

Fast Visual Object Tracking using Ellipse Fitting for Rotated Bounding Boxes

Bao Xin Chen John K. Tsotsos

Department of Electrical Engineering and Computer Science, and Centre for Vision Research
York University
Toronto, Canada

{baoxichen, tsotsos}@eecs.yorku.ca

Abstract

In this paper, we demonstrate a novel algorithm that uses **ellipse fitting** to estimate the bounding box rotation angle and size with the segmentation(mask) on the target for online and real-time visual object tracking. Our method, **SiamMask_E**, improves the bounding box fitting procedure of the state-of-the-art object tracking algorithm SiamMask and still retains a fast-tracking frame rate (**80 fps**) on a system equipped with GPU (GeForce GTX 1080 Ti or higher). We tested our approach on the visual object tracking datasets (VOT2016, VOT2018, and VOT2019) that were labeled with rotated bounding boxes. By comparing with the original SiamMask, we achieved an improved Accuracy of **64.5%** and **30.3%** EAO on VOT2019, which is 4.9% and 2% higher than the original SiamMask. The implementation is available on GitHub: https://github.com/baoxinchen/siammask_e.

1. Introduction

Visual object tracking is an important element of many applications such as person-following robots ([6] [5] [31] [18]), self-driving cars ([1] [7] [30] [4]), or surveillance cameras ([9] [23] [41] [39]), etc. The performance of such systems critically depends on a reliable and efficient object tracking algorithm. It is especially important to track an object **online** and in **real-time** when the camera is running under challenging situations: illumination, changing pose, motion blurring, partial and full occlusion, etc. These two fundamental features are the core requirements for human-robot interactions (e.g., person-following robots).

To address the visual object tracking problems, many benchmarks have been developed, such as Object Tracking Benchmark (OTB50 [36] and OTB100 [37]), and Visual Object Tracking Challenges (VOT2016 [21], VOT2018 [19], VOT2019 [20]). In OTB datasets, ground truth was labeled by axis aligned bounding boxes and while




Figure 1. Our approach SiamMask_E yields larger IoU between the ground truth (blue) and its prediction (green) than the original SiamMask (magenta). SiamMask_E predicts a higher accuracy on the orientation of the bounding boxes which improves the average overlap accuracy (A) and expected average overlap (EAO).

in VOT datasets rotated bounding boxes were used. Comparing between axis-aligned bounding boxes and rotated bounding boxes, rotated bounding boxes contain a minimal amount of background pixels [21]. Thus, the datasets with rotated bounding boxes have the tighter enclosed boxes than the axis-aligned bounding boxes. As well as, the rotated bounding boxes provide the object orientation in the image plane. The orientation information can be further used to solve many computer vision problems (e.g., action classification).

Despite the advantage of rotated bounding boxes, it is very computationally intensive to estimate the rotation an-

gle and scale of the bounding boxes. Many researchers have developed novel algorithms to settle the problem. But most of them have limitations in terms of tracking speed or accuracy [17], [33]. In the meantime, fully convolutional Siamese networks [2] had become popular in the field of object tracking. However, the original Siamese networks did not solve the rotation problem. Wang *et al.* (SiamMask) [35] have been inspired by the advanced version of Siamese network (SiamRPN [25], SiamRPN++ [24]) and wide range of image datasets (Youtube-VOS [38], COCO [26], ImageNet [34], etc.). SiamMask is able to predict a segmentation mask on the target for tracking and fits a minimum area rotated bounding box in real-time (87 fps).

In this paper, we propose a novel efficient rotated bounding box estimation algorithm when a segmentation/mask of an object is given. Particularly, the masks are generated by SiamMask. The key problem is to predict the rotation angle of the bounding boxes. Inspired by the conic fitting problem described by Fitzgibbon *et al.* [8], we try to fit an ellipse on the mask to compute the rotation angle. Once the rotation angle is known, then we could fit a rotated rectangle on the mask. Our algorithm consists of two parts: (1) rotation angle estimation, and (2) scale calculation. Details will be provided in Section 3.

The contribution of this paper can be summarized in the following three aspects:

1. a new real-time state-of-the-art **object tracking algorithm** on the datasets that are labeled with rotated bounding boxes, e.g., VOT challenge series (2015-2019) ¹.
2. a fast novel **rotated bounding box estimation algorithm** when a segmentation/mask is given. This algorithm can be used to generate rotated bounding box ground truth from any segmentation datasets to train a rotation angle regression model. This is the main contribution of the paper.
3. the source code ² will be released as an additional package to PySOT ³ which is written by SenseTime Video Intelligence Research team.

The paper is structured as follows. The most relevant work will be briefly summarized in Section 2. Then, we will describe our approach in detail in Section 3. The evaluation of the algorithm is in Section 4. Finally, Section 5 concludes the paper and discusses future work.

2. Related Work

In this section, we discuss the history of the Siamese network based tracking algorithms and several trackers that

yield rotated bounding boxes.

2.1. Siamese network based trackers

The first Siamese network based object tracking algorithm (**SiamFC**) was introduced by Bertinetto *et al.* [2] in 2016. The Siamese network is trained offline on a dataset for object detection in videos. The input to the network are two images, one is an exemplar image z , the other one is the search image x . Then, a dense response map is generated from the output of the network. **SiamFC** learns and predicts the similarity between the regions in x and the exemplar image z . In order to handle the object scale variation, SiamFC searches for objects at five scales $1.025^{\{-2,-1,0,1,2\}}$ near the target's previous location. As a result, there will be 5 forward passes on each frame. SiamFC runs at about 58 fps, which is the fastest fully convolutional network (CNN) based tracker comparing to online training and updating networks in 2016. However, SiamFC is an axis-aligned bounding box tracker. It couldn't outperform the online training and updating deep CNN tracker MDNet [27] (1 fps) in terms of average overlap accuracy.

He *et al.* [14] combines two branches (Semantic net and Appearance net) of Siamese network (**SA-Siam**) to improve the generalization capability of **SiamFC**. Two branches are individually trained, and then the two branches are combined to output the similarity score at testing. S-Net is an AlexNet [22] pretrained on an image classification dataset. A-Net is a SiamFC pretrained on an object detection from video dataset. S-Net improves the discrimination power of the SA-Siam tracker because different objects activate different sets of feature channels in the Semantic branch. Due to the complexity of the two branches, SA-Siam runs at 50 fps when tracking with pretrained model.

By modifying the original Siamese net with a Region Proposal Network(RPN) [32], Li *et al.* [25] proposed a Siamese Region Proposal Network (**SiamRPN**) to estimate the target location with the variable bounding boxes. The output of **SiamRPN** contains a set of anchor boxes with corresponding scores. So, the bounding box with the best score is considered as the target location. The benefit of RPN is to reduce the multi-scale testing complexity in the traditional Siamese networks (SiamFC, SA-Siam). An updated version SiamRPN++ [24] has released in 2019. In terms of processing speed, SiamRPN is 160 fps and SiamRPN++ is about 35 fps.

Unlike **SiamFC**, **SA-Siam**, and **SiamRPN** yielding axis-aligned bounding boxes, **SiamMask** [35] uses the advantage from a video object segmentation dataset and trained a Siamese net to predict a set of masks and bounding boxes on the target. The bounding boxes are estimated based on the masks using rotated minimum bounding rectangle (MBR) at the speed of 87 fps. However, the MBR does not always predict the bounding boxes that perfectly

¹<http://www.votchallenge.net/challenges.html>

²https://github.com/baoxinchen/siammask_e

³<https://github.com/STVIR/pysot>

align with the ground truth bounding boxes (see Figure 1). Although the same bounding boxes prediction algorithm used in VOT2016 for generating the ground truth can improve the average overlap accuracy dramatically, the running speed decreases to 5 fps. To address this problem, we present a new method in Section 3 that can process frames in real-time and achieves a better result.

2.2. Rotated bounding boxes

Beside the Siamese network trackers, Nebehey *et al.* [28] (CMT) use a key-point matching approach to scale and rotate the bounding boxes. But, this tracker cannot handle deformable objects. [29] is an update of CMT, and the processing speed dropped to 11 fps.

Hua *et al.* [17] suggest a proposal selection method (optical flow [3] and Hough transform [16]) to filter out a group of locations and orientations that very likely contains the object. Then, they use three cues (detection confidence, objectness measures from object edges and motion boundaries) to determine which location has the highest likelihood. But, this approach also couldn't run in real-time (0.3 fps).

Zhang *et al.* [40] propose a rotation estimation method using Log-Polar transformation. In Log-Polar coordinate, a set of 36 rotation sample are chosen on every $\Delta = \frac{2\pi}{R}$, where $R = 36$. But, the rotation sample set also increases the run-time of KCF [15] tracker by 36 times.

Guo *et al.* [11] build a structure-regularized compressive tracking (SCT) with online update. During the detection stage, SCT samples several candidates with different rotation angles based on integral image and quadtree segmentation. SCT runs on a computer system without GPU at 15 fps.

Recently, a rotation adaptive tracking approach was introduced by Rout *et al.* [33]. The authors assume that the rotation angle is limited within a range (e.g., $\pm 10^\circ$). However, this assumption doesn't always hold. He *et al.* [13] built on top of SA-Siam [14] with angle estimation strategy. Although the method could reduce the processing time, it still limits the rotation angle to some degrees (e.g., $-\pi/8, \pi/8$). In order to find an arbitrary rotation angle, we present our approach in the next section.

3. Approach

In the original SiamMask [35] tracker, Wang *et al.* compared three different bounding boxes estimation algorithms: min-max axis-aligned rectangle (*Min-max*), minimum area rectangle (*MBR*), and optimal bounding box [21] (*Opt*). Due to the computational burden, *Opt* could not perform in real-time (5fps). SiamMask with *MBR* is the real-time (87 fps) state-of-the-art tracker in terms of average overlap Accuracy. Although *MBR* performs better than the other bounding box estimation algorithms, it has a weakness such




Figure 2. This figure shows some examples of minimum area rectangle (magenta); this does not determine bounding boxes according to the geometric shape and point distribution of the segmentation/mask. Thus, the rotation angles are not as accurate as our approach (green)

that minimum area rectangle could not represent the geometric shape and point distribution of the masks (see Figure 2). As a result, most of the estimated bounding boxes are not in the correct orientation. In the following subsections, we will discuss an alternate solution to generate bounding boxes with correct rotation angle and tighter size by post-processing on the output mask from SiamMask. Our method consists of the steps in Figure 3.

3.1. Rotation angle estimation

To estimate the rotation angle, we adopted the `fitEllipse` API provided by OpenCV3⁴ which use a least-squares scheme [10] to solved the ellipse fitting problem. An improved version was described in [12]. This algorithm (B2AC) Algebraic distance with quadratic constraint was first introduced by Fitzgibbon *et al.* [8].

An ellipse can be formulated using a conic equation with a constraint:

$$F(x, y) = ax^2 + bxy + cy^2 + dx + ey + f = 0 \quad (1)$$

where, $b^2 - 4ac < 0$

In Equation 1, a, b, c, d, e, f are the coefficients of the ellipse and x, y are the points on the ellipse. By grouping the coefficients into a vector, we have the following two vectors:

$$\mathbf{a} = [a, b, c, d, e, f]^T \quad (2)$$

$$\mathbf{x} = [x^2, xy, y^2, x, y, 1]$$

So, the conic can be written as:

$$F(\mathbf{x}) = \mathbf{x} \cdot \mathbf{a} = 0 \quad (3)$$

To fit an ellipse on a set of points $A = \{(x_1, y_1), \dots, (x_N, y_N)\}$, where $|A| = N$, we need to find the coefficient vector \mathbf{a} . Halíř *et al.* [12] introduced an improved least squares method to minimize the sum of

⁴<https://opencv.org/>




Figure 3. Our algorithm includes seven steps: (a) take a target mask as input. (b) apply an ellipse fitting algorithm [8] on edge of the mask (Here, we have the points on the edge as a set A in Equation 4), then determine the center of the ellipse and the rotation angle. (c) compute the affine transformation matrix using the rotation angle and the center from the ellipse, then apply the transformation on the ellipse center. (d) apply a rectangular rotated bounding box (green) on the ellipse. (e) draw a min-max axis-aligned bounding box (blue) on the transformed mask. (f) calculate the intersection of the blue box and green box to form a new bounding box (red). (g) calculate the inverse of the affine transformation matrix, then apply transformation to convert back to the original image coordinate, and output the red box.




Figure 4. Ellipse notations

squared error of the following equation:

$$\min_{\mathbf{a}} \sum_{i=1}^N F(x_i, y_i)^2 = \min_{\mathbf{a}} \sum_{i=1}^N F(\mathbf{x}_i)^2 \quad (4)$$

where, $\mathbf{x}_i = [x_i^2, x_i y_i, y_i^2, x_i, y_i, 1]$, and $A_i = (x_i, y_i)$

Let us denote the following terms for the fitted ellipse (also see Figure 4):

m semi-major axis

n semi-minor axis

(x_o, y_o) center coordinate of the ellipse

θ rotation angle

Be aware that, when the ellipse is near-circular (rotational symmetric shapes), θ is not stable. A solution for this case is to force $\theta = 90^\circ$. However, it did not increase the performance of the VOT datasets empirically.

3.2. SiamMask_E

Since we need to rotate the image with respect to the ellipse center, an affine transformation (Translation and Rotation for our case) will be used here to compute the transformed coordinates. After the estimation of rotation angle θ and the center point (x_o, y_o) , then we need to compute the 2D affine transformation matrix M :

$$M = \begin{bmatrix} \cos\theta & \sin\theta & (1 - \cos\theta)x_o - \sin\theta y_o \\ -\sin\theta & \cos\theta & \sin\theta x_o - (1 - \cos\theta)y_o \end{bmatrix} \quad (5)$$

Once the affine transformation matrix is computed, then we apply the rotation on the segmentation/mask about the ellipse's center (x_o, y_o) : Let's denote the mask as a set of points $Mask$ (magenta color in Figure 3(a)), and the transformed mask as $Mask'$ (magenta color in Figure 3(d)).


$$Mask' = M * \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \quad \forall (x, y) \in Mask \quad (6)$$

After this step, our aim is to output the intersection (red in Figure 3(f)) between the min-max axis-aligned bounding box (blue in Figure 3(e)) and the ellipse bounding box (green in Figure 3(e)). The advantage of using the ellipse bounding box is to cut out the unexpected portion of the shape (e.g., protruding limbs). Thus, the output bounding box would be able to focus on the trunk of the human body. After the affine transformation, the ellipse bounding box is trivial, and we denote it as G :

$$G = [x_o - n, y_o - m, x_o + n, y_o + m] \quad (7)$$

The min-max axis-aligned bounding box denote as B :

$$B = [\min(\forall x \in Mask'), \min(\forall y \in Mask'), \max(\forall x \in Mask'), \max(\forall y \in Mask')] \quad (8)$$



(a) Refinement step (b) Refinement output

Figure 5. Refinement step: (a) move four edges toward the bounding box center if the constraint in Subsection 3.3 is not satisfied. (b) the magenta box is the estimated bounding box from SiamMask_E, and the green box is a sample output after the refinement step. Blue box is the ground truth.

The intersection bounding box R (red in Figure 3(f)) can be calculated using the following equation:

$$R = [\max(G_1, B_1), \max(G_2, B_2), \min(G_3, B_3), \min(G_4, B_4)] \quad (9)$$

then, convert R to a polygon

$$R = [[R_1, R_2], [R_3, R_2], [R_3, R_4], [R_1, R_4]] \quad (10)$$

The last step is to convert the transformed coordinate back to the image coordinate using the inverse of the affine transformation matrix M . We denote the output bounding box as R' (red color in Figure 3(g)):

$$R' = M^{-1} * \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}; \forall (x, y) \in R \quad (11)$$

3.3. Refinement step (Ref)

As you can see from Figure 1 at row 3 column 2, our bounding box (green) is not as tight as the ground truth (blue). This problem because the *Mask* generated by SiamMask includes the limbs of the dancer. To manage this problem, we implement a refinement procedure to slim the size of the bounding box by evaluating the amount of *Mask* that an edge is crossing. Let's denote the length of an edge as α , and the portion of the edge intersecting the *Mask* is β . We set a constraint such that:

$$\beta > \alpha * 0.2 \quad (12)$$

otherwise, the edge will gradually move toward the bounding box center (see Figure 5). This constraint is empirically proved in Section 4.

4. Experiments

In this section, we evaluate our proposed methods on the datasets that labeled with rotated bounding boxes: VOT2016, VOT2018, and VOT2019.

4.1. Environment setup

In order to provide a fair comparison, we test our algorithm using the same pretrained Siamese network model and the same parameters in [35]. The reported data is evaluated on a desktop computer with the following hardware:

- GPU: GeForce GTX 1080 Ti
- CPU: Intel Core i5-8400 CPU @ 2.80GHz \times 6
- Memory: 32 GB

4.2. Evaluation metrics

We only evaluation on the VOT challenge series (VOT2015-2019 short term), where VOT2015 has the same data sequences as VOT2016, and VOT2017 has the same sequences as VOT2018. These three datasets contain 60 sequences with different challenging situations (e.g., motion blur, size change, occlusion, illumination change, etc). To the best of our knowledge, VOT2015-2019 are the only object tracking datasets that labeled with rotated bounding boxes. We also adopt the supervised tracking evaluation methods that are used in VOT2016 [21]: Accuracy (A), Robustness (R), and Expected Average Overlap (EAO). The Accuracy is the average overlap between the estimated and the ground truth bounding boxes when the target is successfully being tracked. The Robustness measures the ratio between the number of times the tracker loses the target (fails) and the number of resumed trackings. The Expected Average Overlap (EAO) is considered as the primary measurement in the VOT challenge. According to the official toolkit, the tracker will be reinitialized when the estimated bounding box has no intersection with the ground truth. After five frames, the tracker will restart with the ground truth bounding box.

4.3. Overall results

Table 1 presents the result comparison between the state-of-the-art Siamese based tracking algorithms on VOT2016, VOT2018, and VOT2019 datasets. Our tracker SiamMask_E with Ref has the 0.648 Accuracy and 0.432 EAO on VOT2018 dataset which it a new state-of-the-art comparing the other Siamese trackers and the VOT2018 short term challenge winners [19]. Although SiamMask-Opt has the similar performance as ours, due to the computation complexity, SiamMask-Opt can only run at 5 frames per second. However, our tracker is able to process in real-time with the speed of more than 80 frames per second.

	VOT2019			VOT2018			VOT2016			Speed \uparrow
	A \uparrow	R \downarrow	EAO \uparrow	A \uparrow	R \downarrow	EAO \uparrow	A \uparrow	R \downarrow	EAO \uparrow	
SiamRPN++	0.595	0.467	0.287	0.601	0.234	0.415	0.642	0.196	0.464	46 fps
SiamMask	0.596	0.467	0.283	0.598	0.248	0.406	0.621	0.214	0.436	87 fps
SiamMask-Opt*	-	-	-	0.642	0.295	0.387	0.670	0.233	0.442	5 fps
SiamMask_E (Ours)	0.625	0.487	0.296	0.624	0.253	0.419	0.642	0.219	0.445	85 fps
SiamMask_E_Ref (Ours)	0.645	0.497	0.303	0.648	0.267	0.432	0.668	0.233	0.451	80 fps

Table 1. Comparing with the state-of-the-art Siamese trackers on VOT2019, VOT2018, and VOT2016. Our tracker SiamMask_E with Ref outperforms other trackers in terms of average overlap accuracy (A) and expected average overlap (EAO). \uparrow stands for the higher the best, and \downarrow stands for the lower the best. * the numbers are reported in the original paper.

Factor	A \uparrow	R \downarrow	EAO \uparrow
0.1	0.635	0.262	0.424
0.15	0.641	0.267	0.427
0.2	0.648	0.267	0.432
0.25	0.654	0.276	0.425
0.3	0.656	0.281	0.420
0.35	0.657	0.290	0.417
0.4	0.650	0.290	0.413

Table 2. Comparing seven different factors on VOT2018 for our SiamMask_E with refinement step (Ref) (Subsection 3.3). According to the primary measurement (EAO), we choose factor 0.2.

Similarly, our tracker also forms a new state-of-the-art result on VOT2019.

4.4. Comparing different factors for Ref

We tested our refinement step (Ref) (Subsection 3.3) with seven different values (see Table 2) on VOT2018. Among these seven results, Ref with factor 0.2 outperforms the other factors with the highest EAO (0.432) and a decent Accuracy (0.648). Although Ref with factor 0.1 has the best Robustness (0.262), Ref with factor 0.2 is only 1.9% higher. Ref with factor 0.35 reaches the highest Accuracy 0.657, but its Robustness becomes the worst (0.290). As a result, factor 0.2 is selected for the constraint in Subsection 3.3.

4.5. Ablation studies

The ablation test results are shown in Table 3. In the table, SiamMask_E is our baseline model without the refinement step. We exchange the bounding box orientation between SiamMask_E and SiamMask, where SiamMask_E with Minimum Area Bounding Box angle (SiamMask_E + minABBoxAngle) performs weaker than our baseline SiamMask_E. Similarly, SiamMask with ellipse angle (SiamMask + ellipseAngle) is preferable over the original SiamMask. By adding the refinement step (Ref) to both SiamMask and SiamMask_E, the average overlap Accuracy increase dramatically. Furthermore, we modify the bounding box rotation of SiamMask_E + Ref with the original Minimum Area Bounding Box angle (SiamMask_E +

Ref + minABBoxAngle) which results in slightly decreasing on the primary measurement EAO. It proves that using the ellipse’s angle could improve the tracking performance on the VOT datasets. On the other hand, we also test SiamMask + Ref with changing the angle of the Minimum Area Bounding Box to the ellipse’s angle (SiamMask + ellipseAngle + Ref). The result shows that SiamMask + ellipseAngle + Ref also has some degree of improvement on both VOT2018 and VOT2019 on the primary measurement EAO. Overall, SiamMask_E, which improves the bounding box orientation and scale using ellipse fitting on top of SiamMask, has a similar performance as the original SiamMask with the refinement step (SiamMask + Ref). And, SiamMask_E with the refinement step (SiamMask_E + Ref) outperforms any other combinations on the ablation study table.

4.6. Qualitative results

To analysis the improvement, we show several results computed on VOT2019 [20] dataset. We compare the state-of-the-art algorithms SiamMask [35] and SiamRPN++ [24] along with our approach SiamMask_E in Figure 6.

5. Conclusion

In this paper, we updated the **SiamMask** tracker to achieve the next level of state-of-the-art performance. Our new tracker **SiamMask_E** retains real-time processing speed at 80 fps. We show that the bounding box using ellipse fitting outperforms the minimum area rectangle bounding box in terms of better rotation angle and tighter bounding box scale. Our results show the strength of SiamMask network tracking model such that it can outperform the other state-of-the-art trackers.

Future work: Our approach focused on an efficient bounding box refinement algorithm. On a different aspect, if a proper motion model is employed, we believe the result could move to the next level. To attain this, a real-time algorithm is needed to differentiate the camera the target motion in order to estimate the real target motion. As well as, we need to beware the other dynamic distractors in the scene.

	VOT2019			VOT2018		
	A \uparrow	R \downarrow	EAO \uparrow	A \uparrow	R \downarrow	EAO \uparrow
SiamMask_E (Ours)	0.625	0.487	0.296	0.624	0.253	0.419
SiamMask	0.596	0.467	0.283	0.598	0.248	0.406
SiamMask_E + minABoxAngle	0.618	0.472	0.292	0.621	0.253	0.418
SiamMask + ellipseAngle	0.594	0.477	0.284	0.595	0.243	0.409
SiamMask_E + Ref (Ours)	0.645	0.497	0.303	0.648	0.267	0.432
SiamMask + Ref	0.644	0.497	0.297	0.645	0.272	0.416
SiamMask_E + minABoxAngle + Ref	0.647	0.487	0.302	0.649	0.262	0.428
SiamMask + ellipseAngle + Ref	0.639	0.502	0.299	0.643	0.267	0.422

Table 3. Ablation studies: SiamMask_E is our baseline tracker with the ellipse angle and ellipse box, and SiamMask is the original tracker with the minimum area bounding box. Ref stands for the refinement step in Subsection 3.3. minABoxAngle stands for the orientation of the minimum area bounding box. ellipseAngle stands for the orientation of the best fitting ellipse. The result shows that the effectiveness of ellipse orientation and refinement step significantly improve the performance of SiamMask.

Acknowledgement

We acknowledge the financial support of the Natural Sciences and Engineering Research Council of Canada (NSERC), the NSERC Canadian Robotics Network (NCRN), and the Canada Research Chairs Program through grants to John K. Tsotsos. The first author also would like to thank Dekun Wu provided hardware support for testing.

References

- [1] N. Agarwal, C.-W. Chiang, and A. Sharma. A study on computer vision techniques for self-driving cars. In *International Conference on Frontier Computing*, pages 629–634. Springer, 2018.
- [2] L. Bertinetto, J. Valmadre, J. F. Henriques, A. Vedaldi, and P. H. Torr. Fully-convolutional siamese networks for object tracking. In *European conference on computer vision*, pages 850–865. Springer, 2016.
- [3] T. Brox and J. Malik. Large displacement optical flow: descriptor matching in variational motion estimation. *IEEE transactions on pattern analysis and machine intelligence*, 33(3):500–513, 2010.
- [4] A. Buyval, A. Gabdullin, R. Mustafin, and I. Shimchik. Real-time vehicle and pedestrian tracking for didi udacity self-driving car challenge. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2064–2069. IEEE, 2018.
- [5] B. X. Chen, R. Sahdev, and J. K. Tsotsos. Integrating stereo vision with a cnn tracker for a person-following robot. In *International Conference on Computer Vision Systems*, pages 300–313. Springer, 2017.
- [6] B. X. Chen, R. Sahdev, and J. K. Tsotsos. Person following robot using selected online ada-boosting with stereo camera. In *Computer and Robot Vision (CRV), 2017 14th Conference on*, pages 48–55. IEEE, 2017.
- [7] H. Cho, Y.-W. Seo, B. V. Kumar, and R. R. Rajkumar. A multi-sensor fusion system for moving object detection and tracking in urban driving environments. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1836–1843. IEEE, 2014.
- [8] A. W. Fitzgibbon, R. B. Fisher, et al. *A buyer’s guide to conic fitting*. University of Edinburgh, Department of Artificial Intelligence, 1996.
- [9] V. Gajjar, A. Gurnani, and Y. Khandhediya. Human detection and tracking for video surveillance: A cognitive science approach. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2805–2809, 2017.
- [10] W. Gander. Least squares with a quadratic constraint. *Numerische Mathematik*, 36(3):291–307, 1980.
- [11] Q. Guo, W. Feng, C. Zhou, C.-M. Pun, and B. Wu. Structure-regularized compressive tracking with online data-driven sampling. *IEEE Transactions on Image Processing*, 26(12):5692–5705, 2017.
- [12] R. Halir and J. Flusser. Numerically stable direct least squares fitting of ellipses. In *Proc. 6th International Conference in Central Europe on Computer Graphics and Visualization. WSCG*, volume 98, pages 125–132. Citeseer, 1998.
- [13] A. He, C. Luo, X. Tian, and W. Zeng. Towards a better match in siamese network based visual object tracker. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018.
- [14] A. He, C. Luo, X. Tian, and W. Zeng. A twofold siamese network for real-time object tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4834–4843, 2018.
- [15] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista. High-speed tracking with kernelized correlation filters. *IEEE transactions on pattern analysis and machine intelligence*, 37(3):583–596, 2014.
- [16] P. V. Hough. Method and means for recognizing complex patterns, Dec. 18 1962. US Patent 3,069,654.
- [17] Y. Hua, K. Alahari, and C. Schmid. Online object tracking with proposal selection. In *Proceedings of the IEEE international conference on computer vision*, pages 3092–3100, 2015.
- [18] K. Koide and J. Miura. Convolutional channel features-based person identification for person following robots. In *International Conference on Robotics and Automation (ICRA)*, pages 1836–1843. IEEE, 2014.



Figure 6. Qualitative results: We show some sample outputs on eight sequences selected from VOT2019 [20], where the red box is SiamMask [35], the cyan box is SiamRPN++ [24], the green box is SiamMask_E(ours), and the blue box is the ground truth.

tional Conference on Intelligent Autonomous Systems, pages 186–198. Springer, 2018.

- [19] M. Kristan, A. Leonardis, J. Matas, M. Felsberg, R. Pflugfelder, L. Čehovin Zajc, T. Vojir, G. Häger, A. Lukežič, A. Eldesokey, G. Fernandez, and et al. The sixth visual object tracking vot2018 challenge results, 2018.
- [20] M. Kristan, A. Leonardis, J. Matas, M. Felsberg, R. Pflugfelder, L. Čehovin Zajc, T. Vojir, G. Häger, A. Lukežič, A. Eldesokey, G. Fernandez, and et al. The seventh visual object tracking vot2019 challenge results, 2019.
- [21] M. Kristan, A. Leonardis, J. Matas, M. Felsberg, R. Pflugfelder, L. Čehovin Zajc, T. Vojir, G. Häger, A. Lukežič, and G. Fernandez. The visual object tracking vot2016 challenge results. Springer, Oct 2016.
- [22] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In

- Advances in neural information processing systems*, pages 1097–1105, 2012.
- [23] Y.-G. Lee, Z. Tang, and J.-N. Hwang. Online-learning-based human tracking across non-overlapping cameras. *IEEE Transactions on Circuits and Systems for Video Technology*, 28(10):2870–2883, 2017.
- [24] B. Li, W. Wu, Q. Wang, F. Zhang, J. Xing, and J. Yan. Siamrpn++: Evolution of siamese visual tracking with very deep networks. *arXiv preprint arXiv:1812.11703*, 2018.
- [25] B. Li, J. Yan, W. Wu, Z. Zhu, and X. Hu. High performance visual tracking with siamese region proposal network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8971–8980, 2018.
- [26] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [27] H. Nam and B. Han. Learning multi-domain convolutional neural networks for visual tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4293–4302, 2016.
- [28] G. Nebel and R. Pflugfelder. Consensus-based matching and tracking of keypoints for object tracking. In *IEEE Winter Conference on Applications of Computer Vision*, pages 862–869. IEEE, 2014.
- [29] G. Nebel and R. Pflugfelder. Clustering of static-adaptive correspondences for deformable object tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2784–2791, 2015.
- [30] A. Petrovskaya and S. Thrun. Model based vehicle detection and tracking for autonomous urban driving. *Autonomous Robots*, 26(2-3):123–139, 2009.
- [31] Q. Ren, Q. Zhao, H. Qi, and L. Li. Real-time target tracking system for person-following robot. In *2016 35th Chinese Control Conference (CCC)*, pages 6160–6165. IEEE, 2016.
- [32] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [33] L. Rout, D. Mishra, R. K. S. S. Gorthi, et al. Rotation adaptive visual object tracking with motion consistency. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1047–1055. IEEE, 2018.
- [34] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [35] Q. Wang, L. Zhang, L. Bertinetto, W. Hu, and P. H. Torr. Fast online object tracking and segmentation: A unifying approach. *arXiv preprint arXiv:1812.05050*, 2018.
- [36] Y. Wu, J. Lim, and M.-H. Yang. Online object tracking: A benchmark. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [37] Y. Wu, J. Lim, and M.-H. Yang. Object tracking benchmark. In *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, volume 37, pages 1834–1848, 2015.
- [38] N. Xu, L. Yang, Y. Fan, J. Yang, D. Yue, Y. Liang, B. Price, S. Cohen, and T. Huang. Youtube-vos: Sequence-to-sequence video object segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 585–601, 2018.
- [39] R. Xu, S. Y. Nikouei, Y. Chen, A. Polunchenko, S. Song, C. Deng, and T. R. Faughnan. Real-time human objects tracking for smart surveillance at the edge. In *2018 IEEE International Conference on Communications (ICC)*, pages 1–6. IEEE, 2018.
- [40] M. Zhang, J. Xing, J. Gao, X. Shi, Q. Wang, and W. Hu. Joint scale-spatial correlation tracking with adaptive rotation estimation. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 32–40, 2015.
- [41] Y. Zhou, S. Zlatanova, Z. Wang, Y. Zhang, and L. Liu. Moving human path tracking based on video surveillance in 3d indoor scenarios. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 3:97, 2016.