Learning to Adversarially Blur Visual Object Tracking Supplementary Material

Anonymous ICCV submission

Paper ID 1406

In this material, we report more results of OP-ABA for attacking KYS [4] and DiMP [3]. We also provide extra ablation study by showing the loss values of OP-ABA w/o $\mathcal{A}$, OP-ABA w/o $\mathcal{W}$, and OP-ABA during optimization. In addition, we provide seven attacking cases with 'gif' files in the zip.

1. More Results of OP-ABA against Other Trackers

In this material, we report more results of OP-ABA for attacking KYS and DiMP on the OTB100, UAV123, LaSOT, and VOT2018 datasets. As shown in Table I and II, OP-ABA can reduce the precisions and success rates of KYS, DiMP(ResNet50), and DiMP(ResNet18) on OTB100, VOT2018, UAV123, and LaSOT, significantly. When we compare the attack results of DiMP with those of KYS under the same backbone (*i.e.*, ResNet50), it is easier for OP-ABA to attack DiMP since we achieve much higher precision or success rate drop. Compared DiMP(ResNet50) with DiMP(ResNet18), we see that the DiMP with deeper backbone is harder to be attacked since OP-ABA has lower performance drops on the DiMP(ResNet50), which is consist with the results reported in Table 1 and 2 in the main manuscript.

Table I: Attacking results of OP-ABA against KYS and DiMP with ResNet50 and ResNet18 as backbones on OTB100 and VOT2018, respectively.

Backbones	Trackers	OTB100				VOT2018	
		Org. Prec.	Prec. Drop ↑	Org. Succ.	Succ. Drop ↑	Org. EAO	EAO Drop ↑
ResNet50	KYS	89.5	18.8	68.6	13.8	0.405	0.289
	DiMP	89.2	30.9	68.9	23.6	0.423	0.405
ResNet18	DiMP	87.1	37.3	66.7	27.8	0.351	0.332

Table II: Attacking results of OP-ABA against SiamRPN++ with ResNet50 and MobileNetv2 on UAV123 and LaSOT.

Backbones	Trackers	UAV123				LaSOT			
		Org. Prec.	Prec. Drop. ↑	Org. Succ.	Succ. Drop ↑	Org. Prec.	Prec. Drop ↑	Org. Succ.	Succ. Drop ↑
ResNet50	KYS	82.2	15.4	62.6	11.7	52.7	9.5	55.2	9.5
	DiMP	84.4	32.4	63.9	24.6	54.4	21.1	55.3	18.7
MobNetv2	DiMP	81.0	39.0	61.5	29.9	51.5	25.2	53.1	22.8

2. Visualization of OP-ABA Optimization Process

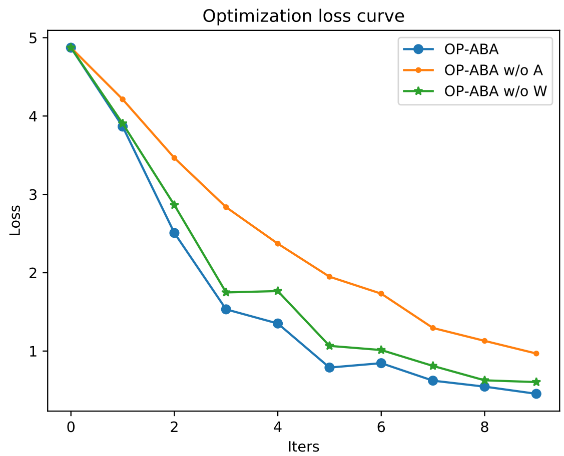

Figure I: The optimization loss during the iteration of OP-ABA w/o $\mathcal{A}$, OP-ABA w/o $\mathcal{W}$, and OP-ABA.

In addition to the ablation study in the Sec 4.3 and Table 4, we further show the loss values of OP-ABA w/o $\mathcal{A}$, OP-ABA w/o $\mathcal{W}$, and OP-ABA during the iterative optimization in Fig. I. Clearly, the loss of OP-ABA considering both $\mathcal{A}$ and $\mathcal{W}$ reduces more quickly than other two variants. When we do not tune the $\mathcal{A}$ (*i.e.*, OP-ABA w/o $\mathcal{A}$), the optimization process of OP-ABA w/o $\mathcal{A}$ becomes less effectively since the loss decreases slowly, demonstrating tunable $\mathcal{A}$ is significantly important for high attack success rate, which is consistent with the conclusion of Sec 4.3 and Table 4 in the main manuscript.

3. Visualization of Attacking Cases

We provide seven cases and compare the tracking results of SiamRPN++(ResNet50) under the original frames, normal blurred frames, and adversarially blurred frames, respectively. Please find the 'gif' files in the zip. Clearly, the adversarially blurred frames have similar appearance with the normal one. The tracker is robust to normal blur but is easily affected by the proposed adversarilly blurred frames.