PVT++: A Simple End-to-End Latency-Aware Visual Tracking Framework

Bowen Li[1,*], Ziyuan Huang[2,*], Junjie Ye[3], Yiming Li[4], Sebastian Scherer[1], Hang Zhao[5], Changhong Fu[3, ✉]

[1]Carnegie Mellon University, [2]National University of Singapore, [3]Tongji University
[4]New York University, [5]Tsinghua University

{bowenli2, basti}@andrew.cmu.edu, ziyuan.huang@u.nus.edu
{ye.jun.jie, changhongfu}@tongji.edu.cn, yimingli@nyu.edu, hangzhao@tsinghua.edu.cn

Abstract

Visual object tracking is essential to intelligent robots. Most existing approaches have ignored the online latency that can cause severe performance degradation during real-world processing. Especially for unmanned aerial vehicles (UAVs), where robust tracking is more challenging and onboard computation is limited, the latency issue can be fatal. In this work, we present a simple framework for end-to-end latency-aware tracking, i.e., end-to-end predictive visual tracking (PVT++). Unlike existing solutions that naively append Kalman Filters after trackers, PVT++ can be jointly optimized, so that it takes not only motion information but can also leverage the rich visual knowledge in most pre-trained tracker models for robust prediction. Besides, to bridge the training-evaluation domain gap, we propose a relative motion factor, empowering PVT++ to generalize to the challenging and complex UAV tracking scenes. These careful designs have made the small-capacity lightweight PVT++ a widely effective solution. Additionally, this work presents an extended latency-aware evaluation benchmark for assessing an any-speed tracker in the online setting. Empirical results on a robotic platform from the aerial perspective show that PVT++ can achieve significant performance gain on various trackers and exhibit higher accuracy than prior solutions, largely mitigating the degradation brought by latency. Our code is public at https://github.com/Jaraxxus-Me/PVT_pp.git.

1. Introduction

Visual object tracking[1] is fundamental for many robotic applications like navigation [49], cinematography [5], and multi-agent cooperation [9]. Most existing trackers are developed and evaluated under an *offline* setting [38, 29, 34, 33, 6, 8], where the trackers are assumed to have zero processing time. However, in real-world deployment, the on-
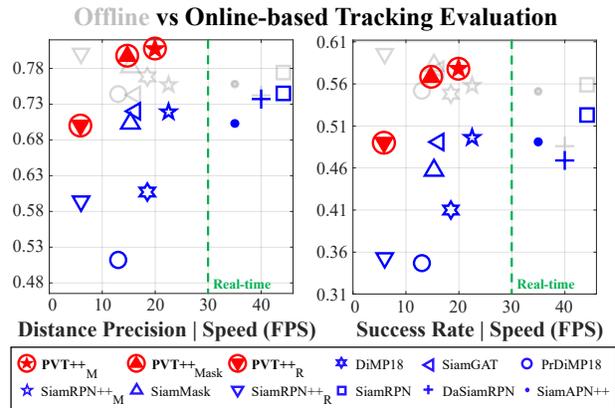
---

[1]We focus on single object tracking in this work.

Figure 1. Distance precision and success rate of the trackers on UAVDT dataset [16]. Compared with offline evaluation, the trackers suffer a lot from their onboard latency in the **online** setting (30 frames/s (FPS)). Coupled with **PVT++**, the predictive trackers achieve significant performance gain with very little extra latency, obtaining on par or better results than the offline setting.

line latency caused by the trackers' processing time cannot be ignored, since the world would have already changed when the trackers finish processing the captured frame. In particular, with limited onboard computation, this issue is more critical in the challenging unmanned aerial vehicle (UAV) tracking scenes [21, 38, 20]. As shown in Fig. 1, compared with *offline* setting (gray markers), the latency can cause severe performance degradation during *online* processing (colored markers). If not handled well, this can easily lead to the failure of robotic applications such as UAV obstacle avoidance [1] and self-localization [61].

To be more specific, the latency hurts online tracking due to: (1) The tracker outputs are always outdated, so there will be mismatch between the tracker result and world state. (2) The trackers can only process the latest frame, so that the non-real-time ones may skip some frames, which makes object motion much larger (see Fig. 2(a) right).

The existence of the latency in real-world applications calls for trackers with prediction capabilities, *i.e.,* predic-
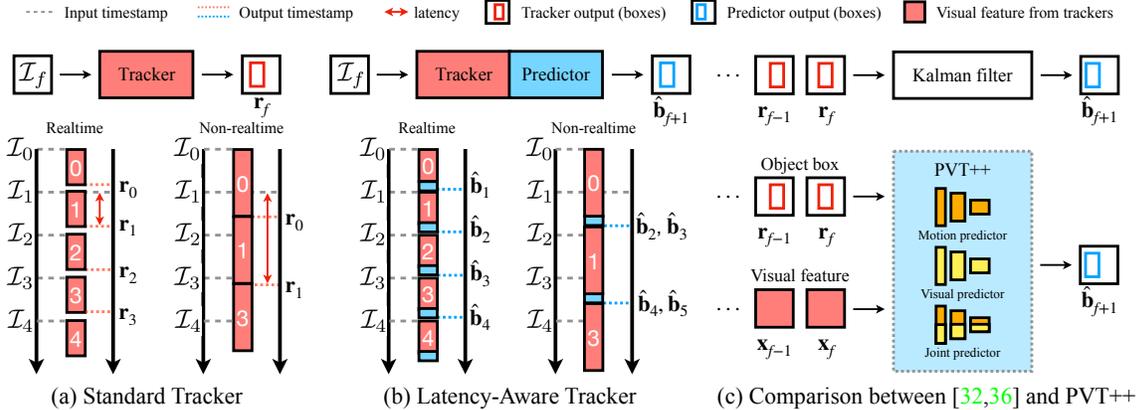
Figure 2. (a) Standard tracker suffers from onboard latency (height of the red boxes). Hence, its result lags behind the world, *i.e.*, $\mathbf{r}_f$ is always obtained after $\mathcal{I}_f$ on the timestamp. (b) Latency-aware trackers introduce predictors to compensate for the latency, which predict the word state, $\hat{\mathbf{b}}_{f+1}$, when finishing the processed frame. (c) Compared with prior KF-based solutions [32, 36], our end-to-end framework for latency-aware tracking PVT++ leverages both motion and visual feature for prediction.

tive trackers. While a standard tracker yields the objects' location in the input frame (*i.e.*, when it *starts* processing the input frame, as in Fig. 2(a)), a predictive tracker predicts where the objects could be when it *finishes* processing the input frame, as illustrated in Fig 2(b).

Existing solutions [32, 36] directly append a Kalman filter (KF) [30] after trackers to estimate the potential object's location based on its motion model (see Fig 2(c)). However, the rich and readily available visual knowledge from trackers is primarily overlooked, including the object's appearance and the surrounding environments, which can be naturally exploited to predict the objects' future paths [51].

To this end, we present a simple framework PVT++ for end-to-end predictive visual tracking. Composed of a tracker and a predictor, PVT++ is able to convert most off-the-shelf trackers into effective predictive trackers. Specifically, to avoid extra latency brought by the predictor, we first design a lightweight network architecture, consisting of a feature encoder, temporal interaction module, and predictive decoder, that leverage both historical motion information and visual cues. By virtue of joint optimization, such a small-capacity network can directly learn from the visual representation provided by most pre-trained trackers for an efficient and accurate motion prediction, as in Fig. 2(c). However, learning this framework is non-trivial due to the training-evaluation domain gap in terms of motion scales. To solve this, we develop a relative motion factor as training objective, so that our framework is independent of the motion scales in training data and can generalize well to the challenging aerial tracking scenes. The integration of lightweight structure and training strategy yields an effective, efficient, and versatile solution.

Beyond methodology, we found that the existing latency-aware evaluation benchmark (LAE) [32] is unable to provide an effective latency-aware comparison for real-time

trackers, since it evaluates the result for each frame as soon as it is given. In this case, the latency for any real-time trackers is one frame. Hence, we present an extended latency-aware evaluation benchmark (e-LAE) for *any-speed* trackers. Evaluated with various latency thresholds, real-time trackers with different speeds can be distinguished.

Empirically, we provide a more general, comprehensive, and practical aerial tracking evaluation for state-of-the-art trackers using our new e-LAE. Converting them into predictive trackers, PVT++ achieves up to **60%** improvement under the *online* setting. As shown in Fig. 1, powered by PVT++, the predictive trackers can achieve comparable or better results than the *offline* setting. Extensive experiments on multiple tracking models [33, 57, 22] and datasets [47, 16, 37] show that PVT++ works generally for latency-aware tracking, which, to the best of our knowledge, is also the first end-to-end framework for *online* visual tracking.

2. Related Work

2.1. Visual Tracking and its Aerial Applications

Visual trackers basically fall into two paradigms, respectively based on discriminative correlation filters [4, 27, 15, 13] and Siamese networks [2, 34, 64, 33, 23, 58]. Compared with general scenarios, aerial tracking is more challenging due to large motions and limited onboard computation resources. Hence, for efficiency, early approaches focus on correlation filters [38, 29, 39, 35]. Later, the development of onboard computation platforms facilitates more robust and applicable Siamese network-based approaches [20, 7, 6, 8].

Most of them are designed under *offline* settings, ignoring the online latency onboard UAVs, which can lead to severe accuracy degradation. We aim to solve the more practical and challenging latency-aware perception problem, which goes beyond tracking and is often neglected.
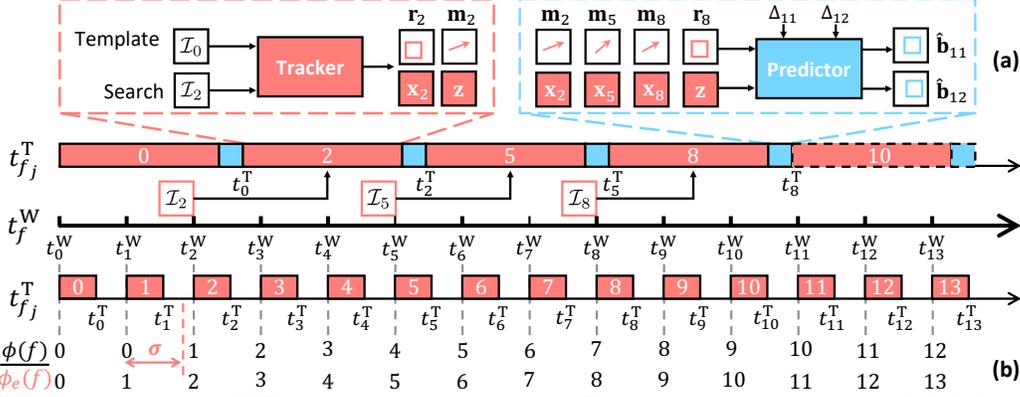
Figure 3. (a) Framework overview of PVT++ for a non-real-time tracker. The tracker has processed frame $0, 2, 5, 8$ and obtained corresponding motions $\mathbf{m}$ and visual features $\mathbf{x}, \mathbf{z}$. The predictor needs to predict future box $\hat{\mathbf{b}}_{11}, \hat{\mathbf{b}}_{12}$ based on tracker result $\mathbf{r}_8$. (b) Comparison between LAE ($\phi(f)$) [32] and our e-LAE ($\phi_e(f)$). For real-time trackers, the mismatch between output and input frames will always be one in LAE ($\phi(f) - f \equiv 1$) regardless of the trackers' various latency. Differently, e-LAE introduces permitted latency thresholds $\sigma \in [0, 1)$, which effectively distinguishes the latency difference of distinct models.

2.2. Latency-Aware Perception

Latency of perception systems is first studied in [36], which introduces a baseline based on the Kalman-filter [30] to compensate for the online latency of object detectors. Inspired by this, [59] converts a real-time detector into a latency-aware one. More recent work [54] also leverage context scenario properties to select best configurations. While most previous work aims at object detection, [32] is more related to us in the field of visual tracking, which introduces dual KFs [30] similar to [36]. Overall, most existing works didn't address latency-aware tracking through a deeply coupled framework. In this work, we target aerial tracking and present an end-to-end structure.

2.3. Tracking with/by Prediction

Integrating prediction into trackers has been widely adopted for robust perception [41, 51, 46, 19, 42, 44].

Some previous works [42, 44] introduce a predictor to correct the tracker results under object occlusion, which potentially share similar design with PVT++. However, they are still developed for *offline* tracking, *i.e.*, the predictor works only for current frame. Yet predicting the future state for aerial tracking is non-trival due to the training-evaluation gap. PVT++ aims to address this gap via a relative motion factor, thus working beyond [42, 44].

Others [41, 51, 46, 19] focus on trajectory prediction after tracking. While PVT++ predicts both the trajectory and the object scale. Moreover, as a result of fast camera motion and viewpoint change, UAV tracking scenes are much more complex than the ordered environments in autonomous driving [41, 46]. PVT++ discovers the valuable visual knowledge in pre-trained trackers, yielding its capability for such challenging condition with small capacity.

2.4. Visual Tracking Benchmarks

Various benchmarks are built for large-scale tracking evaluation [18, 48, 28, 17, 43, 47, 31, 62] with different challenges such as first-person perspective [17], aerial scenes [47], illumination conditions [31, 62], and thermal infrared inputs [43]. Since they all adopt *offline* evaluation, the influence of the trackers' latency is ignored. A recent benchmark targets online evaluation [32], but it falls short in real-time trackers and we aim to improve it in this work.

3. Preliminary

We first introduce the latency-aware tracking task here. The input is an image sequence broadcasting with a certain framerate κ, denoted as $(\mathcal{I}_f, t_f^{\mathrm{W}}), f \in \{0, 1, 2, \cdots\}$, where $t_f^{\mathrm{W}} = \frac{f}{\kappa}$ is the world timestamp and f is the frame index. Provided with the ground truth box $\mathbf{b}_0 = [x_0, y_0, w_0, h_0]$ at initial 0-th frame, the tracker estimates the boxes in the following frames. Detailed notation table see Appendix C.
Inference. During inference, the tracker finds the *latest* frame to process when finishing the previous one. Due to the latency, for the j-th frame that the tracker processes, its index j may differ from its frame index f_j in the image sequence. The frame to be processed (frame f_j) is determined by the tracker timestamp $t_{f_{j-1}}^{\mathrm{T}}$ when the model finishes frame f_{j-1} as follows:

$$f_j = \begin{cases} 0 & , \quad j = 0 \\ \arg\max_f t_f^{\mathrm{W}} \le t_{f_{j-1}}^{\mathrm{T}} & , \quad \text{others} \end{cases} . \quad (1)$$

With the frame index f_j, the tracker processes frame $\mathcal{I}_{f_j}$ to obtain the corresponding box $\mathbf{r}_{f_j} = [x_{f_j}, y_{f_j}, w_{f_j}, h_{f_j}]$, forming the raw result of the tracker on the frame $(\mathbf{r}_{f_j}, t_{f_j}^{\mathrm{T}})$. Since tracker may be non-real-time, input frame ids $f_j, j \in \{0, 1, 2, \cdots\}$ may not be consecutive numbers. For example, in Fig. 3 (a), considering a non-real-time tracker, the processed frames are $f_j = 0, 2, 5, 8, \cdots$.
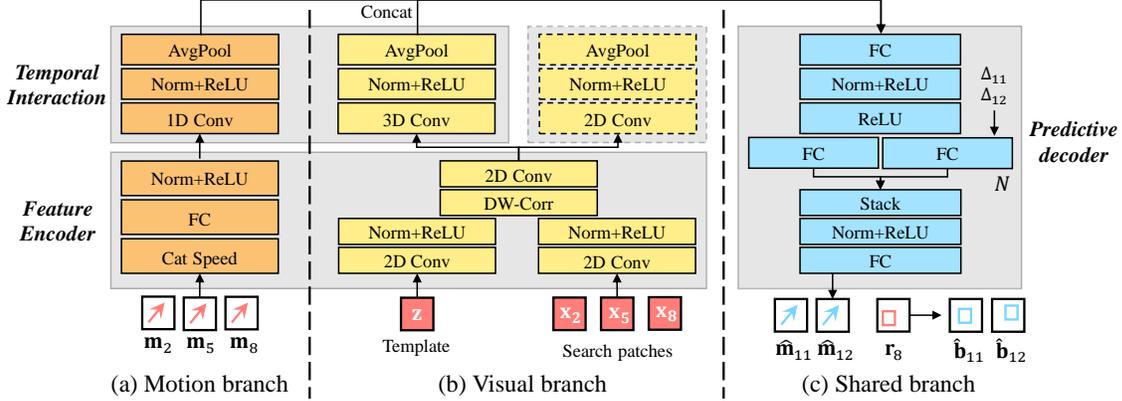
Figure 4. Detailed model structure of the predictor modules in PVT++. The models shares similar architecture, *i.e.*, feature encoder, temporal interaction, and predictive decoder. We present the motion branch, visual branch, and share decoding branch in (a), (b), and (c), respectively. Note that the dashed blocks denote auxiliary branch, which only exists in training. The input and output are in correspondence to the case in Fig. 3 (a).

Evaluation. Latency-aware evaluation (LAE) [32] compares the ground-truth $\mathbf{b}_f$ in frame $\mathcal{I}_f$ with the *latest* result $\hat{\mathbf{b}}_f$ from the tracker at t_f^{W} for evaluation. For standard trackers, the latest result $\hat{\mathbf{b}}_f$ to be compared with the ground-truth is obtained as $\hat{\mathbf{b}}_f = \mathbf{r}_{\phi(f)}$, where $\phi(f)$ is defined as:

$$\phi(f) = \begin{cases} 0 & , \quad t_f^{\mathrm{W}} < t_{f_0}^{\mathrm{T}} \\ \arg\max_{f_j} t_{f_j}^{\mathrm{T}} \le t_f^{\mathrm{W}} & , \quad \text{others} \end{cases} . \quad (2)$$

For instance, in Fig. 3 (b), LAE compares the ground truth $\mathbf{b}_3$ with the raw tracker result $\mathbf{r}_2$.

4. Extended Latency-Aware Benchmark

Existing latency-aware evaluation [36, 32] adopt Eq. (2) to match the raw output $(\mathbf{r}_{f_j}, t_{f_j}^{\mathrm{T}})$ to every input frame f. However, such a policy fails to reflect the latency difference among real-time trackers. As shown in Fig. 3, since the real-time methods is faster than frame rate, every frame will be processed, *i.e.*, $[f_0, f_1, f_2, \cdots] = [0, 1, 2, \cdots]$. In this case, the *latest* results will always be from the previous one frame, *i.e.*, using Eq. (2), $\phi(f) \equiv f - 1$. Differently, we extend Eq. (2) to:

$$\phi(f)_{\mathrm{e}} = \begin{cases} 0 & , \quad t_f^{\mathrm{W}} < t_{f_0}^{\mathrm{T}} \\ \arg\max_{f_j} t_{f_j}^{\mathrm{T}} \le t_f^{\mathrm{W}} + \sigma & , \quad \text{others} \end{cases} , \quad (3)$$

where $\sigma \in [0, 1)$ is the variable permitted latency. Under e-LAE, $\phi(f)_{\mathrm{e}}$ can be $f - 1$ or f for real-time trackers depending on σ. For instance, $\phi(f)_{\mathrm{e}}$ would turn from $f - 1$ to f at larger σ for slower real-time trackers. This extension distinguishes different real-time trackers (see Section 6.2).

5. Predictive Visual Tracking

Because of the unavoidable latency introduced by the processing time, there is always a mismatch between $\phi(f)$ (or $\phi(f)_{\mathrm{e}}$) and f (when σ is small), where $\phi(f)$ is always

smaller than f, *i.e.*, $\phi(f) < f, f > 0$. To compensate for the mismatch, we resort to predictive trackers that predicts possible location of the object in frame f. For the evaluation of f-th frame, prior attempts [36, 32] adopt traditional KF [30] to predict the result based on the raw tracking result $\mathbf{r}_{\phi(f)}$ in $\mathcal{I}_\phi(f)$ [36], *i.e.*, $\hat{\mathbf{b}}_f = \mathrm{KF}(\mathbf{r}_{\phi(f)})$. Since previous work [36, 32] are not learnable, neither existing large-scale datasets nor the visual feature are leveraged. Differently, our predictive visual tracking framework PVT++ aims for an end-to-end predictive tracker, which takes both the historical motion and visual features for a more robust and accurate prediction. Note that we use $\hat{\cdot}$ to represent the prediction (results for evaluation) and others are from the tracker output or ground-truth in the following subsections.

5.1. General Framework

As in Fig. 3 (a), PVT++ consists of a tracker $\mathcal{T}$ and a predictor $\mathcal{P}$. For the f-th frame at world time t_f^{W}, the latest result from the tracker is $\mathbf{r}_{\phi(f)}$ obtained from frame $\mathcal{I}_{\phi(f)}$, *i.e.*, $\mathbf{r}_{\phi(f)} = \mathcal{T}(\mathbf{x}_{\phi(f)}, \mathbf{z})$, where $\mathbf{x}_{\phi(f)}$ is the search feature from $\mathcal{I}_{\phi(f)}$ and $\mathbf{z}$ is the template feature.

After this, the predictor $\mathcal{P}$ takes input from the information generated during tracking of the k past frames (including $\mathcal{I}_{\phi(f)}$), denoted as $\mathrm{Input}_{\phi(f)}$, and predict the position offset normalized by object's scale, *i.e.*, motion $\hat{\mathbf{m}}_f = [\frac{\Delta_{\hat{x}}(f)}{w_{\phi(f)}}, \frac{\Delta_{\hat{y}}(f)}{h_{\phi(f)}}, \mathrm{Log}(\frac{\hat{w}_f}{w_{\phi(f)}}), \mathrm{Log}(\frac{\hat{h}_f}{h_{\phi(f)}})]$, where $\Delta_{\hat{x}}(f)$ and $\Delta_{\hat{y}}(f)$ denote the predicted box center distance between the f-th and $\phi(f)$-th frame. $w_{\phi(f)}$ and $h_{\phi(f)}$ are the tracker's output box scale in frame $\phi(f)$ and $\hat{w}_f, \hat{h}_f$ are the predicted scale in f-th frame. With the raw output $\mathbf{r}_{\phi(f)}$ at $\phi(f)$ and the motion $\hat{\mathbf{m}}_f$ from $\mathcal{I}_{\phi(f)}$ to the f-th frame, the predicted box $\hat{\mathbf{b}}_f$ can be easily calculated.

Relative Motion Factor: Due to the large domain gap between the training [52] and evaluation [37] in terms of the absolute motion scale, we find directly using the absolute
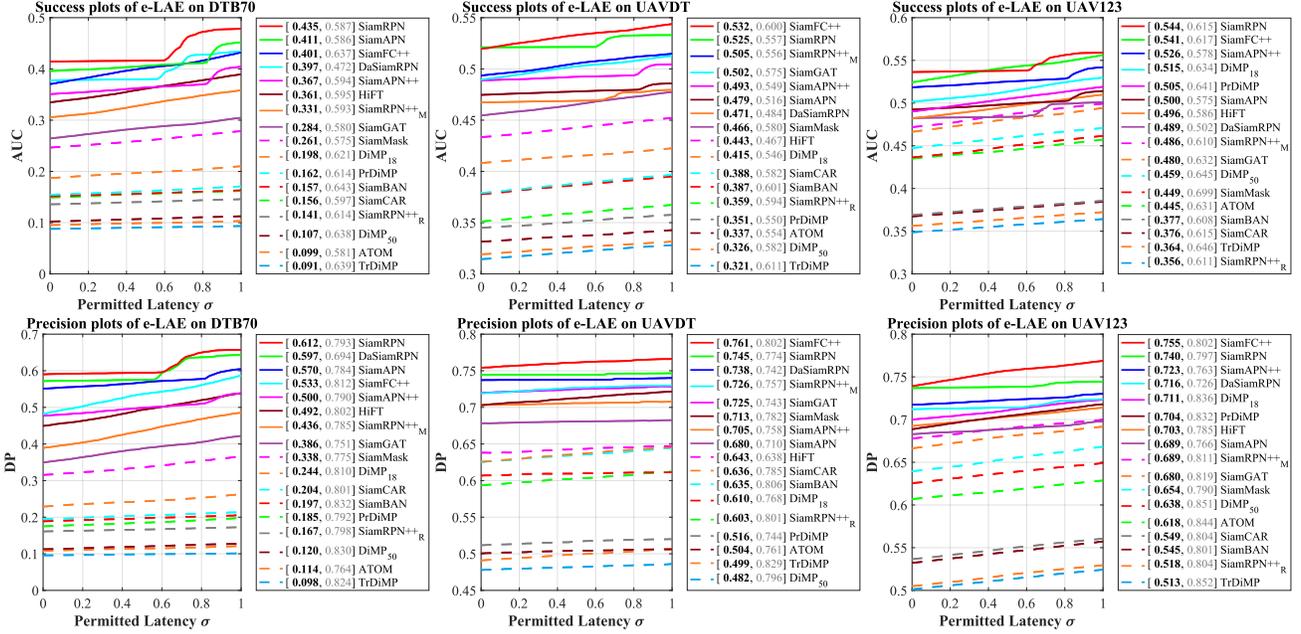
Figure 5. The performance of the SOTA trackers in authoritative UAV tracking benchmarks under our e-LAE benchmark. We report [online mAUC and mDP, offline AUC and DP] in the legend. All trackers struggle to overcome onboard latency in online tracking.

motion value $\hat{\mathbf{m}}_f$ as the objective can result in poor performance (see Table 4). Therefore, we define the output of predictor $\mathcal{P}$ to be the relative motion factor based on the average moving speed $\mathbf{p}_{f_j}$ from the past k frames, which we find is easier to generalize after training:

$$\hat{\mathbf{m}}_f = \mathcal{P}\big(\text{Input}_{\phi(f)}, \Delta_f\big) \odot \mathbf{p}_{f_j}, \quad \mathbf{p}_{f_j} = \frac{1}{k} \sum_{i=j-k+1}^{j} \frac{\mathbf{m}_{f_i}}{\Delta_{f_i}}, \ (4)$$

where $\Delta_f = f - \phi(f)$ denotes the frame interval between current and target frame, and f_j is the latest processed frame, e.g., $\phi(f)$. $\Delta_{f_i} = f_i - f_{i-1}$ denotes the frame interval between $(i-1)$ and i-th processed frame. $\odot$ indicates element-wise multiplication. $\mathbf{m}_{f_i}$ is the normalized input motion defined as $\mathbf{m}_{f_i} = [\frac{\Delta_x(f_i)}{w_{f_{i-1}}}, \frac{\Delta_y(f_i)}{h_{f_{i-1}}}, \text{Log}(\frac{w_{f_i}}{w_{f_{i-1}}}), \text{Log}(\frac{h_{f_i}}{h_{f_{i-1}}})]$, where $\Delta_x(f_i) = x_{f_i} - x_{f_{i-1}}$ and $\Delta_y(f_i) = y_{f_i} - y_{f_{i-1}}$ are the distance from tracker results $\mathbf{r}_{f_i}$ and $\mathbf{r}_{f_{i-1}}$. Such design has made PVT++ agnostic to the specific motion of a dataset, which is crucial for its generalization capability.

We next present the predictor of PVT++ step by step as motion-based $\mathcal{P}_\text{M}$, visual-appearance-based $\mathcal{P}_\text{V}$ and multi-modal-based $\mathcal{P}_\text{MV}$. All the predictors share the same training objective (Eq. (4)) and a similar structure, consisting of feature encoding, temporal interaction, and predictive decoding as in Fig. 4. In practice, a predictor may need to predict N results, depending on the tracker's latency.

5.2. Motion-based Predictor

The motion-based predictor $\mathcal{P}_\text{M}$ only relies on the past motion, i.e., $\text{Input}_{\phi(f)} = \mathbf{m}_{f_{j-k+1}:f_j}$,

$$\hat{\mathbf{m}}_{f,\text{M}} = \mathcal{P}_\text{M}\big(\mathbf{m}_{f_{j-k+1}:f_j}, \Delta_f\big) \odot \mathbf{p}_{f_j}, \quad (5)$$

where $\mathbf{m}_{f_{j-k+1}:f_j} = [\mathbf{m}_{f_{j-k+1}}, \cdots, \mathbf{m}_{f_j}] \in \mathbb{R}^{k \times 4}$.

The detailed model structure of the motion predictor $\mathcal{P}_\text{M}$ is presented in Fig. 4(a). For pre-processing, the motion data $\mathbf{m}_{f_{j-k+1}}, \cdots, \mathbf{m}_{f_j}$ are first concatenated. Then we apply a fully connected (FC) layer with non-linearity for feature encoding and a 1D convolution followed by activation and global average pooling to obtain the temporally interacted motion feature. In the predictive decoding head, a share FC layer with non-linearity is used for feature mapping. N independent FCs map the feature to N future latent spaces. Finally, the latency features are stacked and transformed to 4 dimension output using a shared FC.

For training, we adopt $\mathcal{L}_1$ loss between prediction and ground-truth $\mathcal{L}_\text{M} = \mathcal{L}_1(\hat{\mathbf{m}}_{f,M}, \mathbf{m}_f)$.

5.3. Visual Appearance-based Predictor

For efficiency, our visual predictor $\mathcal{P}_\text{V}$ takes search and template features directly from the tracker backbone as input. Besides, we also find the strong representation in the pre-trained tracker models can boost the small-capacity predictor network. Specifically, template feature $\mathbf{z} \in \mathbb{R}^{1 \times C_\text{V} \times a \times a}$ is extracted from the given object template patch in the initial frame and search feature $\mathbf{x}_{f_j} \in \mathbb{R}^{1 \times C_\text{V} \times s \times s}$ is obtained from the f_j-th frame patch

Table 1. The effect of PVT++ on the four SOTA trackers with different inference speeds and backbones. Our models work generally for different tracker structures and can achieve up to **60**% performance gain. The best scores are marked out in gray for clear reference. We present some qualitative visualization in Appendix D and the supplementary video.

Tracker	Dataset PVT++	DTB70 AUC@La0$_{\triangle\%}$	DP@La0$_{\triangle\%}$	UAVDT AUC@La0$_{\triangle\%}$	DP@La0$_{\triangle\%}$	UAV20L AUC@La0$_{\triangle\%}$	DP@La0$_{\triangle\%}$	UAV123 AUC@La0$_{\triangle\%}$	DP@La0$_{\triangle\%}$
SiamRPN++$_M$ (~22FPS) (MobileNet [53])	N/A	$0.305_{+0.00}$	$0.387_{+0.00}$	$0.494_{+0.00}$	$0.719_{+0.00}$	$0.448_{+0.00}$	$0.619_{+0.00}$	$0.472_{+0.00}$	$0.678_{+0.00}$
	$\mathcal{P}_M$	$0.385_{+26.2}$	$0.523_{+35.1}$	$0.529_{+7.10}$	$0.745_{+3.60}$	$0.481_{+7.40}$	$0.647_{+4.50}$	$0.537_{+13.8}$	$0.737_{+8.70}$
	$\mathcal{P}_V$	$0.352_{+15.4}$	$0.472_{+22.0}$	$0.564_{+14.2}$	$0.799_{+11.1}$	$0.488_{+8.90}$	$0.675_{+9.00}$	$0.504_{+6.80}$	$0.703_{+3.70}$
	$\mathcal{P}_{MV}$	$0.399_{+30.8}$	$0.536_{+38.5}$	$0.576_{+16.6}$	$0.807_{+12.2}$	$0.508_{+13.4}$	$0.697_{+12.6}$	$0.537_{+13.8}$	$0.741_{+9.30}$
SiamRPN++$_R$ (~6FPS) (ResNet50 [26])	N/A	$0.136_{+0.00}$	$0.159_{+0.00}$	$0.351_{+0.00}$	$0.594_{+0.00}$	$0.310_{+0.00}$	$0.434_{+0.00}$	$0.349_{+0.00}$	$0.505_{+0.00}$
	$\mathcal{P}_M$	$0.199_{+46.3}$	$0.258_{+62.3}$	$0.449_{+27.9}$	$0.684_{+15.2}$	$0.404_{+30.3}$	$0.560_{+29.0}$	$0.442_{+26.6}$	$0.627_{+24.2}$
	$\mathcal{P}_V$	$0.179_{+31.6}$	$0.225_{+41.5}$	$0.403_{+14.8}$	$0.665_{+12.0}$	$0.398_{+28.4}$	$0.548_{+26.3}$	$0.398_{+14.0}$	$0.559_{+10.7}$
	$\mathcal{P}_{MV}$	$0.205_{+50.7}$	$0.256_{+61.0}$	$0.488_{+39.0}$	$0.726_{+22.2}$	$0.416_{+34.2}$	$0.568_{+30.9}$	$0.442_{+26.6}$	$0.619_{+22.6}$
SiamMask (~14FPS) (ResNet50 [26])	N/A	$0.247_{+0.00}$	$0.313_{+0.00}$	$0.455_{+0.00}$	$0.703_{+0.00}$	$0.405_{+0.00}$	$0.571_{+0.00}$	$0.436_{+0.00}$	$0.639_{+0.00}$
	$\mathcal{P}_M$	$0.370_{+49.8}$	$0.508_{+62.3}$	$0.531_{+16.7}$	$0.760_{+8.10}$	$0.449_{+10.9}$	$0.607_{+6.30}$	$0.532_{+22.0}$	$0.743_{+16.9}$
	$\mathcal{P}_V$	$0.292_{+18.2}$	$0.405_{+29.4}$	$0.532_{+16.9}$	$0.777_{+10.5}$	$0.430_{+6.20}$	$0.601_{+5.30}$	$0.503_{+15.4}$	$0.705_{+10.3}$
	$\mathcal{P}_{MV}$	$0.342_{+29.5}$	$0.463_{+47.9}$	$0.566_{+24.4}$	$0.797_{+13.4}$	$0.469_{+15.8}$	$0.644_{+12.8}$	$0.536_{+22.9}$	$0.749_{+17.2}$
SiamGAT (~16FPS) (GoogleNet [55])	N/A	$0.264_{+0.00}$	$0.347_{+0.00}$	$0.489_{+0.00}$	$0.720_{+0.00}$	$0.475_{+0.00}$	$0.663_{+0.00}$	$0.466_{+0.00}$	$0.666_{+0.00}$
	$\mathcal{P}_M$	$0.396_{+50.0}$	$0.520_{+49.9}$	$0.549_{+12.3}$	$0.777_{+7.90}$	$0.519_{+9.26}$	$0.705_{+6.33}$	$0.530_{+13.7}$	$0.739_{+11.0}$
	$\mathcal{P}_V$	$0.392_{+48.5}$	$0.536_{+54.4}$	$0.575_{+17.6}$	$0.801_{+11.3}$	$0.503_{+5.90}$	$0.686_{+3.50}$	$0.514_{+10.3}$	$0.713_{+7.10}$
	$\mathcal{P}_{MV}$	$0.415_{+57.2}$	$0.561_{+61.7}$	$0.583_{+19.2}$	$0.796_{+10.6}$	$0.519_{+9.3}$	$0.708_{+6.80}$	$0.531_{+13.9}$	$0.733_{+10.1}$

cropped around $(x_{f_{j-1}}, y_{f_{j-1}})$. Given k past search features $\mathbf{x}_{f_{j-k+1}:f_j} \in \mathbb{R}^{k \times C_V \times s \times s}$ and $\mathbf{z}$, we have:

$$\hat{\mathbf{m}}_{f,V} = \mathcal{P}_V\big(\mathbf{x}_{f_{j-k+1}:f_j}, \mathbf{z}, \Delta_f\big) \odot \mathbf{p}_{f_j} . \tag{6}$$

The detailed model structure of $\mathcal{P}_V$ is shown in Fig. 4(b). Inspired by Siamese trackers [33], the feature encoding stage adopts 1×1 convolution before depth-wise correlation (DW-Corr) to produce the similarity map $\mathbf{x}^e_{f_{j-k+1}:f_j} \in \mathbb{R}^{k \times C_V \times s' \times s'}$. For temporal interaction, we apply 3D convolution and global average pooling.

We find directly training $\mathcal{P}_V$ meets convergence difficulty (See Section 6.3). We hypothesize this is because the intermediate similarity map $\mathbf{x}^e_{f_{j-k+1}:f_j}$ fails to provide explicit motion information. To solve this, we introduce an auxiliary branch $\mathcal{A}$, which takes $\mathbf{x}^e_{f_{j-k+1}:f_j}$ as input to obtain the corresponding motion $\mathbf{m}^e_{f_{j-k+1}:f_j}$,

$$\mathbf{m}^e_{f_{j-k+1}:f_j} = \mathcal{A}(\mathbf{x}^e_{f_{j-k+1}:f_j}). \tag{7}$$

During training, we supervise both the auxiliary branch and the predictive decoder, i.e., $\mathcal{L}_V = \mathcal{L}_1(\hat{\mathbf{m}}_{f,V}, \mathbf{m}_f) + \mathcal{L}_1(\mathbf{m}^e_{f_{j-k+1}:f_j}, \mathbf{m}_{f_{j-k+1}:f_j})$.

5.4. Multi-Modality-based Predictor

The final predictor $\mathcal{P}_{MV}$ is constructed as a combination of motion $\mathcal{P}_M$ and visual predictors $\mathcal{P}_V$, which takes both visual feature $\mathbf{x}_{f_{j-k+1}:f_j}, \mathbf{z}$ and motion information $\mathbf{m}_{f_{j-k+1}:f_j}$ as input, i.e.,

$$\hat{\mathbf{m}}_{f,MV} = \mathcal{P}_{MV}\big(\mathbf{m}_{f_{j-k+1}:f_j}, \mathbf{x}_{f_{j-k+1}:f_j}, \mathbf{z}, \Delta_f\big) \odot \mathbf{p}_{f_j} . \tag{8}$$

As shown in Fig. 4, the encoding and temporal interaction parts of $\mathcal{P}_M$ and $\mathcal{P}_V$ run in parallel to form the first two stages of $\mathcal{P}_{MV}$. We concatenate the encoded feature vectors to obtain the multi-modal feature. The predictive decoder follows the same structure to obtain future motions $\hat{\mathbf{m}}_{f,MV}$. We also tried different fusion strategy in Appendix H.

For training, we add the two additional predictive decoders respectively after motion and visual predictors to help them predict $\hat{\mathbf{m}}_{f,M}$ and $\hat{\mathbf{m}}_{f,V}$, which yields the loss $\mathcal{L}_{MV} = \alpha_M \mathcal{L}_M + \alpha_V \mathcal{L}_V + \mathcal{L}_1(\hat{\mathbf{M}}_{f,MV}, \mathbf{M}_f)$. During inference, we only use the joint predictive decoder.

Remark 1: The predictors $\mathcal{P}_M, \mathcal{P}_V$ and $\mathcal{P}_{MV}$ can be jointly optimized with tracker $\mathcal{T}$, i.e., during training, both input motion and visual feature are from the tracker module with gradient, so that the two modules are deeply coupled.

6. Experiments

6.1. Implementation Details

Platform and Datasets. PVT++ is trained on VID [52], LaSOT [18], and GOT10k [28] using one Nvidia A10 GPU. The evaluation takes authoritative UAV tracking datasets, UAV123, UAV20L [47], DTB70 [37], and UAVDT [16] on typical UAV computing platform, Nvidia Jetson AGX Xavier, for realistic robotic performance. Since the online latency can fluctuate, we run three times and report the average performance. For simplicity, we only consider the tracker's processing latency during evaluation.

Metrics. Following [21], we use two basic metrics, the distance precision (DP) based on center location error (CLE) and area under curve (AUC) based on intersection over union. Under e-LAE, different permitted latency σ corresponds to different DP and AUC, i.e., DP@Laσ and AUC@Laσ. We use mDP and mAUC to indicate the area under cure for DP@Laσ and AUC@La$\sigma, \sigma \in [0:0.02:1]$.

Parameters. For e-LAE, all the evaluated trackers use their official parameters for fairness. To represent the most common case, the image frame rate is fixed to $\kappa = 30$ frames/s (FPS) in all the online evaluation. For the PVT++ models, we use $k = 3$ past frames. To determine N for different models, we pre-run the trackers 3 times and record the maximum number of skipped frames, so that when the latency

Table 2. Attribute-based analysis of PVT++ in UAVDT [16]. We found different modality has their specific advantage. Together, the joint model can utilize both and is the most robust under complex UAV tracking challenges. Gray denotes best results.

Tracker	Metric Att.	AUC@La0							
		BC	CR	OR	SO	IV	OB	SV	LO
SiamRPN++$_M$	N/A	0.448	0.450	0.438	0.494	0.539	0.525	0.490	0.422
	$\mathcal{P}_M$	0.461	0.495	0.481	0.549	0.578	0.542	0.505	0.521
	$\mathcal{P}_V$	0.504	0.520	0.538	0.525	0.588	0.568	0.584	0.436
	$\mathcal{P}_{MV}$	0.505	0.535	0.549	0.545	0.599	0.589	0.586	0.511
SiamMask	N/A	0.404	0.425	0.404	0.468	0.475	0.471	0.438	0.389
	$\mathcal{P}_M$	0.465	0.503	0.491	0.536	0.558	0.542	0.526	0.421
	$\mathcal{P}_V$	0.488	0.498	0.504	0.495	0.563	0.527	0.541	0.494
	$\mathcal{P}_{MV}$	0.520	0.522	0.541	0.540	0.596	0.560	0.566	0.520

Tracker	Metric Att.	DP@La0							
		BC	CR	OR	SO	IV	OB	SV	LO
SiamRPN++$_M$	N/A	0.659	0.643	0.638	0.779	0.777	0.772	0.680	0.569
	$\mathcal{P}_M$	0.666	0.684	0.681	0.815	0.811	0.778	0.691	0.717
	$\mathcal{P}_V$	0.733	0.720	0.753	0.793	0.835	0.822	0.796	0.585
	$\mathcal{P}_{MV}$	0.727	0.732	0.764	0.814	0.848	0.846	0.794	0.694
SiamMask	N/A	0.628	0.620	0.612	0.803	0.743	0.756	0.650	0.571
	$\mathcal{P}_M$	0.672	0.702	0.709	0.818	0.797	0.802	0.729	0.590
	$\mathcal{P}_V$	0.718	0.696	0.723	0.787	0.817	0.801	0.763	0.696
	$\mathcal{P}_{MV}$	0.731	0.712	0.752	0.819	0.829	0.813	0.783	0.711

of one specific frame fluctuates, PVT++ can always cover the skipped frame and make sufficient predictions. Detailed training configurations can be found in Appendix B.

6.2. Extended Latency-Aware Evaluation

We evaluate a total of 17 SOTA trackers[2] under e-LAE: SiamRPN [34], SiamRPN++$_M$ [33], SiamRPN++$_R$ [33], SiamMask [57], SiameseFC++ [58], DaSiamRPN [64], SiamAPN [20], SiamAPN++ [7], HiFT [6], SiamGAT [22], SiamBAN [10], SiamCAR [23], ATOM [12], DiMP$_{50}$ [3], DiMP$_{18}$ [3], PrDiMP [14], and TrDiMP [56].

As in Fig. 5, we draw curve plots to reflect their performance in AUC and DP metrics under different permitted latency σ. We report the [online mAUC and mDP, offline AUC and DP] in the legend. Some *offline* highly accurate trackers like SiamRPN++$_R$ [33], SiamCAR [23], SiamBAN [10], and ATOM [12] can degrade by up to **70**% in our online evaluation setting.

Remark 2: e-LAE can better assess the real-time trackers. In DTB70, SiamAPN++ and HiFT are two real-time trackers with HiFT more accurate in success. While since SiamAPN++ is faster, its e-LAE performance will be better.

6.3. Empirical Analysis of PVT++

Overall Effect. To evaluate PVT++, we construct predictive trackers with four well-known methods, *i.e.,* SiamRPN++$_M$ [33], SiamRPN++$_R$ [33], SiamMask [57], and SiamGAT [22]. As in Table 1, with PVT++, their online performance can be significantly boosted by up to **60**%, which sometimes is better than their *offline* performance. PVT++ also works for recent transformer-based trackers [60, 11], the results can be found in Appendix F.

---

[2]Subscripts denote the backbone used, *i.e.,* MobileNet [53], and ResNet 18 or 50 [26].

Table 3. Dimension analysis of different modules in PVT++ on DTB70 [37] and UAVDT [16]. Enc.$_M$ and Enc.$_V$ represent the motion and visual encoders, respectively. Dec.$_{MV}$ denotes the joint decoder. * indicates our default setting. We find the channel dimension of PVT++ can be small, so that it introduces very few extra latency on robotics platforms.

Dim. of Modules			DTB70		UAVDT	
Enc.$_M$	Enc.$_V$	Dec.$_{MV}$	mAUC	mDP	mAUC	mDP
Base Tracker			0.305	0.387	0.494	0.719
16			0.357	0.479	0.565	0.797
32			0.359	0.483	0.575	0.81
64*	64*	32*	**0.399**	**0.536**	**0.576**	**0.807**
128			0.373	0.504	0.571	0.803
	16		0.362	0.487	0.545	0.772
	32		0.363	0.493	0.554	0.784
64*	64*	32*	**0.399**	**0.536**	**0.576**	**0.807**
	128		0.364	0.486	0.558	0.788
		16	0.369	0.496	0.572	0.804
64*	64*	32*	**0.399**	**0.536**	**0.576**	**0.807**
		64	0.373	0.503	0.567	0.807
		128	0.362	0.485	0.561	0.791

Table 4. Ablation studies on DTB70 [37]. Official version of PVT++ is marked out in Blackbody. The subscripts * means predicting raw value instead of motion factor, † denotes training without auxiliary supervision, and ‡ indicates training with tracker fixed. Red denotes improvement and blue represents dropping.

Ablate Module		Motion Factor*		Auxiliary Supervision† and Joint Training‡					
Method	Base	$\mathcal{P}_M$	$\mathcal{P}_M^*$	$\mathcal{P}_V$	$\mathcal{P}_V^\dagger$	$\mathcal{P}_V^\ddagger$	$\mathcal{P}_{MV}$	$\mathcal{P}_{MV}^\dagger$	$\mathcal{P}_{MV}^\ddagger$
AUC@La0	0.305	**0.385**	0.300	**0.352**	0.311	0.278	**0.399**	0.323	0.294
Delta%	0.00	26.2	-1.60	15.4	2.00	-8.90	30.8	5.90	-3.60
DP@La0	0.387	**0.523**	0.383	**0.472**	0.412	0.349	**0.536**	0.429	0.387
Delta%	0.00	35.1	-1.00	22.0	6.50	-9.80	38.5	-10.9	0.00

Remark 3: Real-time trackers [20, 7, 6, 34] perform generally better than non-real-time ones in online evaluation. While we observe that non-real-time trackers empowered by PVT++ can notably outperform real-time ones. *E.g.*, SiamRPN++$_M$ [33] with P_{MV} achieves an amazing **0.807** mDP in UAVDT, better than SiamFC++ [58] (0.761).

Attribute-based Analysis. For a comprehensive evaluation, we follow [16] and evaluate PVT++ on various challenge attributes[3]. From Table 2, We found that motion and vision have advantages in different attributes. $\mathcal{P}_V$ improves CR and OR, while $\mathcal{P}_M$ is good at SO and LO. The joint model $\mathcal{P}_{MV}$ makes use of both and is the most robust under various complex aerial tracking challenges. For the full attribute analysis, please see Appendix I.

Dimension Analysis. In addition to its promising performance, PVT++ can also work with very small capacity, which contributes to its lightweight architecture and high efficiency on low-powered UAVs. We analyse the modules of PVT++ with different feature channels in Table 3, where 64 channels for encoders (Enc.$_M$, Enc.$_V$) and 32 channels for the joint decoder (Dec.$_J$) work best. We present more efficiency and complexity comparisons with other motion predictors [63, 24, 40] in Appendix G

---

[3]Background cluter (BC), camera rotation (CR), object rotation (OR), small object (SO), illumination variation (IV), object blur (OB), scale variation (SV), and large occlusion (LO).

Table 5. Averaged results comparison on four datasets [47, 16, 37]. The motion based PVT++ can achieve on par or better results than prior KF-based solutions [36, 32]. Further introducing visual cues, PVT++ can acquire higher robustness. KF† and ‡ denotes learnable baselines [50, 25], which are still less robust than PVT++.

Type		Tradition Model			Learning-based				
Tracker	Pred.	N/A	KF[36]	PVT[32]	KF† [50]	KF‡ [25]	$\mathcal{P}_M$	$\mathcal{P}_V$	$\mathcal{P}_{MV}$
SiamRPN++$_M$	AUC@La0	0.43	0.462	0.473	0.466	0.481	0.483	0.477	**0.505**
	DP@La0	0.601	0.639	0.651	0.642	0.658	0.663	0.662	**0.695**
SiamMask	AUC@La0	0.386	0.441	0.465	0.458	0.468	0.471	0.439	**0.478**
	DP@La0	0.557	0.607	0.639	0.631	0.638	0.655	0.622	**0.663**
SiamRPN++$_R$	AUC@La0	0.287	0.361	0.374	0.376	0.386	0.374	0.345	**0.388**
	DP@La0	0.423	0.502	0.523	0.527	0.532	0.532	0.499	**0.542**

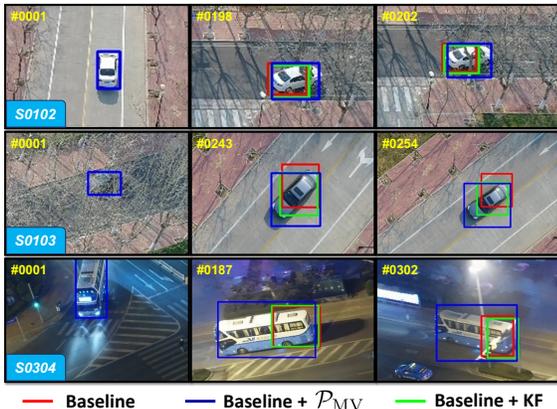— **Baseline**   — **Baseline + $\mathcal{P}_{MV}$**   — **Baseline + KF**

Figure 6. Prediction comparison from UAVDT [16]. We use red lines to demonstrate the original trackers, green for the KF [36] prediction, and blue for PVT++ prediction. Compared to KF, PVT++ is better at handling challenges like rotation, scale variation, and view point change. Best viewed in color.

Ablation Studies. We ablate the effect of motion factor prediction, auxiliary branch, and the joint training of PVT++ on DTB70 [37] with SiamRPN++$_M$ in Table 4. Compared with directly predicting the motion value ($\mathcal{P}_M^\dagger$), using *motion factor* as the prediction target ($\mathcal{P}_M$) can yield much better performance. Removing *auxiliary branch* $\mathcal{A}$ in $\mathcal{P}_V$ and $\mathcal{P}_{MV}$ to be $\mathcal{P}_V^\dagger$ and $\mathcal{P}_{MV}^\dagger$, we observe a significant performance drop due to the difficulty in convergence. *Joint training* the tracker and the predictor ($\mathcal{P}_V$ & $\mathcal{P}_{MV}$) perform much better than fixing the tracker ($\mathcal{P}_V^\ddagger$ and $\mathcal{P}_{MV}^\ddagger$). Training loss of the ablation studies are visualized in Appendix J.

6.4. Comparison with KF-based Solutions

Quantitative Results. Prior attempts to latency-aware perception [36, 32] have introduced model-based approach, *i.e.*, KF [30], as predictors. Based on traditional KF, we also designed stronger learnable baselines, KF† [50] and KF‡ [25], which adopt the same training as PVT++ models. Basically, KF† [50] learns the two noise matrix and KF‡ denotes joint training of KF† and trackers via backpropagation [25]. We compare these KF-based solutions with PVT++ in Table 5, where the same base tracker models are adopted. We present averaged mAUC and mDP in 4 datasets, DTB70 [37], UAVDT [16], UAV20L [47], and UAV123 [47]. Com-
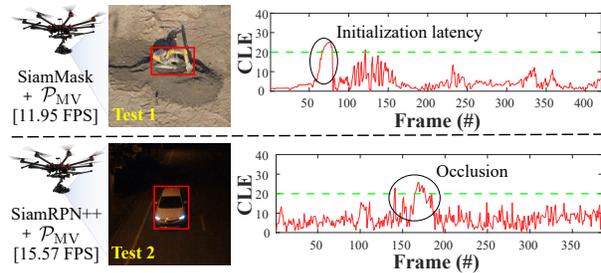

Figure 7. Real-world tests of PVT++. Thanks to PVT++, the non-real-time trackers work effectively under real-world challenges like object deformation in Test 1 and occlusion in Test 2.

pared with the KFs [36, 32], our learning framework holds the obvious advantage in complex UAV tracking scenes. We also observed that PVT++ is very efficient and introduces very little extra latency on the trackers. For specific results per dataset, please refer to Appendix E.

Qualitative Results. To better present the priority of PVT++, some representative scenes are displayed in Fig. 6. Given the same history trajectory, PVT++ holds its advantage against KF-based solution [36]. Especially, when UAV tracking challenges like in-plane rotation (sequence *S0103*) and aspect ration change (sequence *S0304*) appear, PVT++ is capable of fully utilizing the appearance change for robust prediction while simple motion-based KF easily fails.

Remark 4: Different from the KFs, which predict future results one by one, PVT++ outputs N results in a single forward pass, resulting in its high efficiency (especially for the motion predictor).

Apart from the robustness priority, PVT++ is also easier to be deployed. Once trained, no further tuning is needed for PVT++ to fit various scenes. Differently, the noise matrix of the KFs is dependent on the environment, which is hard to tune and may not generalize well to various complex UAV tracking scenes.

6.5. Real-World Tests

We further deploy SiamMask [57] (~11FPS) and SiamRPN++$_M$ [33] (~15FPS) with PVT++ on a UAV with Nvidia Jetson AGX Xavier as onboard processor. The onboard speed and center location error (CLE) results are shown in Fig. 7. Despite that the original tracker is not real-time, our PVT++ framework can convert it into a predictive tracker and achieve a good result (CLE < 20 pixels) in real-world tracking. More tests see Appendix M and the video.

7. Conclusion

In this work, we present a simple end-to-end framework for latency-aware visual tracking, PVT++, which largely compensates for onboard latency. PVT++ integrates a lightweight predictor module that discovers the visual representation from pre-trained trackers for precise and robust future state estimation. To bridge the training-evaluation

domain gap, we propose the relative motion factor, which yields a generalizable framework. In addition to PVT++, we introduce extended latency-aware evaluation benchmark (e-LAE), which assesses an *any-speed* tracker in the *online* setting. Extensive evaluations on robotics platform from the challenging aerial perspective show the effectiveness of our PVT++, which improves the *offline* tracker by up to **60**% in the *online* setting. Real-world tests are further conducted to exhibit the efficacy of PVT++ on physical robots.

References

[1] Wilbert G Aguilar, Leandro Álvarez, Santiago Grijalva, and Israel Rojas. Monocular Vision-Based Dynamic Moving Obstacles Detection and Avoidance. In *Proceedings of the International Conference on Intelligent Robotics and Applications (ICIRA)*, pages 386–398, 2019. 1

[2] Luca Bertinetto, Jack Valmadre, Joao F Henriques, Andrea Vedaldi, and Philip HS Torr. Fully-convolutional Siamese Networks for Object Tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 850–865, 2016. 2

[3] 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 (ICCV)*, pages 6181–6190, 2019. 7

[4] David S Bolme, J Ross Beveridge, Bruce A Draper, and Yui Man Lui. Visual Object Tracking Using Adaptive Correlation Filters. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2544–2550, 2010. 2

[5] Rogerio Bonatti, Cherie Ho, Wenshan Wang, Sanjiban Choudhury, and Sebastian 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] Ziang Cao, Changhong Fu, Junjie Ye, Bowen Li, and Yiming Li. HiFT: Hierarchical Feature Transformer for Aerial Tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15457–15466, 2021. 1, 2, 7

[7] Ziang Cao, Changhong Fu, Junjie Ye, Bowen Li, and Yiming 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 3086–3092, 2021. 2, 7

[8] Ziang Cao, Ziyuan Huang, Liang Pan, Shiwei Zhang, Ziwei Liu, and Changhong Fu. TCTrack: Temporal Contexts for Aerial Tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2022. 1, 2

[9] Yu-Jia Chen, Deng-Kai Chang, and Cheng Zhang. Autonomous Tracking Using a Swarm of UAVs: A Constrained Multi-Agent Reinforcement Learning Approach. *IEEE Transactions on Vehicular Technology*, 69(11):13702–13717, 2020. 1

[10] Zedu Chen, Bineng Zhong, Guorong Li, Shengping Zhang, and Rongrong Ji. Siamese Box Adaptive Network for Visual Tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6668–6677, 2020. 7

[11] Yutao Cui, Cheng Jiang, Limin Wang, and Gangshan Wu. Mixformer: End-to-End Tracking with Iterative Mixed Attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13608–13618, 2022. 7, 12, 13, 14

[12] 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 (CVPR)*, pages 4655–4664, 2019. 7

[13] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. ECO: Efficient Convolution Operators for Tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6638–6646, 2017. 2

[14] 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 (CVPR)*, pages 7181–7190, 2020. 7

[15] Martin Danelljan, Gustav Hager, Fahad Shahbaz Khan, and Michael Felsberg. Learning Spatially Regularized Correlation Filters for Visual Tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4310–4318, 2015. 2

[16] Dawei Du, Yuankai Qi, Hongyang Yu, Yifan Yang, Kaiwen Duan, Guorong Li, Weigang Zhang, Qingming Huang, and Qi Tian. The Unmanned Aerial Vehicle Benchmark: Object Detection and Tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 370–386, 2018. 1, 2, 6, 7, 8, 12, 13, 14, 15, 17

[17] Matteo Dunnhofer, Antonino Furnari, G. Farinella, and C. Micheloni. Is First Person Vision Challenging for Object Tracking? The TREK-100 Benchmark Dataset. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 300–317, 2020. 3

[18] 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 (CVPR)*, pages 5369–5378, 2019. 3, 6, 17

[19] Tharindu Fernando, Simon Denman, Sridha Sridharan, and Clinton Fookes. Tracking by Prediction: A Deep Generative Model for Mutli-Person Localisation and Tracking. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1122–1132, 2018. 3

[20] Changhong Fu, Ziang Cao, Yiming Li, Junjie Ye, and Chen Feng. Siamese Anchor Proposal Network for High-Speed Aerial Tracking. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 510–516, 2021. 1, 2, 7

[21] Changhong Fu, Bowen Li, Fangqiang Ding, Fuling Lin, and Geng Lu. Correlation Filters for Unmanned Aerial Vehicle-

Based Aerial Tracking: A Review and Experimental Evaluation. *IEEE Geoscience and Remote Sensing Magazine*, 10(1):125–160, 2022. 1, 6

[22] Dongyan Guo, Yanyan Shao, Ying Cui, Zhenhua Wang, Liyan Zhang, and Chunhua Shen. Graph Attention Tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9543–9552, 2021. 2, 7, 12, 13

[23] Dongyan Guo, Jun Wang, Ying Cui, Zhenhua Wang, and Shengyong Chen. SiamCAR: Siamese Fully Convolutional Classification and Regression for Visual Tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6269–6277, 2020. 2, 7

[24] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social Gan: Socially Acceptable Trajectories with Generative Adversarial Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2255–2264, 2018. 7, 12, 13

[25] Tuomas Haarnoja, Anurag Ajay, Sergey Levine, and Pieter Abbeel. Backprop KF: Learning Discriminative Deterministic State Estimators. In *Proceedings of the Advances in Neural Information Processing Systems (NIPS)*, volume 29, 2016. 8, 13

[26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2018. 6, 7, 12

[27] 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, 2015. 2

[28] 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*, pages 1–1, 2019. 3, 6, 17

[29] Ziyuan Huang, Changhong Fu, Yiming Li, Fuling Lin, and Peng Lu. Learning Aberrance Repressed Correlation Filters for Real-Time UAV Tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2891–2900, 2019. 1, 2

[30] Rudolph Emil Kalman. A New Approach to Linear Filtering and Prediction Problems. *Transactions of the ASME-Joural of Basic Engineering*, pages 35–45, 1960. 2, 3, 4, 8, 12

[31] Bowen Li, Changhong Fu, Fangqiang Ding, Junjie Ye, and Fuling Lin. ADTrack: Target-Aware Dual Filter Learning for Real-Time Anti-Dark UAV Tracking. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 496–502, 2021. 3

[32] Bowen Li, Yiming Li, Junjie Ye, Changhong Fu, and Hang Zhao. Predictive Visual Tracking: A New Benchmark and Baseline Approach. *arXiv preprint arXiv:2103.04508*, pages 1–8, 2021. 2, 3, 4, 8, 13, 14

[33] 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 (CVPR)*, pages 4277–4286, 2019. 1, 2, 6, 7, 8, 12, 13, 17, 18

[34] Bo Li, Junjie Yan, Wei Wu, Zheng Zhu, and Xiaolin Hu. High Performance Visual Tracking with Siamese Region Proposal Network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8971–8980, 2018. 1, 2, 7

[35] Fan Li, Changhong Fu, Fuling Lin, Yiming Li, and Peng Lu. Training-Set Distillation for Real-Time UAV Object Tracking. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 9715–9721, 2020. 2

[36] Mengtian Li, Yu-Xiong Wang, and Deva Ramanan. Towards Streaming Perception. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 473–488, 2020. 2, 3, 4, 8, 13, 14

[37] Siyi Li and Dit-Yan 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 4140–4146, 2017. 2, 4, 6, 7, 8, 12, 13, 14, 15

[38] Yiming Li, Changhong Fu, Fangqiang Ding, Ziyuan Huang, and Geng Lu. AutoTrack: Towards High-Performance Visual Tracking for UAV with Automatic Spatio-Temporal Regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11920–11929, 2020. 1, 2

[39] Yiming Li, Changhong Fu, Ziyuan Huang, Yinqiang Zhang, and Jia Pan. Keyfilter-Aware Real-Time UAV Object Tracking. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 193–199, 2020. 2

[40] Junwei Liang, Lu Jiang, Juan Carlos Niebles, Alexander G Hauptmann, and Li Fei-Fei. Peeking into the Future: Predicting Future Person Activities and Locations in Videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5725–5734, 2019. 7, 12, 13

[41] Ming Liang, Bin Yang, Wenyuan Zeng, Yun Chen, Rui Hu, Sergio Casas, and Raquel Urtasun. PnPNet: End-to-End Perception and Prediction With Tracking in the Loop. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11553–11562, 2020. 3

[42] Ningxin Liang, Guile Wu, Wenxiong Kang, Zhiyong Wang, and David Dagan Feng. Real-Time Long-Term Tracking with Prediction-Detection-Correction. *IEEE Transactions on Multimedia*, 20(9):2289–2302, 2018. 3

[43] Qiao Liu, Zhenyu He, Xin Li, and Yuan Zheng. PTB-TIR: A Thermal Infrared Pedestrian Tracking Benchmark. *IEEE Transactions on Multimedia*, 22:666–675, 2020. 3

[44] Yuan Liu, Ruoteng Li, Yu Cheng, Robby T Tan, and Xiubao Sui. Object Tracking using Spatio-Temporal Networks for Future Prediction Location. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 1–17, 2020. 3

[45] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *Proceedings of the International Confer-

ence on Learning Representations (ICLR), pages 1–18, 2018. 13

[46] Wenjie Luo, Bin Yang, and Raquel Urtasun. Fast and Furious: Real Time End-to-End 3d Detection, Tracking and Motion Forecasting with a Single Convolutional Net. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3569–3577, 2018. 3

[47] Matthias Mueller, Neil Smith, and Bernard Ghanem. A Benchmark and Simulator for UAV Tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 445–461, 2016. 2, 3, 6, 8, 12, 14, 15, 16

[48] M. Müller, Adel Bibi, Silvio Giancola, Salman Al-Subaihi, 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. 3

[49] Yuya Nishida, Takashi Sonoda, Shinsuke Yasukawa, Kazunori Nagano, Mamoru Minami, Kazuo Ishii, and Tamaki Ura. Underwater Platform for Intelligent Robotics and its Application in Two Visual Tracking Systems. *Journal of Robotics and Mechatronics*, 30(2):238–247, 2018. 1

[50] Nicola A Piga, Ugo Pattacini, and Lorenzo Natale. A Differentiable Extended Kalman Filter for Object Tracking Under Sliding Regime. *Frontiers in Robotics and AI*, 8:686447, 2021. 8, 13

[51] Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M Kitani, Dariu M Gavrila, and Kai O Arras. Human Motion Trajectory Prediction: A Survey. *The International Journal of Robotics Research*, 39(8):895–935, 2020. 2, 3

[52] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 4, 6, 12, 17

[53] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted Residuals and Linear Bottlenecks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4510–4520, 2018. 6, 7, 12

[54] Gur-Eyal Sela, Ionel Gog, Justin Wong, Kumar Krishna Agrawal, Xiangxi Mo, Sukrit Kalra, Peter Schafhalter, Eric Leong, Xin Wang, Bharathan Balaji, et al. Context-Aware Streaming Perception in Dynamic Environments. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 621–638, 2022. 3

[55] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the Inception Architecture for Computer Vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016. 6, 12

[56] 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, 2021. 7

[57] 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 (CVPR)*, pages 1328–1338, 2019. 2, 7, 8, 12, 13, 17, 18

[58] Yinda Xu, Zeyu Wang, Zuoxin Li, Ye Yuan, and Gang Yu. SiamFC++: Towards Robust and Accurate Visual Tracking with Target Estimation Guidelines. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 12549–12556, 2020. 2, 7

[59] Jinrong Yang, Songtao Liu, Zeming Li, Xiaoping Li, and Jian Sun. Real-time Object Detection for Streaming Perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2022. 3

[60] Botao Ye, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Joint Feature Learning and Relation Modeling for Tracking: A One-Stream Framework. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 341–357, 2022. 7, 12, 13, 14

[61] Junjie Ye, Changhong Fu, Fuling Lin, Fangqiang Ding, Shan An, and Geng Lu. Multi-Regularized Correlation Filter for UAV Tracking and Self-Localization. *IEEE Transactions on Industrial Electronics*, 69(6):6004–6014, 2022. 1

[62] Junjie Ye, Changhong Fu, Guangze Zheng, Danda Pani Paudel, and Guang Chen. Unsupervised Domain Adaptation for Nighttime Aerial Tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8896–8905, 2022. 3

[63] Pu Zhang, Wanli Ouyang, Pengfei Zhang, Jianru Xue, and Nanning Zheng. SR-LSTM: State Refinement for LSTM towards Pedestrian Trajectory Prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12085–12094, 2019. 7, 12, 13

[64] Zheng Zhu, Qiang Wang, Bo Li, Wei Wu, Junjie Yan, and Weiming Hu. Distractor-aware Siamese Networks for Visual Object Tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2, 7

Appendix

A. Overview

To make our end-to-end predictive visual tracking framework (PVT++) reproducible, we present the detailed configuration in Appendix B, covering the specific model structure, the training settings (with specific hyper-parameters), and the inference settings. Moreover, we provide the PVT++ code library and official models to ensure reproducability. For clear reference of the notations used in method section, we provide a notation table in Appendix C. In Appendix D, we display representative qualitative visualization results from the authoritative datasets, UAV123 [47], UAV20L [47], DTB70 [37], and UAVDT [16], where the superiority of our PVT++ is clearly shown. In Appendix E, we present detailed results comparison between KF [30] and PVT++ to better demonstrate the superiority of our method. In addition to convolution neural network backbone[26, 53, 55]-based trackers, PVT++ further works for transformer-based ones [60, 11], which is presented in Appendix F. In Appendix G, we show that PVT++ is more efficiency and introduces much less extra latency onboard compared with other trajectory preductors [24, 63, 40]. We also tried to fuse the motion and visual cues earlier in Appendix H, where we give an analysis to the strategy adopted in PVT++. The full attribute-based results from all the four datasets [47, 37, 16] are reported in Appendix I, where we exhaustively analyse the specific advantages of two modalities for prediction under various UAV tracking challenges. The training process of different PVT++ models is visualized in Appendix J, where we present the loss curves to indicate the converging process. The extra latency introduced by the PVT++ predictor modules is unavoidable, which can have some negative effect to online performance. We provide such analysis in Appendix K. We further find PVT++ is capable of converging well in smaller training set (using only 3563 videos from Imagenet VID [52]), which is shown in Appendix L. Finally, we present additional real-world tests in Appendix M, covering more target objects and tracking scenes.

B. Detailed Configuration

Specific Model Structure. Corresponding to Fig. 4 in the paper, we present the detailed model structure of each layer in Table I. Consider B batch inputs and k history frames, the output sizes are also shown in Table I for clear reference. Subscripts are used to distinguish between different layers, *i.e.*, $\cdot_t$ denotes encoding layer for template feature, $\cdot_s$ denotes encoding layer for search feature, $\cdot_e$ denotes encoding layer for the similarity map. $\cdot_a$ represents the auxiliary branch.

Remark 5: These structures are general for all the four implemented base trackers [33, 57, 22].

Table I. Detailed structure and output sizes of PVT++ models. We use subscript to distinguish between different layers. The output sizes correspond to B batch input.

Branch	Layer	Kernel	C_{in}	C_{out}	Out. Size
Motion	FC	-	8	32	$B \times k \times 32$
	1D Conv	3	32	32	$B \times k \times 32$
	Avg. Pool	-	32	32	$B \times 32$
Visual	2D Conv$_t$	3×3	256	64	$B \times k \times 64 \times 29 \times 29$
	2D Conv$_s$	3×3	256	64	$B \times k \times 64 \times 25 \times 25$
	2D Conv$_e$	1×1	64	64	$B \times k \times 64 \times 25 \times 25$
	3D Conv	$3 \times 3 \times 3$	64	64	$B \times k \times 64 \times 25 \times 25$
	Avg. Pool	-	64	64	$B \times 64$
	2D Conv$_a$	1×1	64	64	$B \times k \times 64 \times 25 \times 25$
	2D Conv$_a$	1×1	64	4	$B \times k \times 4 \times 25 \times 25$
	Avg. Pool$_a$	-	4	4	$B \times k \times 4$
Shared	FC	-	[32, 64, 96]	32	$B \times 32$
	FC	-	32	32	$B \times N \times 32$
	FC	-	32	4	$B \times N \times 4$

Table II. List of the important notations in this work.

Symbol	Meaning	Dimension
f	World frame number	$\mathbb{R}$
$\mathcal{I}_f$	f-th image frame	$\mathbb{R}^{W \times H \times 3}$
j	Serial number of the processed frame	$\mathbb{R}$
f_j	World frame id of the processed j-th frame	$\mathbb{R}$
t_f^{W}	World timestamp	$\mathbb{R}$
$t_{f_j}^{\text{T}}$	Tracker timestamp	$\mathbb{R}$
$\phi(f), \phi(f)_e$	Input frame id to be paired with frame f	$\mathbb{R}$
σ	Permitted latency during evaluation	$\mathbb{R}$
$\mathbf{r}_f = [x_f, y_f, w_f, h_f]$	Raw output by the tracker in frame f	$\mathbb{R}^{1 \times 4}$
$\hat{\mathbf{b}}_f = [\hat{x}_f, \hat{y}_f, \hat{w}_f, \hat{h}_f]$	Final output bounding box to be evaluated	$\mathbb{R}^{1 \times 4}$
$\mathcal{T}$	Tracker model	—
$\mathcal{P}$	Predictor model	—
$\mathbf{m}_{f_j}$	Normalized input motion from f_{j-1} to f_j	$\mathbb{R}^{1 \times 4}$
$\mathbf{p}_{f_j}$	Average moving speed from f_{j-k+1} to f_j	$\mathbb{R}^{1 \times 4}$
$\hat{\mathbf{m}}_f$	Predicted motion from $\phi(f)$ to f	$\mathbb{R}^{1 \times 4}$
$\mathbf{m}_f$	Ground-truth motion from $\phi(f)$ to f	$\mathbb{R}^{1 \times 4}$
Δ_f	Frame interval between the latest and the f-th frame	$\mathbb{R}$
$\Delta_{\hat{x}}(f), \Delta_{\hat{y}}(f)$	Predicted distance between the f-th and $\phi(f)$-th frame	$\mathbb{R}$
$\Delta_x(f_j), \Delta_y(f_j)$	Distance from $\mathbf{r}_{f_j}$ to $\mathbf{r}_{f_j-1}$	$\mathbb{R}$
$\mathbf{x}_{\phi(f)}$	Search patch feature in frame $\phi(f)$	$\mathbb{R}^{C \times W \times H}$
$\mathbf{z}$	Template feature	$\mathbb{R}^{C \times a \times a}$
$k(= 3)$	Number of past frames	$\mathbb{R}$
N	Number of the parallel FC layers in the decoder	$\mathbb{R}$

plemented base trackers [33, 57, 22].

Training Settings. All the predictive modules need temporal video data for training. However, to our disappointment, existing training pipeline [33] takes a detection-like paragdim. Basically, the raw search patches are *independently* cropped from the object center location, then the random shift, padding are applied to generated the training search patch. In this case, the training patches from consecutive frames actually contain no temporal information.

To solve this, we construct a new pipeline termed as dynamic temporal training. The search patch from f_j-th frame is cropped around the object's center location in the previous frame $\mathcal{I}_{f_{j-1}}$, so that past motion $\mathbf{M}_{\phi(f)}$ and past search patch $\mathbf{X}_{\phi(f)}$ correspond to each other and contain real temporal information from $\mathcal{I}_{f_{j-k+1}}$ to $\mathcal{I}_{f_j}$.

Remark 6: The new training pipeline is dynamic, *i.e.*, $[f_{j-k}, f_{j-k+1}, \cdots, f_j]$ can be adjusted as hyper-parameters to fit different models' different latency.

Table III. Per dataset results of different predictor modules. For all the three base trackers in various datasets, our PVT++ generally outperforms previous standard KF solutions [36, 32] and stronger learnable KF baselines, KF† and KF‡.

| Dataset | | DTB70 | | UAVDT | | UAV20L | | UAV123 | |
Tracker	Pred.	AUC@La0	DP@La0	AUC@La0	DP@La0	AUC@La0	DP@La0	AUC@La0	DP@La0
SiamRPN++$_\text{M}$ (21FPS)	N/A	0.305	0.387	0.494	0.719	0.448	0.619	0.472	0.678
	KF [36]	0.349	0.482	0.527	0.737	0.458	0.624	0.515	0.712
	PVT [32]	0.377	0.518	0.533	0.740	0.458	0.624	0.522	0.722
	KF† [50]	0.367	0.504	0.519	0.732	0.466	0.630	0.511	0.703
	KF‡ [25]	0.365	0.496	0.563	0.780	0.483	0.658	0.513	0.598
	$\mathcal{P}_\text{M}$ (**Ours**)	0.385	0.523	0.529	0.745	0.481	0.647	0.537	0.737
	$\mathcal{P}_\text{V}$ (**Ours**)	0.352	0.472	0.564	0.799	0.488	0.675	0.504	0.703
	$\mathcal{P}_\text{MV}$ (**Ours**)	**0.399**	**0.536**	**0.576**	**0.807**	**0.508**	**0.697**	**0.537**	**0.741**
SiamMask (12FPS)	N/A	0.247	0.313	0.455	0.703	0.405	0.571	0.436	0.639
	KF [36]	0.294	0.407	0.535	0.758	0.436	0.582	0.499	0.679
	PVT [32]	0.362	0.504	0.539	0.751	0.443	0.598	0.514	0.701
	KF† [50]	0.349	0.486	0.530	0.749	0.440	0.588	0.513	0.702
	KF‡ [25]	0.348	0.468	0.558	0.775	0.465	0.629	0.502	0.683
	$\mathcal{P}_\text{M}$ (**Ours**)	**0.370**	**0.508**	0.531	0.760	0.449	0.607	0.532	0.743
	$\mathcal{P}_\text{V}$ (**Ours**)	0.292	0.405	0.532	0.777	0.430	0.601	0.503	0.705
	$\mathcal{P}_\text{MV}$ (**Ours**)	0.342	0.463	**0.566**	**0.797**	**0.469**	**0.644**	**0.536**	**0.749**
SiamRPN++$_\text{M}$ (21FPS)	N/A	0.136	0.159	0.351	0.594	0.310	0.434	0.349	0.505
	KF [36]	0.189	0.232	0.451	0.667	0.387	0.528	0.415	0.582
	PVT [32]	0.201	0.254	0.467	0.687	0.396	0.547	0.434	0.605
	KF† [50]	0.200	0.254	0.460	0.680	0.412	0.572	0.433	0.603
	KF‡ [25]	0.204	0.252	0.504	0.728	0.406	0.549	0.432	0.599
	$\mathcal{P}_\text{M}$ (**Ours**)	0.199	**0.258**	0.449	0.684	0.404	0.560	**0.442**	**0.627**
	$\mathcal{P}_\text{V}$ (**Ours**)	0.179	0.225	0.403	0.665	0.398	0.548	0.398	0.559
	$\mathcal{P}_\text{MV}$ (**Ours**)	**0.205**	0.256	**0.488**	**0.726**	**0.416**	0.568	0.442	0.619

All the PVT++ models are optimized by AdamW [45]. The motion predictor is trained for 100 epochs with a base learning rate equalling to 0.03, which is multiplied by 0.1 at epoch 30 and 80. The visual and multi-modal predictors are trained for 300 epochs with a base learning rate of 0.003, which is multiplied by 0.1 at epoch 200. In all the four base trackers [33, 57, 22], $\mathcal{P}_\text{V}$ and $\mathcal{P}_\text{MV}$ both take the visual feature from the neck to implement vision-aided prediction. During joint training, the tracker backbone is fixed and the tracker neck, together with the head are freed in the first 20 epochs with a small learning rate of 10^{-5}.

A "fast" tracker may only need to predict future three frames to compensate for its latency, while a "slow" one may have to output ten future state. To make this possible, the second last layer of PVT++ predictive decoder is N parallel fully connected layers for predicting N future state, *i.e.*, future $1 \sim N$ frames. Therefore, different models vary in the pre-defined N and Δ_f during training. we set $N = 3, \Delta_f = [1 : 3]$ for SiamRPN++$_\text{M}$ [33], $N = 12, \Delta_f = [1 : 12]$ for SiamRPN++$_\text{R}$ [33], $N = 6, \Delta_f = [1 : 6]$ for SiamMask [57], and $N = 4, \Delta_f = [1 : 3]$ for SiamGAT [22]. Note that these hyper-parameter are roughly determined by the averaged latency of the base trackers.

Inference Settings. During inference, when f_{j+1}-th frame comes, the predictor $\mathcal{P}$ first conducts $(f_{j+1} - f_j)$ to $f_{j+1} + N$ frames prediction with $k = 3$ past frames information, then the tracker processes $f_{j+1} - th$ frame and

Table IV. Efficiency and complexity comparison between PVT++ and other motion predictors [24, 63, 40]. Our framework is 10x~100x faster than other works.

Input	Traj.			Traj. + RGB	
Model	Social GAN [24]	SR-LSTM [63]	$\mathcal{P}_\text{M}$	NEXT [40]	$\mathcal{P}_\text{MV}$
MACs	5.6M	51.7M	**0.05M**	2.7G	**1.2G**
Latency (ms)	50.2	652.0	**4.2**	181.6	**8.6**

Table V. Effect of PVT++ on transformer-based trackers [60, 11]. Our framework can boost the performance by up to **40**%.

Dataset	DTB70				UAVDT			
Metric	AUC@La0		DP@La0		AUC@La0		DP@La0	
PVT++	✗	✓	✗	✓	✗	✓	✗	✓
OSTrack [60]	0.306	**0.400**	0.375	**0.535**	0.533	**0.626**	0.789	**0.839**
MixFormer [11]	0.198	**0.250**	0.242	**0.320**	0.413	**0.516**	0.644	**0.719**

updates the history information (motion and visual).

Note that we take the latency of both tracker and predictor modules into account in the online evaluation.

C. Complete Notation Reference Table

We provide the important notations, their meaning, and dimension in Table II, for clear reference.

D. Visualization

We present some typical tracking visualization in Fig. I. The sequences, *ManRunning2*, *Paragliding5*, *Wakeboarding1*, and *Wakeboarding2* are from DTB70 [37].*S0303*, *S0304*, *S0310*, and *S1604* are from UAVDT [16]. In
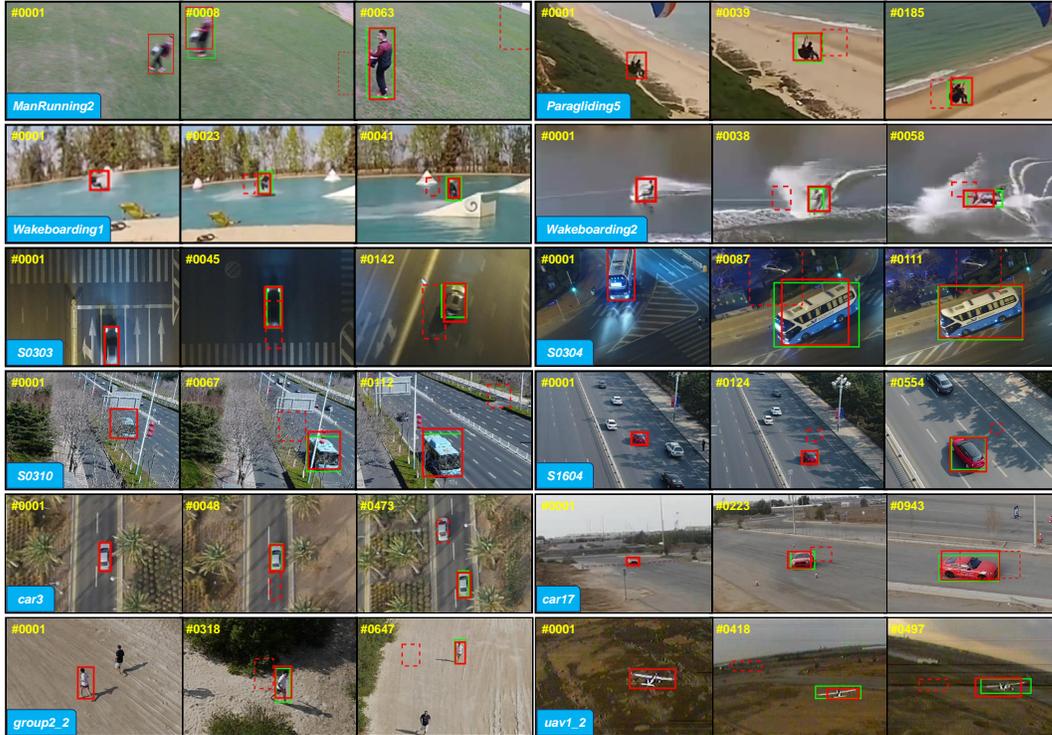
- - - Baseline   **——— Baseline + $\mathcal{P}_{MV}$**   **——— Ground-truth**

Figure I. Representative sequences from authoritative UAV tracking datasets, DTB70 [37], UAVDT [16], UAV20L [47], and UAV123 [47]. We use dashed red lines to demonstrate the original trackers, which are severely affected by onboard latency. Coupled with our PVT++ ($\mathcal{P}_{MV}$), the robustness can be significantly improved (solid red boxes). Green boxes denote ground-truth. Some typical sequences are also made into supplementary video for better reference.

Table VI. Results comparison between two fusion strategy. $\mathcal{P}_{MV}$ denotes our default PVT++, the modalities fuse after independent temporal interaction (late fusion). $\mathcal{P}_{MV}^{\dagger}$ indicates that the two cues fuse before temporal interaction (early fusion).

| | DTB70 | | UAVDT | |
Pred.	AUC@La0	DP@La0	AUC@La0	DP@La0
N/A	0.305	0.387	0.494	0.719
$\mathcal{P}_{MV}$ (late fuse)	**0.399**	**0.536**	**0.576**	**0.807**
$\mathcal{P}_{MV}^{\dagger}$ (early fuse)	0.370	0.498	0.571	0.800

UAV20L and UAV123 [47], we also present *car3*, *car17*, *group2_2*, and *uav1_2*. With extremely limited onboard computation, the original trackers (red dashed boxes) will easily fail due to high latency. Once coupled with our PVT++ ($\mathcal{P}_{MV}$), the models (solid red boxes) are much more robust. We use greed boxes to denote ground-truth.

E. Prediction Quantitative Comparison

To provide a thorough quantitative comparison of the predictor performance, we reported the results per dataset in Table III. We observe that for different tracker models in various benchmarks, PVT++ is more robust than prior solutions [32, 36]. Compared with learnable KFs, KF† and KF$\ddagger$, our PVT++ holds obvious advantage by virtue of the visual cue and joint learning.

F. Effect on Transformer-based Trackers

For transformer-based trackers, MixFormer [11] and OS-Track [60] (~6 and ~10 FPS onboard), PVT++ yields up to 40% improvement as shown in Table V.

G. Efficiency and Complexity Comparison

PVT++ is a lightweight plug-and-play framework designed for latency-aware tracking, while most existing trajectory predictors are computationally heavy. As in Table IV, PVT++ is 10x~100x faster than existing trajectory predictors and introduces much less extra latency onboard.

H. Fusion Strategy Comparison

As introduced in the paper, inside PVT++, the three modules, *Feature encoder*, *temporal interaction*, and *predictive decoder* run one after another. For the default setting, the fusion of the motion and visual cues happens after *temporal interaction*, using the concatenate function. Here, we also tried to integrate the two modality earlier before *temporal interaction* and right after *feature encoder*,

Table VII. Attribute-based analysis of the three trackers with PVT++ models in DTB70 [37] dataset.

Tracker		SiamRPN++$_M$ (21FPS)				SiamRPN++$_R$ (5FPS)				SiamMask (12FPS)			
Metric	Att.	N/A	$\mathcal{P}_M$	$\mathcal{P}_V$	$\mathcal{P}_{MV}$	N/A	$\mathcal{P}_M$	$\mathcal{P}_V$	$\mathcal{P}_{MV}$	N/A	$\mathcal{P}_M$	$\mathcal{P}_V$	$\mathcal{P}_{MV}$
	ARV	0.330	0.386	0.349	0.418	0.156	0.233	0.214	0.253	0.247	0.375	0.291	0.393
	BC	0.257	0.330	0.276	0.319	0.079	0.077	0.102	0.102	0.168	0.264	0.202	0.167
	DEF	0.357	0.410	0.358	0.438	0.144	0.217	0.198	0.241	0.253	0.398	0.287	0.364
	FCM	0.277	0.373	0.333	0.376	0.091	0.144	0.122	0.138	0.195	0.327	0.258	0.301
	IPR	0.302	0.368	0.324	0.387	0.133	0.187	0.169	0.204	0.217	0.346	0.256	0.316
AUC@La0	MB	0.198	0.305	0.277	0.321	0.056	0.073	0.069	0.085	0.147	0.236	0.187	0.254
	OCC	0.280	0.337	0.281	0.304	0.149	0.214	0.204	0.224	0.233	0.290	0.285	0.274
	OPR	0.278	0.314	0.334	0.439	0.161	0.158	0.208	0.225	0.202	0.360	0.265	0.362
	OV	0.292	0.405	0.372	0.399	0.054	0.099	0.076	0.102	0.168	0.227	0.258	0.289
	SV	0.354	0.470	0.419	0.489	0.145	0.187	0.192	0.220	0.278	0.435	0.347	0.418
	SOA	0.238	0.301	0.261	0.302	0.140	0.196	0.184	0.200	0.227	0.326	0.275	0.315
	ARV	0.340	0.466	0.385	0.498	0.101	0.220	0.171	0.234	0.247	0.474	0.333	0.472
	BC	0.352	0.477	0.396	0.498	0.118	0.106	0.141	0.139	0.228	0.385	0.291	0.237
	DEF	0.374	0.512	0.398	0.525	0.083	0.203	0.144	0.214	0.246	0.509	0.326	0.449
	FCM	0.363	0.517	0.470	0.525	0.106	0.188	0.156	0.171	0.241	0.456	0.353	0.414
	IPR	0.349	0.475	0.398	0.495	0.124	0.212	0.170	0.224	0.236	0.454	0.310	0.400
DP@La0	MB	0.246	0.418	0.379	0.453	0.051	0.110	0.090	0.088	0.167	0.349	0.248	0.327
	OCC	0.408	0.496	0.426	0.459	0.223	0.327	0.316	0.344	0.361	0.439	0.458	0.404
	OPR	0.213	0.312	0.317	0.453	0.083	0.083	0.113	0.127	0.128	0.382	0.224	0.357
	OV	0.413	0.590	0.564	0.586	0.062	0.166	0.101	0.161	0.222	0.363	0.385	0.439
	SV	0.366	0.569	0.467	0.569	0.123	0.186	0.180	0.208	0.287	0.528	0.402	0.492
	SOA	0.333	0.432	0.379	0.447	0.217	0.306	0.295	0.302	0.340	0.479	0.429	0.462

Table VIII. Attribute-based analysis of the three trackers with PVT++ models in UAVDT [16] dataset.

Tracker		SiamRPN++$_M$ (21FPS)				SiamRPN++$_R$ (5FPS)				SiamMask (12FPS)			
Metric	Att.	N/A	$\mathcal{P}_M$	$\mathcal{P}_V$	$\mathcal{P}_{MV}$	N/A	$\mathcal{P}_M$	$\mathcal{P}_V$	$\mathcal{P}_{MV}$	N/A	$\mathcal{P}_M$	$\mathcal{P}_V$	$\mathcal{P}_{MV}$
	BC	0.448	0.461	0.504	0.505	0.332	0.410	0.375	0.445	0.404	0.465	0.488	0.520
	CR	0.450	0.495	0.520	0.535	0.296	0.371	0.402	0.452	0.425	0.503	0.498	0.522
	OR	0.438	0.481	0.538	0.549	0.318	0.389	0.416	0.477	0.404	0.491	0.504	0.541
AUC@La0	SO	0.494	0.549	0.525	0.545	0.318	0.420	0.361	0.457	0.468	0.536	0.495	0.540
	IV	0.539	0.578	0.588	0.599	0.382	0.495	0.459	0.537	0.475	0.558	0.563	0.596
	OB	0.525	0.542	0.568	0.589	0.382	0.460	0.408	0.498	0.471	0.542	0.527	0.560
	SV	0.490	0.505	0.584	0.586	0.366	0.422	0.406	0.484	0.438	0.526	0.541	0.566
	LO	0.422	0.521	0.436	0.511	0.320	0.379	0.368	0.429	0.389	0.421	0.494	0.520
	BC	0.659	0.666	0.733	0.727	0.591	0.637	0.647	0.671	0.628	0.672	0.718	0.731
	CR	0.643	0.684	0.720	0.732	0.462	0.585	0.572	0.645	0.620	0.702	0.696	0.712
	OR	0.638	0.681	0.753	0.764	0.515	0.619	0.606	0.688	0.612	0.709	0.723	0.752
DP@La0	SO	0.779	0.815	0.793	0.814	0.645	0.711	0.706	0.759	0.803	0.818	0.787	0.819
	IV	0.777	0.811	0.835	0.848	0.657	0.747	0.755	0.801	0.743	0.797	0.817	0.829
	OB	0.772	0.778	0.822	0.846	0.676	0.714	0.700	0.766	0.756	0.802	0.801	0.813
	SV	0.680	0.691	0.796	0.794	0.581	0.618	0.622	0.684	0.650	0.729	0.763	0.783
	LO	0.569	0.717	0.585	0.694	0.504	0.554	0.566	0.608	0.571	0.590	0.696	0.711

still adopting concatenation. The results comparison of two strategies is shown in Table VI, where we find both are effective and the late fusion is better.

I. Full Attribute-based Analysis

We present full attribute-based analysis in Table VII, Table VIII, Table IX, and Table X. Following the previous work [37], we report results on aspect ratio variation (ARV), background clutter (BC), deformation (DEF), fast camera motion (FCM), in-plane rotation (IPR), motion blur (MB), occlusion (OCC), out-of-plane rotaTion (OPR), out-of-view (OV), scale variation (SV), and similar object around (SOA) in Table VII. As shown in Table VIII, results on back-

ground clutter (BC), camera rotation (CR), object rotation (OR), small object (SO), illumination variation (IV), object blur (OB), scale variation (SV), and large occlusion (LO), are reported for UAVDT [16]. For UAV20L and UAV123 [47], we present results on scale variation (SV), aspect ratio change (ARC), low resolution (LR), fast motion (FM), full occlusion (FOC), partial occlusion (POC), out-of-view (OV), background clutter (BC), illumination variation (IV), viewpoint change (VC), camera motion (CM), and similar object (SO) in Table IX and Table X.

We observe that the two modalities has their own advantage in different UAV tracking challenges. For example, consider UAVDT dataset [16] (Table VIII), the visual

Table IX. Attribute-based analysis of the three trackers with PVT++ models in UAV20L [47] dataset.

Tracker		SiamRPN++$_M$ (21FPS)				SiamRPN++$_R$ (5FPS)				SiamMask (12FPS)			
Metric	Att.	N/A	$\mathcal{P}_M$	$\mathcal{P}_V$	$\mathcal{P}_{MV}$	N/A	$\mathcal{P}_M$	$\mathcal{P}_V$	$\mathcal{P}_{MV}$	N/A	$\mathcal{P}_M$	$\mathcal{P}_V$	$\mathcal{P}_{MV}$
AUC@La0	SV	0.437	0.470	0.483	0.500	0.300	0.395	0.392	0.410	0.395	0.437	0.420	0.461
	ARC	0.425	0.411	0.438	0.451	0.291	0.352	0.360	0.371	0.373	0.409	0.392	0.438
	LR	0.267	0.354	0.344	0.352	0.215	0.295	0.276	0.279	0.244	0.263	0.275	0.290
	FM	0.410	0.357	0.394	0.418	0.269	0.304	0.325	0.315	0.319	0.375	0.329	0.442
	FOC	0.256	0.272	0.234	0.241	0.170	0.227	0.184	0.164	0.221	0.237	0.231	0.255
	POC	0.418	0.480	0.463	0.478	0.286	0.379	0.380	0.396	0.378	0.417	0.430	0.441
	OV	0.438	0.512	0.476	0.492	0.272	0.356	0.394	0.405	0.377	0.428	0.448	0.462
	BC	0.225	0.258	0.229	0.250	0.119	0.215	0.153	0.159	0.189	0.198	0.210	0.210
	IV	0.452	0.414	0.470	0.491	0.303	0.393	0.379	0.403	0.426	0.437	0.382	0.443
	VC	0.472	0.450	0.466	0.488	0.302	0.339	0.377	0.384	0.395	0.436	0.420	0.475
	CM	0.431	0.463	0.475	0.491	0.297	0.393	0.388	0.406	0.391	0.432	0.412	0.452
	SO	0.482	0.519	0.557	0.567	0.399	0.531	0.477	0.491	0.487	0.519	0.438	0.492
DP@La0	SV	0.600	0.630	0.662	0.683	0.417	0.544	0.536	0.556	0.552	0.588	0.581	0.627
	ARC	0.591	0.562	0.606	0.624	0.408	0.487	0.486	0.503	0.524	0.558	0.550	0.603
	LR	0.444	0.545	0.539	0.548	0.388	0.483	0.465	0.456	0.422	0.414	0.458	0.465
	FM	0.631	0.548	0.595	0.625	0.417	0.464	0.495	0.476	0.518	0.573	0.524	0.667
	FOC	0.469	0.473	0.436	0.428	0.358	0.423	0.358	0.324	0.425	0.420	0.431	0.459
	POC	0.585	0.654	0.648	0.669	0.410	0.530	0.531	0.548	0.540	0.570	0.606	0.613
	OV	0.597	0.683	0.658	0.679	0.356	0.473	0.518	0.540	0.529	0.578	0.618	0.630
	BC	0.426	0.440	0.399	0.434	0.284	0.398	0.304	0.295	0.378	0.349	0.390	0.385
	IV	0.628	0.560	0.649	0.686	0.428	0.551	0.503	0.539	0.595	0.590	0.545	0.617
	VC	0.616	0.571	0.605	0.631	0.364	0.420	0.452	0.477	0.518	0.551	0.546	0.611
	CM	0.599	0.629	0.660	0.681	0.417	0.544	0.534	0.553	0.550	0.588	0.580	0.626
	SO	0.604	0.652	0.719	0.734	0.498	0.645	0.594	0.609	0.610	0.648	0.559	0.619

Table X. Attribute-based analysis of the three trackers with PVT++ models in UAV123 [47] dataset.

Tracker		SiamRPN++$_M$ (21FPS)				SiamRPN++$_R$ (5FPS)				SiamMask (12FPS)			
Metric	Att.	N/A	$\mathcal{P}_M$	$\mathcal{P}_V$	$\mathcal{P}_{MV}$	N/A	$\mathcal{P}_M$	$\mathcal{P}_V$	$\mathcal{P}_{MV}$	N/A	$\mathcal{P}_M$	$\mathcal{P}_V$	$\mathcal{P}_{MV}$
AUC@La0	SV	0.456	0.518	0.488	0.514	0.338	0.423	0.383	0.427	0.420	0.509	0.480	0.518
	ARC	0.413	0.496	0.468	0.491	0.315	0.402	0.365	0.406	0.398	0.498	0.467	0.510
	LR	0.291	0.357	0.328	0.350	0.179	0.264	0.214	0.256	0.257	0.364	0.324	0.257
	FM	0.373	0.430	0.461	0.482	0.261	0.316	0.307	0.341	0.333	0.425	0.422	0.447
	FOC	0.254	0.317	0.270	0.306	0.191	0.251	0.214	0.246	0.242	0.325	0.284	0.303
	POC	0.401	0.436	0.402	0.446	0.284	0.373	0.335	0.374	0.363	0.449	0.426	0.459
	OV	0.442	0.489	0.488	0.516	0.289	0.394	0.368	0.407	0.403	0.504	0.476	0.492
	BC	0.254	0.293	0.247	0.296	0.188	0.258	0.215	0.247	0.248	0.360	0.307	0.309
	IV	0.365	0.421	0.423	0.465	0.310	0.379	0.352	0.381	0.378	0.480	0.441	0.466
	VC	0.459	0.552	0.506	0.558	0.322	0.409	0.387	0.432	0.407	0.534	0.499	0.548
	CM	0.466	0.542	0.514	0.535	0.319	0.421	0.381	0.422	0.420	0.529	0.502	0.522
	SO	0.478	0.497	0.444	0.459	0.362	0.462	0.382	0.435	0.434	0.492	0.464	0.514
DP@La0	SV	0.657	0.714	0.679	0.710	0.488	0.599	0.594	0.537	0.614	0.711	0.671	0.720
	ARC	0.602	0.689	0.651	0.678	0.453	0.575	0.502	0.561	0.588	0.701	0.656	0.715
	LR	0.548	0.595	0.568	0.586	0.392	0.488	0.471	0.438	0.510	0.621	0.554	0.637
	FM	0.517	0.591	0.617	0.646	0.323	0.417	0.368	0.429	0.450	0.588	0.564	0.609
	FOC	0.497	0.550	0.489	0.533	0.387	0.460	0.406	0.448	0.460	0.569	0.505	0.541
	POC	0.614	0.630	0.586	0.640	0.440	0.556	0.497	0.542	0.553	0.653	0.619	0.664
	OV	0.632	0.670	0.674	0.715	0.372	0.533	0.467	0.533	0.556	0.701	0.653	0.685
	BC	0.474	0.475	0.436	0.489	0.407	0.470	0.411	0.444	0.473	0.587	0.512	0.526
	IV	0.546	0.594	0.586	0.644	0.447	0.541	0.521	0.486	0.550	0.674	0.623	0.664
	VC	0.654	0.743	0.681	0.746	0.443	0.575	0.512	0.586	0.587	0.735	0.683	0.744
	CM	0.668	0.748	0.713	0.735	0.440	0.587	0.514	0.573	0.606	0.737	0.699	0.734
	SO	0.714	0.703	0.625	0.647	0.554	0.681	0.568	0.639	0.650	0.691	0.671	0.724

branch is relatively good at challenges like camera rotation (CR), object rotation (OR), and scale variation (SV), where the object motion could be very complex and the visual appearance is helpful in prediction. On the other hand, motion cues are robust when the visual feature is not reliable, for instance, similar object (SO) and large occlusion (LO)

challenge. In general, motion predictor is better than visual predictor, from which we conclude that past motion is still the main cue to inference future motion. While for the challenging UAV tracking, where motion could be extremely random and dynamic, introducing visual cues can significant improve the prediction robustness. Together, the
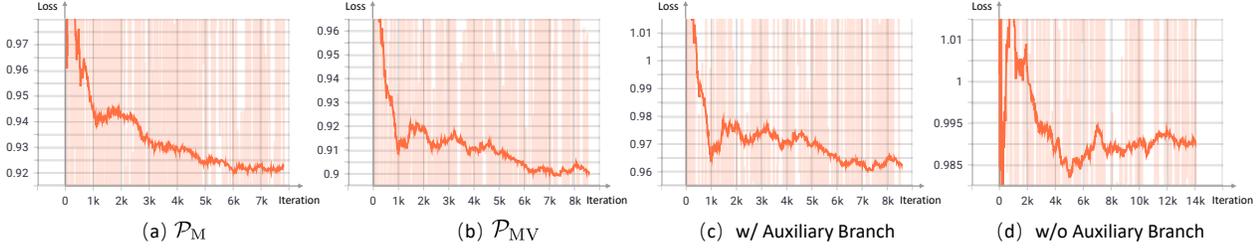
(a) $\mathcal{P}_M$ (b) $\mathcal{P}_{MV}$ (c) w/ Auxiliary Branch   (d) w/o Auxiliary Branch

Figure II. Training loss curves of PVT++ models. Coupled with visual feature, $\mathcal{P}_{MV}$ can better learn to predict than $\mathcal{P}_M$, thus the loss is observed to be smaller. Without auxiliary branch, the loss curve is less smooth, indicating the importance of $\mathcal{A}$.

Table XI. Effect of extra latency brought by PVT++ in UAVDT [16] dataset. Here, the base tracker takes SiamRPN++$_M$, whose original latency is fixed to 44.5 ms/frame (its average onboard latency). We use $\cdot^{\dagger}$ to indicate neglecting the latency. With $\sim$5ms/frame extra time, the performance is slightly lower (2$\sim$3% performance drop), while it is acceptable and still brings upto 15% performance gain.

Model	Tracker			Tracker+$\mathcal{P}_{MV}^{\dagger}$			Tracker+$\mathcal{P}_{MV}$		
Metric	mAUC$_{\Delta\%}$	mDP$_{\Delta\%}$	Latency	mAUC$_{\Delta\%}$	mDP$_{\Delta\%}$	Latency	mAUC$_{\Delta\%}$	mDP$_{\Delta\%}$	Latency
Result	0.494$_{+0.00}$	0.719$_{+0.00}$	44.5ms	0.587$_{+18.8}$	0.825$_{+14.7}$	44.5ms	0.576$_{+16.6}$	0.807$_{+12.2}$	50.0ms

Table XII. Performance of PVT++ models trained with different datasets. Full denotes $\sim$9,000 videos from VID [52], LaSOT [18], and GOT-10k [28]. VID indicates using only $\sim$3,000 videos from VID [52]. AVG means average results on the four test datasets. Since PVT++ utilizes the trained tracking models, We observe the training are not very sensitive to the scale of training set.

Dataset		DTB70		UAVDT		UAV20L		UAV123		AVG	
PVT++	Training	mAUC	mDP	mAUC	mDP	mAUC	mDP	mAUC	mDP	mAUC	mDP
$\mathcal{P}_V$	Full	0.352	0.472	0.564	0.799	0.488	0.675	0.504	0.703	**0.477**	0.662
	VID	0.362	0.483	0.519	0.752	0.497	0.694	0.513	0.731	0.473	**0.665**
$\mathcal{P}_{MV}$	Full	0.399	0.536	0.576	0.807	0.508	0.697	0.537	0.741	**0.505**	**0.695**
	VID	0.405	0.554	0.53	0.757	0.511	0.701	0.534	0.745	0.495	0.689

jointly optimized model $\mathcal{P}_{MV}$ is the most reliable for UAV latency-aware vsiaul tracking.

J. Training Visualization

The training loss curves of PVT++ models with SiamRPN++$_M$ [33] is shown in Fig. II. Compared with motion predictor $\mathcal{P}_M$, the joint predictor $\mathcal{P}_{MV}$ can better learn to predict, resulting in smaller training loss. We also compared the losses from models with (c) or without (d) the auxiliary branch $\mathcal{A}$. Without $\mathcal{A}$, the loss curve fluctuates a lot, indicating that the model can't converge very well.

K. Effect of Extra Latency

PVT++ will bring a bit extra latency during online perception, which is negative for the performance. As shown in Table XI, the latency of original tracker [33] is about 45 ms/frame. Ignoring the predictor's latency, the online performance can reach 0.587 mAUC and 0.825 mDP. Taking the extra latency of $\sim$ 5 ms/frame into account, the result will slightly suffer, decreasing to 0.576 mAUC and 0.807 mDP. Therefore, though PVT++ introduces extra latency, the online performance can still be significantly improved by more than **10**%.

L. Training Set Analysis

Since PVT++ models can make full use of a trained tracker model, we find $\mathcal{P}_V$ and $\mathcal{P}_{MV}$ not very sensitive to the scale of training set. As shown in Table XII, trained with only $\sim$3,000 videos from VID [52], our PVT++ can still converge well and achieve on par performance compared with the fully trained models.

M. More Real-World Tests

In addition to the four real-world tests in Sec. 6.5 of the main paper, we present six more tests (together eight tests) in Fig. III, where we implemented the models on a real UAV and performed several flights. The real-world tests involve two non-real-time trackers, SiamRPN++$_M$ [33] ($\sim$ 15.57 FPS in the tests) and SiamMask [57] ($\sim$ 11.95 FPS in the tests), which are largely affected by their high onboard latency. Coupled with our PVT++ ($\mathcal{P}_{MV}$), the predictive models work well under various tracking scenes, *e.g.*, aspect ratio change in Test 1, dark environment in Test 2, 5, 7, and 8, view point change in Test 3, and occlusion in Test 2. The real-world tests also cover various target objects like person, building, car, and island, as shown in Fig. III. We have made them into videos for clear reference. The robustness of PVT++ in the onboard tests validate its effectiveness in the real-world UAV tracking challenges.
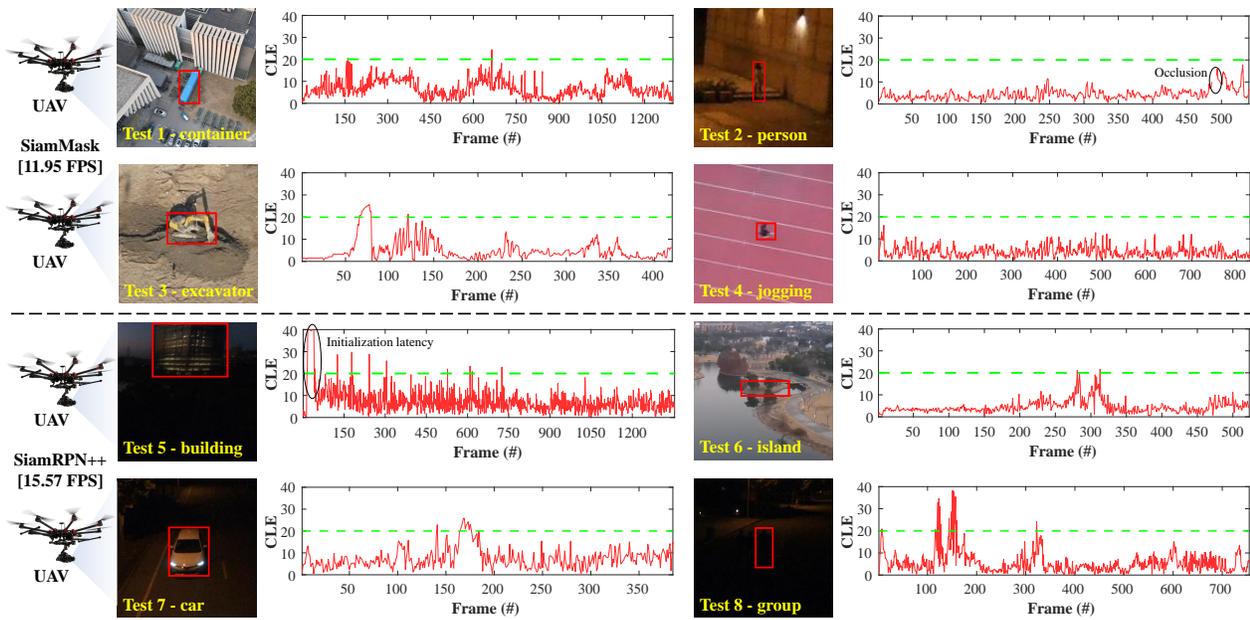
Figure III. Eight real-world tests of PVT++ on non-real-time trackers, SiamMask [57] and SiamRPN++$_\mathrm{M}$ [33]. We present the tracking scenes, the target objects, and center location error (CLE) in the figure. Under various challenges like aspect ration change, illumination variation, low resolution, PVT++ maintains its robustness, with CLE below 20 pixels in most frames.