial Vehicles. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 5441–5446, 2014. 1
- [20] Changhong Fu, Bowen Li, Fangqiang Ding, Fuling Lin, and Geng Lu. Correlation Filter for UAV-Based Aerial Tracking: A Review and Experimental Evaluation. *IEEE Geoscience and Remote Sensing Magazine*, pages 2–37, 2021. 2
- [21] H. K. Galoogahi, A. Fagg, and S. Lucey. Learning Background-Aware Correlation Filters for Visual Tracking. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1144–1152, 2017. 2, 5
- [22] D. Guo, Y. Shao, Y. Cui, Z. Wang, L. Zhang, and C. Shen. Graph Attention Tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–10, 2021. 7, 8
- [23] D. Guo, J. Wang, Y. Cui, Z. Wang, and S. Chen. SiamCAR: Siamese Fully Convolutional Classification and Regression for Visual Tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6268–6276, 2020. 1, 2, 7, 8
- [24] Q. Guo, W. Feng, C. Zhou, R. Huang, L. Wan, and S. Wang. Learning Dynamic Siamese Network for Visual Object Tracking. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1781–1789, 2017. 2, 5
- [25] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 2, 3
- [26] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista. High-Speed Tracking with Kernelized Correlation Filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(3):583–596, 2015. 5

- [27] L. Huang, J. Tan, J. Liu, and J. Yuan. Hand-Transformer: Non-Autoregressive Structured Modeling for 3D Hand Pose Estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 17–33, 2020. 2
- [28] L. Huang, X. Zhao, and K. Huang. GOT-10k: A Large High-Diversity Benchmark for Generic Object Tracking in the Wild. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–17, 2019. 5
- [29] Z. Huang, C. Fu, Y. Li, F. Lin, and P. Lu. Learning Aberrance Repressed Correlation Filters for Real-time UAV Tracking. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2891–2900, 2019. 2, 5, 6
- [30] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1097–1105, 2012. 1, 2, 3, 5
- [31] B. Li, W. Wu, Q. Wang, F. Zhang, J. Xing, and J. Yan. SiamRPN++: Evolution of Siamese Visual Tracking With Very Deep Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4277–4286, 2019. 1, 2, 3, 5, 6, 7, 8
- [32] B. Li, J. Yan, W. Wu, Z. Zhu, and X. Hu. High Performance Visual Tracking with Siamese Region Proposal Network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8971–8980, 2018. 2
- [33] F. Li, C. Tian, W. Zuo, L. Zhang, and M. Yang. Learning Spatial-Temporal Regularized Correlation Filters for Visual Tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4904–4913, 2018. 5, 6
- [34] S. Li and D. Yeung. Visual Object Tracking for Unmanned Aerial Vehicles: A Benchmark and New Motion Models. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 1–7, 2017. 1, 5, 7, 8
- [35] X. Li, C. Ma, B. Wu, Z. He, and M. Yang. Target-Aware Deep Tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1369–1378, 2019. 1, 2, 5, 6
- [36] Y. Li, C. Fu, F. Ding, Z. Huang, and G. Lu. AutoTrack: Towards High-Performance Visual Tracking for UAV With Automatic Spatio-Temporal Regularization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11920–11929, 2020. 2, 5, 6
- [37] T. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *Proceedings of the European conference on computer vision (ECCV)*, pages 740–755, 2014. 5
- [38] T. Meinhardt, A. Kirillov, L. Laura, and C. Feichtenhofer. TrackFormer: Multi-Object Tracking with Transformers. *arXiv preprint arXiv:2101.02702*, 2021. 2
- [39] M. Mueller, N. Smith, and B. Ghanem. A Benchmark and Simulator for UAV Tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 445–461, 2016. 1, 3, 5, 6, 7, 8
- [40] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 91–99, 2015. 2
- [41] O. Russakovsky, J. Deng, H. Su, J. Krause, et al. Imagenet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015. 5
- [42] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. Chen. Mobilenetv2: Inverted Residuals and Linear Bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4510–4520, 2018. 3
- [43] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going Deeper with Convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–9, 2015. 3
- [44] N. Wang, Y. Song, C. Ma, W. Zhou, W. Liu, and H. Li. Unsupervised Deep Tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1308–1317, 2019. 1, 5, 6
- [45] N. Wang, W. Zhou, Q. Tian, R. Hong, M. Wang, and H. Li. Multi-cue Correlation Filters for Robust Visual Tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4844–4853, 2018. 5, 6
- [46] Q. Wang, L. Zhang, L. Bertinetto, W. Hu, and P. H. S. Torr. Fast online object tracking and segmentation: A unifying approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1328–1338, 2019. 7, 8
- [47] F. Yang, H. Yang, J. Fu, H. Lu, and B. Guo. Learning Texture Transformer Network for Image Super-Resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5791–5800, 2020. 2
- [48] J. Ye, C. Fu, F. Lin, F. Ding, S. An, and G. Lu. Multi-Regularized Correlation Filter for UAV Tracking and Self-Localization. *IEEE Transactions on Industrial Electronics*, pages 1–10, 2021. 1
- [49] F. Yu and V. Koltun. Multi-Scale Context Aggregation by Dilated Convolutions. In *Proceedings of the International Conference on Learning Representations (ICLR)*, pages 1–9, 2016. 1
- [50] L. Zhang and P. N. Suganthan. Robust Visual Tracking via Co-Trained Kernelized Correlation Filters. *Pattern Recognition*, 69:82–93, 2017. 2, 5
- [51] L. Zhou, Y. Zhou, J. J. Corso, R. Socher, and C. Xiong. End-to-End Dense Video Captioning with Masked Transformer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8739–8748, 2018. 2
- [52] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai. Deformable DETR: Deformable Transformers for End-to-End Object Detection. *arXiv preprint arXiv:2010.04159*, 2020. 2
- [53] Z. Zhu, Q. Wang, B. Li, W. Wu, J. Yan, and W. Hu. Distractor-Aware Siamese Networks for Visual Object Tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 101–117, 2018. 1, 2, 5, 6