

FMODetect: Robust Detection of Fast Moving Objects

Denys Rozumnyi^{1,4} Jiří Matas⁴ Filip Šroubek² Marc Pollefeys^{1,3} Martin R. Oswald¹

¹Department of Computer Science, ETH Zurich

³Microsoft Mixed Reality and AI Zurich Lab

²UTIA, Czech Academy of Sciences

⁴Visual Recognition Group, Czech Technical University in Prague

Abstract

We propose the first learning-based approach for fast moving objects detection. Such objects are highly blurred and move over large distances within one video frame. Fast moving objects are associated with a deblurring and matting problem, also called deblatting. We show that the separation of deblatting into consecutive matting and deblurring allows achieving real-time performance, i.e. an order of magnitude speed-up, and thus enabling new classes of application. The proposed method detects fast moving objects as a truncated distance function to the trajectory by learning from synthetic data. For the sharp appearance estimation and accurate trajectory estimation, we propose a matting and fitting network that estimates the blurred appearance without background, followed by an energy minimization based deblurring. The state-of-the-art methods are outperformed in terms of recall, precision, trajectory estimation, and sharp appearance reconstruction. Compared to other methods, such as deblatting, the inference is of several orders of magnitude faster and allows applications such as real-time fast moving object detection and retrieval in large video collections.

1. Introduction

Fast moving objects (FMOs) are objects that look significantly blurred in images. Within the aperture of a single frame, they move at high speed over distances larger than their size. FMOs typically appear as blurred streaks formed as a composition of the background and their blurred appearance. Other scenarios include fast rotations of objects like ventilators or propellers. FMOs are common in everyday scenarios like fast cars, falling objects, flying insects, rain, and hailstorm. They are mainly present in sports videos such as tennis, football, badminton, or other games with moving objects. Applications of FMO detection methods are temporal super-resolution, speed measurement, FMO retrieval, tracking, removal, and highlighting.




Figure 1. Examples of automatically detected fast moving objects (FMOs) by the proposed method on soccer (top) and volleyball (bottom) videos.

FMOs were introduced in [24], where the authors created the first dataset containing fast moving objects and proposed a proof-of-concept algorithm to detect and track such objects. The problem of detecting and tracking FMOs has been unnoticed by the research community, and such objects are not present in standard tracking datasets such as VOT [15], OTB [28], and ALOV [26]. More precisely, they may or may not be present, but FMO is never the object intended for tracking by the ground truth annotation. We highlight this by running FMO retrieval by the proposed method and detect FMOs in very recent standard large datasets like GOT-10k [8] and LaSOT [5] and find unannotated FMOs (Fig. 1, bottom). By processing the ground truth bounding boxes in more than 5M frames from both datasets, we noticed that there is a significant overlap between consecutive bounding boxes in almost all cases, and the speed is close to zero (Fig. 3). In contrast, the recently proposed TbD dataset [13] with fast moving objects has almost zero intersection over union between consecutive objects, and the speed is higher.

Apart from [24], there are a few methods that deal with FMOs. Tracking by Deblatting (TbD) [13] jointly solves an inverse problem of *deblatting* (*deblurring* and *matting*) and decouples the input images into trajectory, appearance, and



Figure 2. Examples of fast moving object (FMO) retrieval on YouTube videos by our method running in real-time. **Top:** frames from a 10-minute surveillance footage [youtu.be/CkVJyAKwByw]. **Bottom:** frames from a 5-minute trick-shot video [youtu.be/A2FsgKoGD04]. For each image, the detected FMO and its trajectory is shown as a red line in a close-up below. The object description contains a digitally clickable link to the timestamp when the event happened, in [minutes:seconds] format.

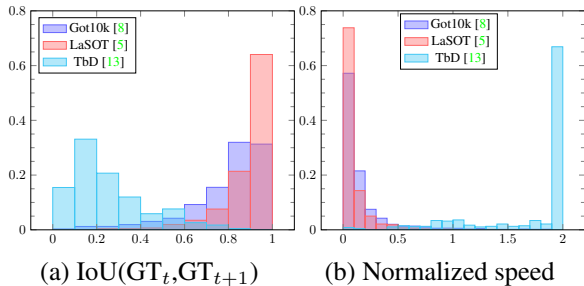


Figure 3. Histograms of (a) Intersection over Union (IoU) scores between consecutive ground truth bounding boxes and (b) speed normalized by the object size for datasets GOT-10k [8], LaSOT [5], and TbD [13] that focuses on FMOs. FMOs are present in GOT-10k and LaSOT datasets, but never annotated, *i.e.* trackers are not tested on fast moving objects.

mask. All previous methods are based on classical energy minimization. While powerful, they still require many parameter tuning and are slow. In most cases, they have to be tuned for different sequences separately. Also, their performance is not entirely satisfactory, and they force several limiting assumptions, *e.g.* static camera, high contrast between the object and the background, and spherical shape.

This paper proposes a novel method for FMO detection and trajectory and appearance estimation using a combination of deep learning and energy minimization based deblurring. Instead of solving an ill-posed and complex deblat-

ting [13] problem, we first solve the matting problem that separates the background from the input image by learning from synthetic data. Then, an easier deblurring problem is solved by energy minimization. The proposed FMODetect is real-time capable and achieves state-of-the-art performance, which allows us to tackle challenging applications such as FMO retrieval and tracking in long videos (Fig. 2).

In summary, we make the following *contributions*:

- We present the first learned approach to the problem of FMO detection. Compared to the previous methods, the proposed one is simpler, does not require extensive parameter tuning, and works with the same settings for a wide range of scenarios, *e.g.* objects of different complexities, moving among various trajectories, and captured by moving or zooming cameras. The method is trained on a new synthetic dataset and generalizes well to unseen and more difficult real-world data.
- By separating deblatting into matting and deblurring, computational and problem complexities are simplified, and the method achieves state-of-the-art results. Our method is an order of magnitude faster than the previous methods, and the detection part is real-time capable. This makes applications such as robust FMO retrieval from large video collections realistic.
- For the sharp appearance estimation, we propose a novel energy minimization-based deblurring.

2. Related work

The pioneering work in [24] introduced FMOs but was limited by many assumptions such as a linear trajectory, high contrast, and no occlusions. Some of them have been addressed in the follow-up method called Tracking by Deblatting (TbD) [13]. TbD solves an inverse problem of *deblurring* and *matting* by alternating iterative minimization, followed by fitting piecewise parabolic curves. As a by-product, TbD estimates sharp 2D object appearance, which is used for long-term tracking of fast moving objects. Both methods introduced datasets with fast moving objects. The ground truth in the FMO dataset [24] was only a binary mask with rough locations of FMOs in each frame. The TbD dataset [13] provides more ground truth information from a high-speed camera, such as object templates and sub-frame object trajectories.

The above-mentioned methods assumed causal processing of video frames, which means that the trajectory at the current frame is estimated using only information from previous frames. However, some scenarios involving fast moving objects do not require causal processing, and some latency is acceptable. Non-causal Tracking by Deblatting (TbD-NC) [23] is the first method to exploit this fact. TbD-NC takes detections in each frame from a causal method as input and produces a single trajectory for the whole sequence. Then, the final trajectory is a continuous, piecewise polynomial curve. TbD-NC improves the accuracy and recall of trajectory estimation of fast moving objects. Also, it allows for accurate speed estimation, which is on par with expensive radar guns used in professional tennis matches. Applications such as temporal super-resolution become significantly better with non-causal estimation.

The method proposed in this paper is causal, and computations are done for each frame independently. Therefore, the method in TbD-NC can be applied to improve the trajectories estimated by the proposed method.

The authors in [24] introduced fast moving objects together with the corresponding image formation model

$$I = H * F + (1 - H * M)B, \quad (1)$$

where the image is formed as a mixture of blurred object appearance and the background. Motion blur is modeled by the convolution of the object trajectory, represented by the blur kernel H , and the sharp object appearance F . The influence of background depends on the indicator function M of appearance F . For more details, the reader is referred to [24].

Many deblurring methods have been proposed. They all assume small motions like in [10, 18, 21]. The method in [20] considers larger motions but only in a trivial case. We tried several of these methods, and none was able to deblur or reconstruct fast moving objects. Special deblurring

of fast moving objects is also proposed in [12, 14, 27], but only a simple case is studied. The only competitive method is TbD [13]. TbD-3D [25] is an extension towards 3D objects with 3D rotations, which can be directly applied to the output of the proposed method.

Network architectures similar to the proposed one have been used in other methods. Hou *et al.* [7] proposed a two-encoder-two-decoder architecture but matting of sharp images is considered, as well as in [29, 32]. The architecture in [3] is similar, but the task is inverse – to synthesize the motion blur given sharp images. Similar U-Net [22] architectures have been applied in many image-to-image translation tasks [6, 9].

3. Method

The proposed method detects all fast moving objects in an image, recovers their trajectory, and removes the influence of the background, *i.e.* matting. Then an optional step is applied to get sharp object appearance, based on energy minimization based deblurring. An overview of the architecture is shown in Fig. 4.

First, an FMO detection network proposes potential regions with fast moving objects using a truncated distance function (TDF) to represent the trajectory. Second, for each region found by the detection network, the matting and fitting network estimates the parametric curve describing FMO trajectory and performs matting – removing the background from the input image and leaving only the blurred object appearance $H * F$ and $H * M$ from image formation model (1). This representation can be seen as an FMO moving in front of a completely black background. Estimating just $H * F$ is not sufficient since the shape of the object should be estimated besides F , which models only the surface color. Consider an object of black color; here $H * M$, unlike $H * F$, is non-zero. The final step, deblurring, is done by energy minimization similar to deblatting [13]. However, by first removing the blended background, it is easier and faster.

3.1. Detection network

The inputs to the detection network are the input image I and the background B , estimated as a median of the last three frames in a video sequence. The median image approximation worked best in our experiments. Changing the number of frames for the background calculation is essentially a trade-off between sensitivity to camera motion (lower number is better) and noise reduction (higher is better). Our experiments showed that using 3 frames was the most robust setting for camera-motion which provided a sufficiently clean background.

The output is the truncated distance function D , measuring the distance to the object trajectory. If we assume

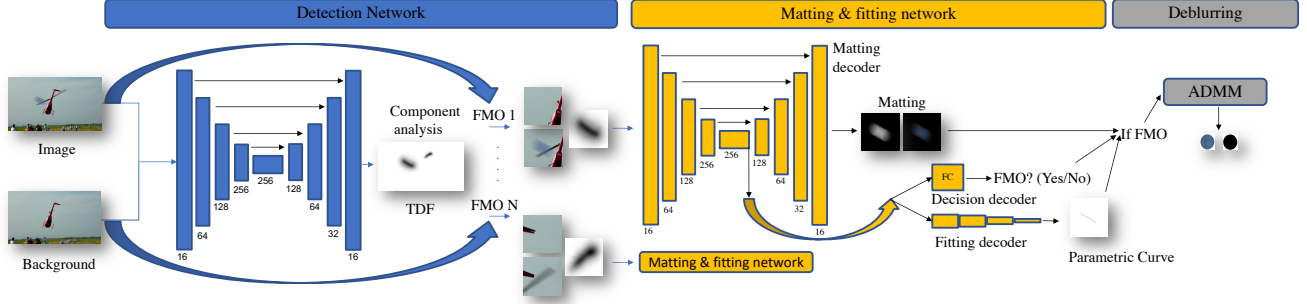


Figure 4. The FMO detection network recovers all fast moving objects (FMOs) in the input image. Then, the matting and fitting network is applied on each detected FMO. The outputs, blurred object without background and its trajectory, are then further used in the proposed deblurring (6) to establish sharp object appearance and refine the trajectory.

that pixels are in some domain $\mathbb{X} \subset \mathbb{R}^2$, then for each pixel $x \in \mathbb{X}$, D is defined as

$$D(x) = 1 - \min \left(1, \min_t \frac{\|x - C(t)\|_2}{2r} \right), \quad (2)$$

where $C(t) : [0, 1] \rightarrow \mathbb{R}^2$ is the ground truth continuous trajectory, and r is the ground truth object radius.

The ground truth for the training sequences is prepared as follows. First, a distance function to the trajectory is computed. Then, it is divided by the object diameter – the maximum distance between any two points on the object. Note that pixels influenced by the object are located at half diameter distance to the trajectory. However, the network converges much faster when a larger neighborhood is considered. Experiments confirmed that it is beneficial to invert the TDF output, *i.e.* 0 means maximum distance (the truncation level), and 1 means ‘on the trajectory.’ The network benefits from the fact that the output is expected to be zero at most locations, which is easier to predict.

The architecture is similar to U-Net [22], and it is fully convolutional. The encoder contains four blocks as in [3], each followed by a max-pooling layer. Each block is built from three convolution layers with filter size 3×3 , followed by a Leaky ReLU [17] with slope 0.1. The decoder part also contains four blocks, each followed by a transposed convolution layer [30, 31]. Inputs and outputs are rescaled to the resolution of 256×512 pixels, with zero mean and unit variance. The full-resolution images are used for testing.

The loss function for the detection network is defined as

$$\mathcal{L}_d = \frac{1}{N_1} \sum_{D(x)>0} \|D(x) - \hat{D}(x)\|_1 + \frac{1}{N_0} \sum_{D(x)=0} \|\hat{D}(x)\|_1, \quad (3)$$

where D corresponds to the ground truth TDF, and \hat{D} to the estimated. N_1 is the number of positive pixels in the ground truth TDF, and N_0 is the number of zero value pixels where a fast moving object is not present in the neighborhood.

3.2. Matting and fitting network

The output D of the detection network is binarized by the threshold value ϵ , and all connected components are extracted. For each connected component, we take crops from the input image, the background, and the non-binarized TDF. They are all scaled to the same resolution of 256×256 pixels and are inputs to the network. The matting and fitting network output is three-fold: blurred object appearance without the influence of the background (matting), parametric trajectory, and a binary decision whether the crop corresponds to an FMO. The last binary output is used to correct the mistakes of the detection network. Experimentally, we have confirmed that the binary predictor removes $\approx 70\%$ of false positives (measured on the TbD dataset [13]).

The matting and fitting network has a one-encoder-three-decoder architecture. The architecture is similar to the detection network. The transposed convolution operators are replaced by upsampling operators to avoid checkerboard artifact [19]. The output of the matting decoder is a 4-channel image ($H * F, H * M$) since F is an RGB image and M is a gray-scale image. The fitting decoder is a multi-layer fully connected network with the final output dimension of 8 channels: the degree of freedom of the estimated parametric trajectory. The last decoder contains only one fully connected layer attached to the latent space and is followed by a sigmoid activation with a single output.

We assume that the trajectory within a frame is one of the following: a line, a parabola, a piece-wise line with one bounce. Then, a common parametric representation of these three cases is

$$\mathcal{C}(t) = c_0 + c_1 \min(2t, 1) + c_2 \min(2t, 1)^2 + c_3 \max(2t - 1, 0), \quad \text{s.t. } 0 \leq t \leq 1, \quad (4)$$

where parameters $c_k \in \mathbb{R}^2$. In case of a line, parameters c_2 and c_3 are equal to zero vectors. In case of a parabola, only c_3 is zero, and for a piece-wise line c_2 is equal to zero. This parametrization can also represent a parabolic curve with a linear tail. However, to avoid this case, whenever $\|c_3\| > 1$, we set c_2 to zero, otherwise c_3 is set to zero.

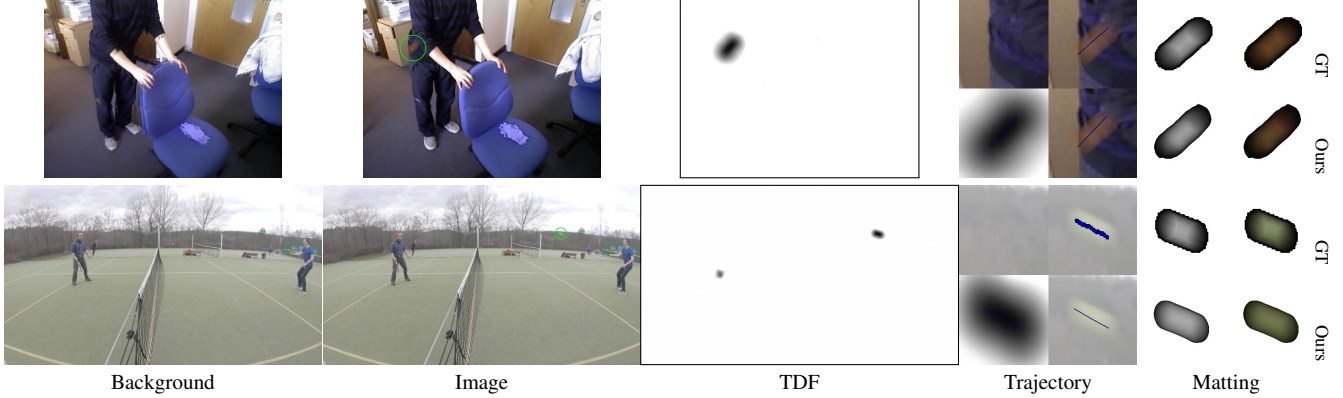


Figure 5. Step by step data processing by our method. We include frames from the synthetic dataset (top row) and real-world TbD dataset [13] (bottom row). *Truncated distance function* (TDF) is an intermediate output of our method for a detected trajectory. Trajectory reconstruction is plotted on the crops. The detected object is marked by a green circle. Examples of deblurring after matting are in Fig. 6.

The loss function of the matting and fitting network is

$$\mathcal{L} = \alpha_a b (\|H * F - \hat{H}_F\|_1 + \|H * M - \hat{H}_M\|_1) + \alpha_c b \mathcal{L}_c(\mathcal{C}, \hat{\mathcal{C}}) + \alpha_b \text{BCE}(b, \hat{b}), \quad (5)$$

which combines three terms. The first one corresponds to blurred appearance reconstruction represented by the ground-truth pair $(H * F, H * M)$. Estimated blurred appearance without background is denoted by (\hat{H}_F, \hat{H}_M) . The second term penalizes deviations from the ground truth trajectory by the curve loss \mathcal{L}_c that samples the trajectory at three points and computes an average distance between each pair. Since the direction is not known, we compute both options and choose the one with the minimal distance. The last term computes binary cross entropy (BCE) between the estimated and the ground truth binary variable b that indicates whether the crop contains an FMO. All previous terms are multiplied by the ground truth binary variable to allow the network to output the best guess of trajectory and appearance even when there is no fast moving object, penalizing only for the incorrectly estimated binary variable b .

3.3. Deblurring

We formulate deblurring as an optimization problem. The output $\hat{\mathcal{C}}$ of the fitting decoder is used for initialization. We also include optimization with respect to H in deblurring to account for small inaccuracies in the blur kernel. The object appearance and mask are then estimated by solving

$$\min_{F, M, H} \frac{1}{2} \left(\|H * F - \hat{H}_F\|_2^2 + \|H * M - \hat{H}_M\|_2^2 \right) + \alpha_F \|\nabla F\|_1 + \alpha_M \|\nabla M\|_1 \quad (6)$$

s.t. $0 \leq F \leq M \leq 1$, $H \geq 0$ and $\sum_i H_i = 1$. The deblurring inputs are \hat{H}_F and \hat{H}_M , which are the outputs of the matting decoder. All constraints are convex sets and we solve the problem using the alternating direction method of

multipliers [2]. The variables are split into blur kernel H and appearance model (F, M) , and both are updated iteratively. In case a parametric curve from H is needed, e.g. for evaluation, we run curve fitting from [13].

The deblurring problem (6) is much easier to solve than the deblatting problem [13]. Deblurring provides more accurate results, as illustrated in several examples in Fig. 6. It is twice as fast as deblatting, and if minimization with respect to H is omitted from (6), it is 10 times faster.

Among other limitations of deblatting, which are solved by the proposed approach, are also failures to separate the background or a shadow correctly (Fig. 6, golf) or to deal with low contrast objects (Fig. 6, softball).

3.4. Training

The standard datasets with FMOs are the TbD [13] and FMO [24] datasets, with 12 and 16 sequences respectively. Ground truth trajectories are estimated by running a state-of-the-art tracking algorithm on high-speed camera footage where the blur is not present. Since the ground truth of matting is not available and is not realistic to get, we created a new synthetic dataset with fast moving objects, called VOT-FMO. The VOT-FMO dataset is based on the VOT dataset [16] – a standard dataset for tracking. As previously shown in [24] and Fig. 1, standard tracking datasets do not label FMOs, even though they still may be present there. Thus, the synthetic dataset is created by applying the FMO formation model (1) with backgrounds taken from the VOT dataset and with artificially inserted FMOs. Artificial FMOs are random discs of up to 100 pixels in radius with textures generated by GeoPatterns¹. Training with objects of various sizes results in good scale invariance. We generate 5000 random frames for training and 500 for validation.

Weights for the loss are set empirically: $\alpha_a = 15$, $\alpha_b = 0.4$. Parameter α_c is set to $\frac{4}{256}$, which normalizes the

¹<https://github.com/jasonlong/geo-pattern>




Figure 6. Matting and deblurring of fast moving objects in the TbD dataset [13] using the proposed method, compared to deblurring [13]. From top to bottom: (a) input image and background by median, (b) left: estimated trajectory plotted on the input image, right: estimated truncated distance function (TDF), (c) output of the matting decoder, (d) ground truth (GT), (e) the proposed deblurring, (f) deblurring [13].

curve loss to produce one unit when the average distance between trajectories is $\frac{1}{4}$ of the crop width of 256 pixels. Regularization weights for deblurring were experimentally set to $\alpha_F = 0.001$ and $\alpha_M = 0.05$. We found that the value of the binarization threshold ϵ between 0.1 and 0.4 gives similar results. For experiments, we set it to 0.3.

The network is trained by minimizing the joint loss function using Adam optimizer [11] with learning rate 2×10^{-5} . The implementation is written in Python using Keras [4] framework with Tensorflow [1] backend. The detection network contains 4.8M parameters, and the matting and fitting network 5.7M parameters. Training takes around three days on Nvidia GTX 1080 Ti GPU with 11 GB memory. Inference takes on average 0.05 seconds on a single image of 600 by 960 pixels, which shows the real-time capability of the proposed method. In comparison, the previous methods [13, 23, 24] require a couple of seconds per frame.

Fig. 5 shows step-by-step visualizations of method outputs. Deblurring results are shown in Fig. 6 and Fig. 7.

4. Experiments

We evaluate the FMO detection accuracy on several real-world datasets. Detection accuracy is measured by precision and recall. A detection is called true positive if intersection over union between the predicted region influenced by the FMO and the ground truth one is greater than 0.1. Then, precision is the percentage of true positives from all detections. Recall is defined as the ratio between the num-




Figure 7. Additional deblurring results on a non-spherical object. From left to right: input image I , background B , estimated H_F and H_M by the matting decoder (removed background), the ground truth from a high-speed camera, object appearance F and mask M estimated by the proposed deblurring, and compared to reconstructions by the deblurring method [13].

ber of true positives and the number of ground truth annotations. If the ground truth sub-frame trajectory \mathcal{C}^* is available, we measure the accuracy of the estimated trajectory \mathcal{C} by Trajectory Intersection over Union (TIOU) [13]:

$$\text{TIOU}(\mathcal{C}, \mathcal{C}^*) = \int_t \text{IoU}(M_{\mathcal{C}(t)}^*, M_{\mathcal{C}^*(t)}^*) dt, \quad (7)$$

where object mask is placed on a trajectory point $\mathcal{C}(t)$, denoted by $M_{\mathcal{C}(t)}^*$, and the standard Intersection over Union (IoU) is computed. Then the average is calculated along the whole trajectory. We compare to the FMO [24] and TbD [13] methods, both designed to detect and track FMOs.

4.1. Ablation study

Ablation study is performed on the real-world TbD dataset [13]. The first ablated version is fully learning-based

	Causal processing				Non-causal processing	
	FMO [24]	TbD [13]	(a) FMODetect	(b) + deblurring	(c) + NC	TbD-NC [23]
Recall \uparrow	0.56	0.96	0.97	0.97	0.99	0.99
TIoU \uparrow	0.352	0.713	0.519	0.715	0.781	0.779
Runtime [1/s] \uparrow	1 fps	0.2 fps	20 fps	0.4 fps	N/A	N/A

Table 1. Recall and trajectory accuracy TIoU (7) on the TbD dataset – comparison of the proposed methods to: FMO method [24], TbD [13], and TbD-NC [23] methods. Versions of the proposed method: (a) real-time with trajectories estimated by the network, (b) with the proposed deblurring, and (c) with non-causal post-processing from [23]. The best performing causal method is highlighted in **bold**. Non-causal methods require expensive processing of all frames simultaneously. Qualitative comparison is in Fig. 8.

Dataset	Metric	FMO	TbD	FMODetect
FMO [24]	Recall \uparrow	35.5	41.1	48.8
	Precision \uparrow	59.2	81.6	84.9
TbD-3D [13]	Recall \uparrow	63.8	69.8	83.5
	Precision \uparrow	81.7	98.2	87.9
Falling [12]	Recall \uparrow	51.5	77.2	82.0
	Precision \uparrow	84.0	92.2	98.9
YouTube	Recall \uparrow	36.4	27.3	90.1
Runtime [1/s] \uparrow		1 fps	0.2 fps	20 fps

Table 2. Precision and recall, averaged on the FMO [24], TbD-3D [13], and the falling objects [12] datasets. We additionally report recall of detected fast moving objects in a trick shots video (Fig. 2) with manually annotated ground truth. FMOdetect is compared to the TbD [13] and FMO [24] methods. The best performing method is highlighted in bold.

prediction. As shown in Table 1 (a), FMOdetect is very accurate in terms of FMO detection (measured by recall). However, if we further evaluate trajectories predicted by the network, the trajectory accuracy (TIoU) is not on par with the state-of-the-art. If we instead of solving the deblatting problem jointly by the network (Table 1 (a), TIoU) solve only the matting problem by the network and run the proposed deblurring to further refine the trajectories (Table 1 (b)), we achieve state-of-the-art trajectory accuracy.

We also show that non-causal post-processing [23] can complete missing detections and make trajectories more consistent within the whole sequence. Compared to TbD-NC [23], TIoU is marginally better.

4.2. Detection accuracy

The TbD dataset captures simplistic scenario of spherical almost uniform objects moving in a constant scene. In this case, the difference between FMOdetect and the other methods is marginal. Further experiments demonstrate that in more challenging cases, the proposed method is significantly more robust. We evaluate detection accuracy on the FMO dataset [24] (various small objects), TbD-3D dataset [25] (spherical highly textured objects moving in

3D), and the falling objects dataset [12] (arbitrarily shaped objects, *e.g.* key, pen, marker). Additionally, we gather 11 short sub-videos from YouTube trick shots video (some are in Fig. 2) and manually annotate the object of interest. Since other FMOs are not annotated in the YouTube dataset, we measure only recall here. In terms of detection accuracy, FMOdetect outperforms all other methods by wide margin on all datasets (Table 2). Only on the TbD-3D dataset, the precision is lower since our method also detected the shadow of the object, and it is still questionable if such detection should be considered a correct detection. Compared to other methods, FMOdetect is faster and runs in real-time.

4.3. Applications

The main application is FMO retrieval as shown on YouTube videos in Fig. 2. Processing both videos (total 15 minutes) took us 15 minutes, whereas the TbD method took around 25 hours. Another application is temporal super-resolution where we generate videos with higher frame rates. After $F, M, H, \mathcal{C}(t)$ are extracted, we split the fitted trajectory $\mathcal{C}(t)$ into the desired number of pieces and apply the formation model (1) to each temporal super-resolution frame separately with the same F, M . Examples are shown in the supplementary material.

4.4. Model generalization

As shown in previous experiments, the method works well on real data, even when trained on synthetic. The method has only seen synthetic FMOs with spherical shape in the training examples, but FMOs of more complex shapes can also be detected. Examples are in Fig. 2 (a hand, a cap, and a keychain), in Fig. 6 (thrown wooden cube), and in Fig. 7 (marker). Since FMOdetect is trained with dynamic cameras and backgrounds from the VOT dataset, it is robust to the camera and background motion (Fig. 2). All previous methods fail to detect or reconstruct objects with low contrast (softball in Fig. 6), while the proposed method succeeds in this task as there was no restriction on the contrast in the synthetic dataset. Among drawbacks of the proposed method are failures on irregular motions, *e.g.* the ceiling fan in Fig. 2 since the synthesized dataset has no such motions. FMO deblurring from a single input image




Figure 8. Trajectory estimation on four sequences from the TbD dataset. From top to bottom: real-time version of the proposed method, with deblurring, with deblurring and non-causal post-processing from [23], TbD [13], TbD-NC [23], the ground truth.

without the background is an open problem – the need for the background image is currently a limitation, but it can be computed for many applications.

5. Conclusion

We proposed the first learning-based method for fast moving object detection. FMODetect improves precision, robustness, and computation time compared to previous methods. In terms of detection accuracy, the proposed methods achieves significantly better results on challenging datasets. By simplifying deblatting into matting and deblur-

ring, we estimate appearance using energy minimization and produce qualitatively better results than deblatting. The proposed method is real-time capable and an order of magnitude faster than previous FMO detection methods. This allows for realistic applications such as efficient FMO scanning or retrieval in large video collections.

Acknowledgements. This research was supported by Research Center for Informatics (project CZ.02.1.01/0.0/0.0/16_019/0000765 funded by OP VVV), Google Focused Research Award, the Czech Science Foundation grant GA21-03921S, the Praemium Academiae awarded by the Czech Academy of Sciences, International Federation of Association Football (FIFA), and Innosuisse grant No. 34475.1 IP-ICT.

References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org. 6
- [2] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, Jan. 2011. 5
- [3] T. Brooks and J. T. Barron. Learning to synthesize motion blur. In *CVPR*, pages 6833–6841, June 2019. 3, 4
- [4] François Chollet. keras. <https://github.com/fchollet/keras>, 2015. 6
- [5] H. Fan, L. Lin, F. Yang, P. Chu, G. Deng, S. Yu, H. Bai, Y. Xu, C. Liao, and H. Ling. Lasot: A high-quality benchmark for large-scale single object tracking. In *CVPR*, pages 5369–5378, June 2019. 1, 2
- [6] Abel Gonzalez-Garcia, Joost van de Weijer, and Yoshua Bengio. Image-to-image translation for cross-domain disentanglement. In *NeurIPS*, page 1294–1305, Red Hook, NY, USA, 2018. Curran Associates Inc. 3
- [7] Qiqi Hou and Feng Liu. Context-aware image matting for simultaneous foreground and alpha estimation. In *ICCV*, October 2019. 3
- [8] Lianghua Huang, Xin Zhao, and Kaiqi Huang. Got-10k: A large high-diversity benchmark for generic object tracking in the wild. *IEEE TPAMI*, page 1–1, 2019. 1, 2
- [9] P. Isola, J. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, pages 5967–5976, July 2017. 3
- [10] Meiguang Jin, Givi Meishvili, and Paolo Favaro. Learning to extract a video sequence from a single motion-blurred image. In *CVPR*, June 2018. 3
- [11] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *ICLR*, 2015. 6
- [12] J. Kotera, J. Matas, and F. Šroubek. Restoration of fast moving objects. *IEEE TIP*, 29:8577–8589, 2020. 3, 7
- [13] J. Kotera, D. Rozumnyi, F. Šroubek, and J. Matas. Intra-frame object tracking by deblatting. In *ICCVW*, Oct 2019. 1, 2, 3, 4, 5, 6, 7, 8
- [14] J. Kotera and F. Šroubek. Motion estimation and deblurring of fast moving objects. In *ICIP*, pages 2860–2864, Oct 2018. 3
- [15] Matej Kristan, Aleš Leonardis, Jiří Matas, Michael Felsberg, Roman Pflugfelder, Luka Čehovin, Tomáš Vojtř, Gustav Hager, Alan Lukežič, Gustavo Fernández, Abhinav Gupta, Alfredo Petrosino, Alireza Memarmoghdam, Alvaro Garcia-Martin, Andrés Solís Montero, Andrea Vedaldi, Andreas Robinson, and other. The visual object tracking vot2016 challenge results. In Gang Hua and Hervé Jégou, editors, *ECCVW*, pages 777–823, Cham, 2016. Springer International Publishing. 1
- [16] Matej Kristan, Ales Leonardis, Jiri Matas, Michael Felsberg, Roman Pflugfelder, Luka Čehovin Zajc, Tomas Vojtř, Goutam Bhat, Alan Lukežic, Abdelrahman Eldesokey, Gustavo Fernandez, and et al. The sixth visual object tracking vot2018 challenge results. In *ECCVW*, Sep 2018. 5
- [17] Andrew L. Maas, Awni Y. Hannun, and Andrew Y. Ng. Rectifier nonlinearities improve neural network acoustic models. In *ICML Workshop on Deep Learning for Audio, Speech and Language Processing*, 2013. 4
- [18] S. Nah, T. H. Kim, and K. M. Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *CVPR*, pages 257–265, July 2017. 3
- [19] Augustus Odena, Vincent Dumoulin, and Chris Olah. Deconvolution and checkerboard artifacts. *Distill*, 2016. 4
- [20] Liyuan Pan, Richard Hartley, Miaomiao Liu, and Yuchao Dai. Phase-only image based kernel estimation for single image blind deblurring. In *CVPR*, June 2019. 3
- [21] Kuldeep Purohit, Anshul Shah, and A. N. Rajagopalan. Bringing alive blurred moments. In *CVPR*, June 2019. 3
- [22] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, editors, *MICCAI*, pages 234–241, Cham, 2015. Springer International Publishing. 3, 4
- [23] D. Rozumnyi, J. Kotera, F. Šroubek, and J. Matas. Non-causal tracking by deblatting. In Gernot A. Fink, Simone Frintrop, and Xiaoyi Jiang, editors, *GCPR*, pages 122–135, Cham, 2019. Springer International Publishing. 3, 6, 7, 8
- [24] D. Rozumnyi, J. Kotera, F. Šroubek, L. Novotný, and J. Matas. The world of fast moving objects. In *CVPR*, pages 4838–4846, July 2017. 1, 3, 5, 6, 7
- [25] D. Rozumnyi, J. Kotera, F. Šroubek, and J. Matas. Sub-frame appearance and 6d pose estimation of fast moving objects. In *CVPR*, pages 6777–6785, 2020. 3, 7
- [26] A. W. M. Smeulders, D. M. Chu, R. Cucchiara, S. Calderara, A. Dehghan, and M. Shah. Visual tracking: An experimental survey. *IEEE TPAMI*, 36(7):1442–1468, July 2014. 1
- [27] F. Šroubek and J. Kotera. Motion blur prior. In *ICIP*, pages 928–932, 2020. 3
- [28] Yi Wu, Jongwoo Lim, and Ming-Hsuan Yang. Online object tracking: A benchmark. In *CVPR*, 2013. 1
- [29] N. Xu, B. Price, S. Cohen, and T. Huang. Deep image matting. In *CVPR*, pages 311–320, July 2017. 3
- [30] M. D. Zeiler, D. Krishnan, G. W. Taylor, and R. Fergus. Deconvolutional networks. In *CVPR*, pages 2528–2535, June 2010. 4
- [31] M. D. Zeiler, G. W. Taylor, and R. Fergus. Adaptive deconvolutional networks for mid and high level feature learning. In *ICCV*, pages 2018–2025, Nov 2011. 4
- [32] Y. Zhang, L. Gong, L. Fan, P. Ren, Q. Huang, H. Bao, and W. Xu. A late fusion cnn for digital matting. In *CVPR*, pages 7461–7470, June 2019. 3