

Multi-Modal Fusion for End-to-End RGB-T Tracking

Lichao Zhang¹, Martin Danelljan², Abel Gonzalez-Garcia¹, Joost van de Weijer¹, Fahad Shahbaz Khan³

¹ Computer Vision Center, Universitat Autònoma de Barcelona, Spain

² Computer Vision Laboratory, ETH Zürich, Switzerland

³ Inception Institute of Artificial Intelligence, UAE

{lichao, agonzalez, joost}@cvc.uab.es, martin.danelljan@vision.ee.ethz.ch, fahad.khan@inceptioniai.org

Abstract

We propose an end-to-end tracking framework for fusing the RGB and TIR modalities in RGB-T tracking. Our baseline tracker is DiMP (Discriminative Model Prediction), which employs a carefully designed target prediction network trained end-to-end using a discriminative loss. We analyze the effectiveness of modality fusion in each of the main components in DiMP, i.e. feature extractor, target estimation network, and classifier. We consider several fusion mechanisms acting at different levels of the framework, including pixel-level, feature-level and response-level. Our tracker is trained in an end-to-end manner, enabling the components to learn how to fuse the information from both modalities. As data to train our model, we generate a large-scale RGB-T dataset by considering an annotated RGB tracking dataset (GOT-10k) and synthesizing paired TIR images using an image-to-image translation approach. We perform extensive experiments on VOT-RGBT2019 dataset and RGBT210 dataset, evaluating each type of modality fusing on each model component. The results show that the proposed fusion mechanisms improve the performance of the single modality counterparts. We obtain our best results when fusing at the feature-level on both the IoU-Net and the model predictor, obtaining an EAO score of 0.391 on VOT-RGBT2019 dataset. With this fusion mechanism we achieve the state-of-the-art performance on RGBT210 dataset.

1. Introduction

As an important task in computer vision, visual object tracking, especially RGB tracking [7, 22, 4, 16, 12, 17, 10, 30, 61, 38, 11], has undergone profound changes in recent years. Researchers mainly focus on RGB tracking as large datasets are available [53, 28, 49]. However, RGB tracking obtains unsatisfactory performance in bad environmental conditions, e.g. low illumination, rain, and smog. It was found that thermal infrared sensors provide a more stable signal for these scenarios. Therefore, RGB-T tracking has

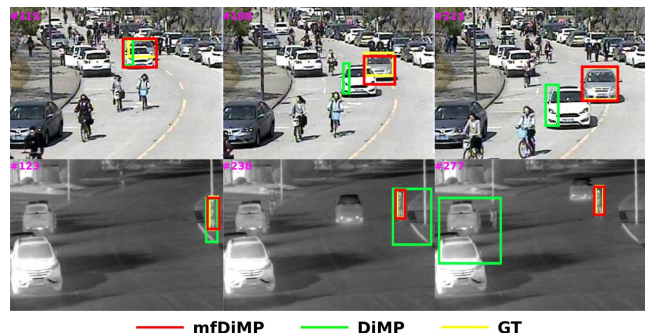


Figure 1. **Qualitative comparison between ‘mfDiMP’ and ‘DiMP’.** Two exemplar videos from RGB modality and TIR modality on the top and bottom separately, where DiMP performs on each of them with single modality input. Our *mfDiMP* can effectively track the object by fusing both modalities.

drawn more research attention recently [31, 34, 32, 35].

As multi-modal data, i.e. from the RGB and TIR modalities, can provide complementary information for tracking, multi-modal tracking is a promising research direction. Images from the RGB modality have the advantage that they contain high-frequency texture information and provide rich representations for describing objects. Images from the TIR modality have the advantage that they are not influenced by illumination variations and shadows. Moreover, objects with elevated temperature can be distinguished from the background as the background is normally colder. Therefore, fusing the information from multi-modal data could benefit the tracker because it can exploit the complementary information of the modalities to improve tracking performance.

There exists relatively little research on multi-modal tracking [52, 37, 31, 33, 35]. Most of these works are still using the sparse representation, normally with the hand-crafted features, for multi-modal tracking [37, 31, 33, 35]. Later on in [35], for comparison they design some baseline RGB-T trackers by extending the single modal tracker to a multi-modal tracker. This extension is done by directly concatenating the features from the RGB and the TIR

modalities into a single vector, which is then fed into the tracker. They also use some deep features for concatenation, but they are still off-the-shelf features pre-trained for other tasks. Therefore, there is still no previous work which investigates end-to-end training. We mention two main reasons for this. First, it is not obvious in what part of the tracking pipeline the fusion should be done. Ideally, we should fuse the information of the different modalities in such a way that it allows for optimal end-to-end training. Second, data scarcity of multi-modal tracking is a major obstacle to end-to-end training. Currently there are no large-scale aligned multi-modal datasets. These two issues, i.e. no specific fusion scheme and lack of data, limit the progress of end-to-end multi-modal training.

To tackle this problem, in this paper we investigate how to effectively fuse multi-modal data in an end-to-end manner, which enables the optimal use of information from both modalities (see Figure 1). We propose three end-to-end fusion architectures, consisting of pixel-level fusion, feature-level fusion, and response-level fusion. We use as baseline tracker the RGB tracker DiMP [5]. To ensure that the proposed fusion tracker can be trained in an end-to-end manner, we also generate a large-scale paired synthetic RGB-T dataset with the method proposed in [58]. We perform extensive experiments on two commonly used benchmarks for RGB-T tracking: VOT-RGBT2019 [1] and RGBT210 [34]. Our multi-modal fusion tracker sets a new state-of-the-art on both datasets, achieving an EAO score of 0.391 on VOT-RGBT2019 and 55.5% success rate on RGBT210.

This paper is organized as follows. In section 2, we discuss the successful single modality tracking methods of recent years and the situation of current multi-modal tracking. In section 3, we introduce the baseline tracker and analyze its components. In section 4, we describe the proposed methods and formulations for the fusion of multi-modal tracking and provide the synthetic data for end-to-end training. In section 5, we present our extensive experiments on the VOT-RGBT2019 dataset and RGBT210 dataset. Finally, in section 6, we conclude our work and propose future research directions.

2. Related work

2.1. Single modality tracking

Most current tracking algorithms focus on RGB images [7, 22, 4, 16, 12, 17, 10, 30, 61, 38, 11], although several approaches track in the TIR modality instead [58, 54, 55, 35]. Despite the development of deep learning in many computer vision tasks, object tracking continued to use hand-crafted features during the first stage of deep learning [39, 44, 17, 10, 47]. Later on, some trackers [17, 10, 39] pioneered in the involvement of deep features in tracking by using the pre-trained models for an image classification task [46]. The main reasons for only us-

ing pre-trained models were the lack of large-scale training datasets and the difficulty of designing a suitable end-to-end training framework for tracking. Bertinetto *et al.* [4] proposed to train a network end-to-end by using a video object detection dataset [46]. Recently, several large-scale tracking datasets [18, 43, 23, 49], e.g. GOT-10k [23], have been released with millions of images and various categories for training. Therefore, some current tracking approaches [11, 29, 5] perform end-to-end training by leveraging these large-scale datasets.

RGB trackers. Bertinetto *et al.* [4] proposed to use a fully-convolutional architecture to learn a similarity metric offline, i.e. a Siamese network. After training, the Siamese network is deployed for online tracking with high efficiency. To learn attention on the cross correlation, Wang *et al.* [51] include additional attention components in the Siamese network and learn the spatial and channel weights for the *exemplar* model. Li *et al.* [30] utilize a proposal network to estimate the score maps and bounding boxes using two branches, which provides more accurate object scales than the traditional multi-resolution scale estimation. Later it is extended to use deeper and wider networks achieving significant improvement [29].

An alternative approach to Siamese networks is correlation filter (CF) based tracking [7, 22, 20, 16, 56, 59, 57, 12, 14, 15, 41, 26, 38, 40], which has occupied top positions for many years given its discriminative abilities and efficient tracking speed. The core part of CF trackers is the calculation of a filter that is later applied to detect the object in the search region of next frame. The calculation is performed in the Fourier domain, which makes it highly efficient. To overcome the issue of boundary effect in correlation filter in tracking, Danelljan *et al.* [14] proposed to regularize the filter with a Gaussian window and Kiani *et al.* [26] proposed to use a mask formulated in the correlation filter. Some CF trackers [39, 13, 17, 10] also benefited from pre-trained convolutional features. CFNet [48] added end-to-end training by formulating CF as one layer of the network, although this only gives a marginal gain with respect to the baseline model, SiamFC [4]. Park *et al.* [45] proposed to learn an initial model for the correlation filter offline, accelerating the convergence speed for the filter optimization.

TIR trackers. Contrarily to RGB tracking, most of the top performing TIR trackers still use hand-crafted features in their models. For example, SRDCFir [19] extends the SRDCF [14] tracker for TIR data by combining motion features with hand-crafted visual features, e.g. HOG [9], color names [50], intensity, etc. EBT [60] uses edge features to devise an objectiveness measure that generates high quality object proposals. Yu *et al.* [54, 55] propose structural learning on dense samples around the object, using edge and HOG features [9], transferred to the Fourier domain for efficiency. Zhang *et al.* [58] propose using an end-to-end train-

able deep network. They generate a large-scale TIR tracking dataset for training from existing RGB tracking dataset. They use a current image translation approach [24] to synthesize a large amount of TIR images from RGB and they transfer the corresponding object annotations. By training the network with this data, they achieve state-of-the-art results in TIR tracking. Following this idea, we obtain a large-scale RGB-T dataset that enables the use of deep learning for RGB-T tracking.

2.2. Modality fusion tracking

Fusing the RGB and TIR modalities is a promising direction. Some RGB-T trackers once have been proposed. Conaire *et al.* [8] proposed to efficiently combine visible and thermal features by fusing the outputs of multiple spatio-temporal trackers, which is a derivation from mean-shift type algorithm [6]. Wu *et al.* [52] used a sparse representation for the target template by concatenating RGB and TIR image patches. Similarly, Liu *et al.* [37] also use a sparse representation by minimizing the coefficients from each modality. However, these methods provide sub-optimal fusions as both modalities contribute equally, while in practice one modality may have more valuable information than the other. Li *et al.* [31, 33] addressed this with an adaptive fusion scheme to integrate visible and thermal information in the sparse representation by introducing weights to balance the contribution of each modality. In order to limit the effect of background clutter during tracking, Li *et al.* [35] introduced a ranking between the two modalities, which is taken into account in the used patch-based features. They effectively avoided background effects by using the learned features with a structured SVM.

As far as we know, all of the current RGB-T approaches use hand-crafted features, which significantly limits their tracking performance. Although several RGB-T tracking datasets [31, 34, 32] have been recently released, they are only for testing purposes and are not large enough for training a deep learning based RGB-T tracker. We propose adapting a deep RGB tracker for RGB-T by exploring different types of modality fusion, and performing end-to-end training with partly synthesized RGB-T data.

3. Baseline RGB tracker

In this section, we describe the architecture of the tracker we have selected for our multi-modal tracking experiments. We use the Discriminative Model Prediction (DiMP) tracker [5], which was originally proposed for single modality tracking.

Discriminative Model Prediction. DiMP [5] proposed an end-to-end trainable tracking architecture, capable of learning a powerful discriminative filter by embedding the online learning of the target model into itself. DiMP consists of the following components: feature extractor, model predictor,

and target estimation network (IoU-Net [25]). With these carefully designed components and an effective optimization method, they achieve excellent performance on RGB tracking by setting a new state-of-the-art on several RGB tracking datasets [27, 53, 23, 43, 18].

Feature extractor. The backbone feature extractor F normally aims to extract the deep feature representations for the follow-up implementation models. Here, specifically in DiMP [5], the deep representations are extracted for the model predictor and target estimation network.

DiMP [5] employs the ResNet-18 and ResNet-50 architectures, which is trained on ImageNet, as the backbone feature extractors for DiMP-18 and DiMP-50 separately. They implement fine-tuning the backbone for the end-to-end training. After an analysis on the impact of different feature blocks in DiMP [5], they use the features from block3 and block4 for IoU-Net, and only from block4 for the classifier. The feature extractor F is shared and only performed on a single image patch per frame.

For training the feature extractor F , they input data for F with a pair of sets (M_{train}, M_{test}) . Each set $M = \{(I_j, b_j)\}_{j=1}^{N_{frames}}$ contains images I_j along with their object bounding box b_j . The target model is predicted by using M_{train} and then evaluated on the test frames M_{test} . M_{train} and M_{test} are constructed by sampling N_{frames} frames for both from first and second halves of the segment respectively. They pass the images through the feature extractor F , and obtain the train set $S_{train} = \{(x_j, c_j)\}$, where $x_j = F(I_j)$, and c_j is the center coordinate of the box b_j .

Model predictor and response map. The Model predictor is to obtain the final optimized filter f , which consists of model initializer, which is a convolutional layer followed by a precise ROI pooling[25], and model optimizer, which is to solve the final model f by the steepest descent (SD). The model filter f is solved by using multiple samples in S_{train} , which happens in the model initializer. The input of the model predictor is a set of S_{train} , and obtain the model f by training on the model predictor: $f = D(S_{train})$. Then the filter f is evaluated on the test samples S_{test} and finally classification loss for offline training is computed as:

$$L_{cls} = \frac{1}{N_{iter}} \sum_{i=0}^{N_{iter}} \sum_{(x,c) \in S_{test}} \left\| l(x * f^{(i)}, z_c) \right\|^2. \quad (1)$$

Where, z_c is a Gaussian function centered as the target c . $f^{(0)}$ is the output of model initializer. The response map can be calculated as: $s = x * f, x \in S_{test}$.

Bounding box estimation. DiMP uses an IoU-Net based architecture from the ATOM tracker [11]. The function of the IoU-Net model is to predict the IoU between the deep feature x of an image and a bounding box candidate B . Bounding box estimation is then performed by maximizing the IoU prediction. The network has two branches, one

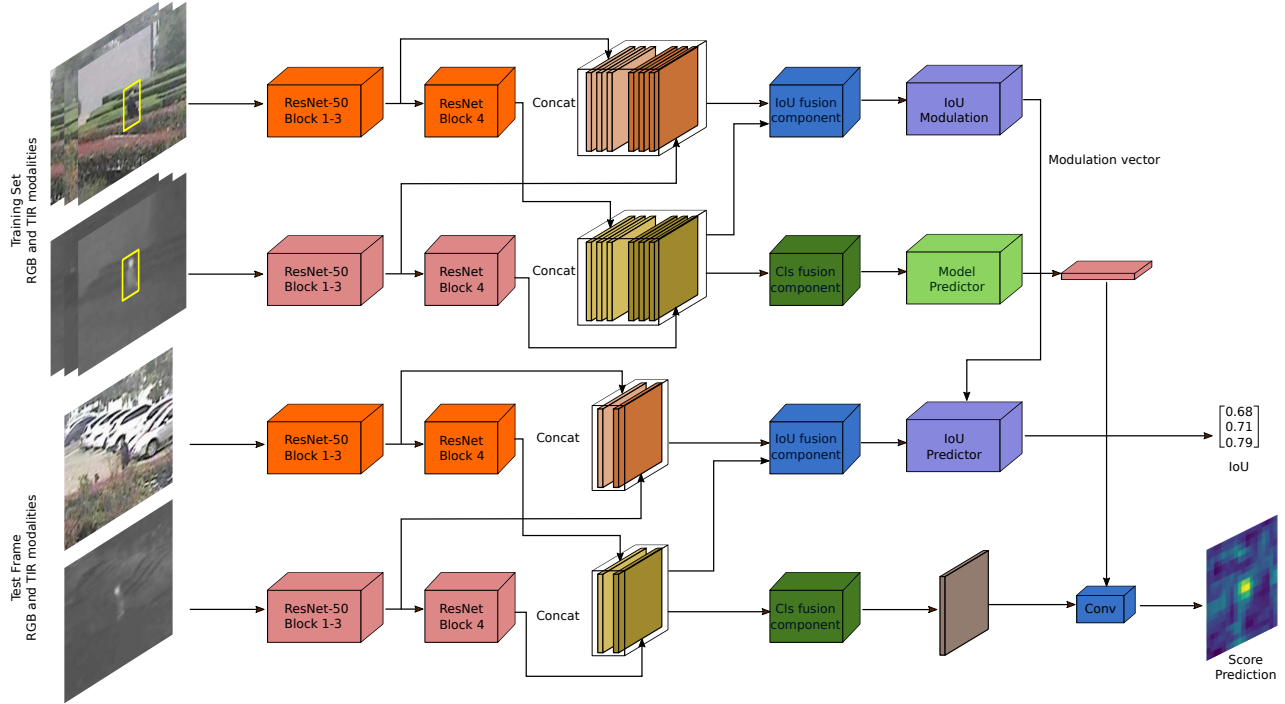


Figure 2. **Overview of our multi-modal fusion framework on feature-level.** We input images from RGB and TIR modalities to their feature extractor separately. Then we fuse the deep features from different blocks of the backbone. Fused features from block3 and block4 are input to IoU modulation and IoU predictor. Fused features from block4 are input to the model predictor for the final response map.

is the *IoU modulation* for calculating the modulation vector from reference image, and the other branch is *IoU predictor* for predicting the IoU values from test image. Then the reference branch is added with a convolutional layer, while the test branch is added with two convolutional layers as it dominates the IoU prediction. Both of them then are followed by PrPool (Precise ROI Pooling)[25] and a fully connected layer. Here the interaction between the two branches is that a precomputed vector in the reference branch is used to modulate the feature representation of the test image via channel-wise correlation. The IoU is predicted in terms of the bounding box B as follows:

$$IoU(B) = g(c(x_0, B_0) \cdot z(x, B)) \quad (2)$$

Where, x_0, B_0 are from the reference image, and x, B are from the test image. z is the feature representation after PrPool layer in test branch. g is the IoU predictor with three fully connected layers. c is a modulation vector.

4. End-to-end multi-modal tracking

There are two main issues when extending state-of-the-art RGB trackers to multi-modal data such as RGB-T. First, a fusion component is not considered as a native designed component for the RGB tracker architecture, since the tracker only considers a single modality as input. Therefore, when extending to multi-modal data, these trackers must be equipped with a fusion strategy. Second, the lack of large-scale paired RGB-T training datasets complicates the

end-to-end training of feature representations, which have been shown to significantly improve results for RGB tracking. To tackle the former, we investigate how to effectively fuse multi-modal data for tracking, aiming to make the best use of all available data modalities, in this case, RGB and TIR. To tackle the latter, we ensure that the proposed multi-modal tracker can be trained in an end-to-end manner by generating a large-scale paired synthetic RGB-T dataset, similarly to the method proposed in [58].

In this section, we first comprehensively explain our three end-to-end multi-modal fusion architectures, namely pixel-level fusion, feature-level fusion, and response-level fusion. We also explain how we apply [58] to generate a large multi-modal dataset.

4.1. Multi-modal fusion for tracking

In this subsection, we investigate three different mechanisms for multi-modal fusion with the aim to find the optimal fusion architecture. We start the fusion work on the input of the network (pixel-level). Then we explore the fusion on the intermediate of the network. In [35], they extended some RGB trackers by concatenating the RGB and TIR features into a single vector, and then used them as off-the-shelf features for the classifiers of various trackers. In contrast, we end-to-end train fused features, which are input to both the model predictor and the target estimation network (feature-level). Moreover, we explore fusion on the final response maps of the network (response-level).

Pixel-level fusion. The first modality fusion we consider is at the input of the network. We propose to fuse the RGB and TIR images by directly concatenating the images along the channel direction and then inputting the fused RGB-T image to the feature extractor. To complete this fusion, we extend the filter size of the first layer in feature extractor from $7 \times 7 \times 3 \times 64$ to $7 \times 7 \times 4 \times 64$. The images that are input to the feature extractor should be concatenated as $I^F = [I^V | I^T]$, where I^V is the RGB image, I^T is the TIR image and the fused image is I^F .

Feature-level fusion. To delay the fusion to a more semantically aware network stage, we evaluate the fusion effectiveness in the intermediate part of the network architecture. Concretely, we implement the fusion after the feature extractor, i.e. fusing the deep feature representations from the RGB and TIR modalities. We pass the RGB and TIR images through the feature extractors separately and extract features from both modalities independently. Then, we concatenate the features from each modality and feed them into the IoU predictor and model predictor. This provides a more expressive representation for the IoU predictor and more discriminative features for the model predictor. The framework of our proposed method for multi-modal fusion on feature-level is shown in Figure 2, where we show how we concatenate the feature representations output by the feature extractors. The feature concatenation can be expressed as intuitive syntax: $x^F = [x^V | x^T]$. Here, x^V is the features from the RGB modality, x^T is the features from the TIR modality, and x^F is the fused features.

Response-level fusion. To evaluate the effectiveness of an ensemble of independently trained trackers on each modality, we perform the multi-modal fusion on the final part of the training architecture in DiMP [5], i.e. response-level fusion. For the response-level fusion, we use a pair of feature extractors and model predictors to process each image from RGB and TIR modalities separately. Finally, we sum their response maps to get the fused response map. We input a single modality to the IoU-Net component and thus there are two cases for training the whole network, one using the RGB modality to fine-tune IoU-Net and one using the TIR modality instead. Assuming that we have two single modality response maps, s^V from the RGB modality and s^T from the TIR modality, we calculate the fused response map by summing them: $s^F = s^V + s^T$.

4.2. RGB-T data generation

The lack of large-scale paired RGB-T training datasets hampers end-to-end tracking in RGB-T datasets. We borrow the method from [58], which proposes to use image-to-image translation methods to generate synthetic TIR data for tracking. In their paper, they show that using such data improves results for end-to-end training of TIR trackers. Here we explain how we generated the training data aiming for fine-tuning the pre-trained DiMP models. We take ad-

vantage of a normal RGB training dataset for RGB tracking, and then generate the TIR images by a well-trained image-to-image translation model [24]. With the above steps, we obtain an aligned synthetic RGB-T training dataset for RGB-T tracking. As a result, our proposed fusion architectures (see section 4.1) can also benefit from end-to-end training. Ideally, this will allow us to obtain a performance gain proportional to that observed for RGB tracking.

After applying the described process, we obtain two datasets for RGB-T tracking, (M_{train}, M_{test}) , both of the form $M = \{(I_j^V, I_j^T, b_j)\}_{j=1}^{N_{frames}}$. Here, I_j^T represents the j -th synthetic TIR image generated from the aligned RGB image I_j^V , and b_j is their identical bounding box.

5. Experiments

In this section, we provide a comprehensive evaluation of the proposed tracker mfDiMP on two benchmarks, VOT-RGBT2019 [1] and RGBT210 [34], and describe all implementation and evaluation details.

5.1. Generating the training RGB-T dataset

We use the recent Generic Object Tracking Benchmark (GOT-10k) [23] to train our fused modality networks. GOT-10k has over 10,000 video segments, covering 563 classes of real-world moving objects and more than 80 motion patterns, amounting to a total of over 1.5 million manually labeled bounding boxes. It also provides additional supervision in terms of attribute labels such as ratio of object visible or motion type. We employ GOT-10k’s training set, which contains 9,335 videos (1,403,359 frames), with 480 object classes and 69 motion classes. We refrain from using the set of 1000 prohibited videos listed in the VOT challenge website [1], so we train our model with the remaining 8,335 videos (1,251,981 frames).

With this reduced version of GOT-10k RGB dataset, we generate a large-scale RGB-T dataset by synthesizing paired TIR images using an image-to-image translation approach, as in [58]. Specifically, we use pix2pix [24] for image-to-image translation given its excellent performance [58]. To train the pix2pix model, we use a total of 87K pairs of aligned images in the RGB and TIR modalities, depicting several different scenarios. These images are carefully collected and arranged from many current existing RGB-TIR datasets [58]. We train the pix2pix model using the default settings described in [58]. After training, we use pix2pix to transfer the selected RGB videos in GOT-10k [23] to synthetic TIR videos, along with the labels.

5.2. Evaluation datasets and protocols

VOT-RGBT2019 dataset [1] contains 60 public testing sequences, with a total of 20,083 frames. It is used as the most recent edition of the VOT challenge. We follow the VOT protocol, which establishes that when the evaluated tracker fails, i.e. when the overlap with the ground-truth

Fusion level	Feature extractor	IoU-Net	Model predictor	Response map	EAO (\uparrow)	A (\uparrow)	R (\downarrow)
Single modality	RGB	RGB	RGB	RGB	0.327	0.586	0.345
	TIR	TIR	TIR	TIR	0.331	0.584	0.332
	TIR	TIR (ft)	TIR (ft)	TIR	0.336	0.587	0.331
	TIR (ft)	TIR	TIR (ft)	TIR	0.339	0.589	0.329
	TIR (ft)	TIR (ft)	TIR (ft)	TIR	0.341	0.590	0.328
	RGB (ft)	RGB (ft)	RGB (ft)	RGB	0.335	0.586	0.331
Pixel-level	RGB+TIR (ft)	RGBT (ft)	RGBT (ft)	RGBT	0.345	0.552	0.281
Response-level	RGB/TIR (ft)	RGB (ft)	RGB/TIR (ft)	RGB+TIR	0.342	0.546	0.309
	RGB/TIR (ft)	TIR (ft)	RGB/TIR (ft)	RGB+TIR	0.349	0.554	0.291
Feature-level	RGB/TIR (ft)	RGB (ft)	RGB+TIR (ft)	RGBT	0.346	0.545	0.266
	RGB/TIR (ft)	TIR (ft)	RGB+TIR (ft)	RGBT	0.359	0.564	0.243
	RGB/TIR (ft)	RGB+TIR (ft)	RGB (ft)	RGB	0.354	0.563	0.276
	RGB/TIR (ft)	RGB+TIR (ft)	TIR (ft)	TIR	0.366	0.601	0.261
	RGB/TIR (ft)	RGB+TIR (ft)	RGB+TIR (ft)	RGBT	0.389	0.605	0.224
	RGB/TIR (ft / ft \times 10)	RGB+TIR (ft)	RGB+TIR (ft)	RGBT	0.391	0.615	0.228

Table 1. **Fusion mechanisms analysis on VOT-RGBT2019 [1]**. We evaluate several fusion mechanisms at different levels of DiMP [5]. The results are reported in terms of EAO, normalized weighted mean of accuracy (A), and normalized weighted mean of robustness score (R). We explicitly show the input modality for each component of the tracker. Here, ‘RGB’ and ‘TIR’ are the single modality, ‘RGB/TIR’ means each modality input separately, ‘RGB+TIR’ means that both modalities are input simultaneously, and ‘RGBT’ indicates fused features from both modalities used in the remaining of network. Finally, (‘ft’) means fine-tuning and (‘ft \times 10’) means fine-tuning with a higher learning rate. The best results are highlighted in bold font.

is below a given threshold, it is re-initialized in the correct location five frames after the failure. The main evaluation measure used to rank the trackers is Expected Average Overlap (EAO), which is a combination of accuracy (A) and robustness (R). We compute all results using the provided toolkit [1].

RGBT210 dataset [34] contains 210 highly-aligned public RGB and TIR video pairs for testing, with 210K frames in total and a maximum of 8K frames per sequence pair. There are a total of 12 representative attributes, such as camera moving, large scale variations and environmental challenges, which are annotated for each video. These facilitate attribute-sensitive evaluation analyses. We compare our results with other trackers using the provided toolkit [2]. We use precision plot and success plot to evaluate the trackers.

5.3. Implementation details

We use DiMP [5] as our base tracker with ResNet-50 [21] as backbone network. The base architecture of DiMP is pre-trained on several large-scale RGB training datasets [18, 43, 23, 36]. To test the single modality versions, we simply input the images from either modality as in traditional RGB trackers [22, 10, 4]. RGB images have 3 image channels while TIR images have 1 channel, and so the pixel-level fusion uses 4-channel images. For the feature-level fusion, we concatenate the convolutional features after the feature extractors. Finally, for the response-level fusion we add together the final confidence maps independently predicted by the RGB and TIR modalities.

We use separate, modality-specific feature extractors for the response-level fusion and feature-level fusion. As hyperparameters for fine-tuning our architecture, we use the

default values used to train each component in DiMP [5], which have been carefully set and described by the authors in section 3.2 of [5]. We keep the default learning rates for each component as in the DiMP model and then decrease them by collaboratively multiplying a small gain learning rate, i.e. 0.001 when fine-tuning. In one of our experiments, we set the learning rate for the TIR feature extractor \times 10 higher than for the RGB feature extractor. We do this considering that the RGB feature extractor was pre-trained with a large-scale RGB dataset, leading to satisfactory RGB features. On the other hand, the TIR feature extractor needs to catch up with that of the RGB modality in terms of learning, and thus it requires a higher learning rate for fine-tuning.

As a result of the stochastic nature of DiMP, the tracker generates different results for every run. Following the procedure employed in [5], we compute the default 15 runs of our mfDiMP tracker for VOT-RGBT2019 dataset and 5 for RGBT210 dataset. Then we obtain the final result by averaging the results of all runs.

5.4. Analysis of fusion mechanisms

Table 1 presents our analysis to determine the best location to fuse the modalities in DiMP. The table is an extensive evaluation of all considered fusions under different configurations. We start with a comprehensive evaluation of the baseline tracker DiMP [5] for single modality. We present several configurations in the upper part of Table 1 (‘Single modality’) using the RGB modality or TIR modality alone. The first two rows are the original DiMP (pre-trained for RGB) using either RGB or TIR images during online tracking. We observe how the TIR images obtain a higher result. For the next three rows, we fine-tune the feature extractor and/or IoU-Net with synthetic TIR images. In this case,

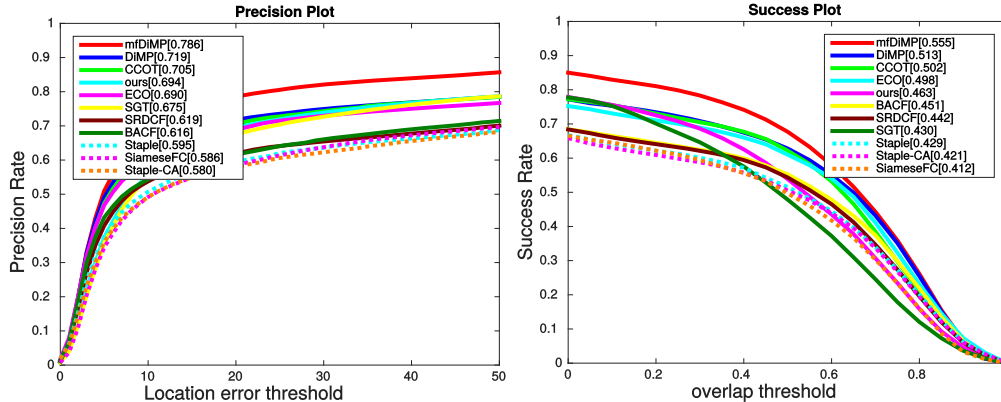


Figure 3. **Precision plot and success plot by comparing our mfDiMP with the top-10 trackers on RGBT210 dataset [34].** We can see our mfDiMP outperforms DiMP with an absolute gain of 6.7% and 4.2% in terms of precision rate and success rate respectively.

	ECO [10]	SiamFC [4]	DaSiamRPN [61]	ATOM [11]	DiMP [5]	mfDiMP
EAO(↑)	0.265	0.254	0.324	0.318	0.327	0.391
A (↑)	0.580	0.594	0.604	0.575	0.586	0.615
R (↓)	0.480	0.533	0.482	0.374	0.345	0.228
FPS (↑)	11.2	38.1	62.4	12.1	13.6	10.3

Table 2. **State-of-the-art comparison on VOT-RGBT2019 dataset.** Our mfDiMP improves the baseline tracker DiMP with an absolute gain of 6.4% in terms of EAO without significantly deteriorating the computational efficiency. The best results are highlighted in bold font.

fine-tuning the single modality network for TIR improves the pre-trained networks with an absolute gain of 1%. Finally, we can see how fine-tuning only on RGB improves the performance of the pre-trained model, but to a lesser extent than using TIR. In the lower part of Table 1 we analyze the effectiveness of each fusion mechanism for DiMP [5], which we discuss in detail in the remainder of this section.

Pixel-level fusion. From Table 1, we can observe that pixel-level fusion improves the performance of the baseline tracker from 0.331 to 0.345 with an absolute gain of 1.4%. The images from the RGB and TIR modalities have complementary information. Therefore, using the fused images to train the network end-to-end can help the model to learn better deep feature representations, which in turn improves tracking performance.

Response-level fusion. In this case, the fusion takes place in the final response map output by the classifier. The signals from both modalities pass through the classifier separately and both compute the response map. Then we sum together the two response maps, obtaining the final fused response map. Meanwhile, we input a single modality for the IoU-Net, either RGB or TIR. Both fusion mechanisms enhance the tracking performance, and using the TIR modality for IoU-Net outperforms using RGB, achieving scores of 0.349 and 0.342, respectively. The results show that fusion on response-level obtains the same effectiveness as pixel-level fusion.

Feature-level fusion. We consider inputting fused feature representations into two different DiMP components, i.e. model predictor and IoU-Net. In the former case, only the model predictor receives fused features, whereas the remaining component (IoU-Net) still uses features from a single modality. In both cases, we improve over single modality tracking, and the version with TIR for IoU-Net obtains a better result with 0.359, an absolute gain of 2.8% with respect to the best single modality model (fine-tuned on TIR, row 5 of Table 1). We obtain an even higher score, 0.366, when fusing features for IoU-Net while using TIR features for the model predictor. This demonstrates that using fused features for IoU-Net outperforms using fused features for the model predictor. We attribute this to the fact that IoU-Net is a more complex network than the model predictor, where the fused feature representations can be more effective to prompt IoU-Net to its ultimate capacity. Finally, we use fused features to feed both model predictor and IoU-Net, which significantly improves the result to 0.389, a substantial absolute gain of 5.8%. Considering that the feature extractor is pre-trained on RGB images, it is natural to assume that the feature extractor for TIR images needs a stronger training signal. For this reason, we propose a variant in which the feature extractor for the TIR modality has a higher learning rate ($\times 10$). This variant achieves 0.391, which is the best result and significantly outperforms the best single modality tracker with a big jump.

From Table 1, we can see that fusion on feature-level with end-to-end training provides significant improvements in tracking performance. Specifically, fusion of the feature representations for both model predictor and IoU-Net achieves the best result on VOT-RGBT2019 dataset [1]. We can also see that as pixel-level fusion and response-level fusion both take place in the extra part of the network, they are easier to implement and only fewer variants need to be evaluated, compared with the fusion on the intermediate of the network. In the following sections, we select this best performing variant as our final tracker, which we call *mfDiMP*.

	ECO [10]	CSR [38]	CMRT [35]	SGT [34]	CNN+KCF+ RGBT[22]	CFnet+ RGBT[48]	mfDiMP
No Occlusion	87.7/64.3	68.1/45.2	86.1/59.4	82.4/50.7	63.7/42.9	69.7/52.2	88.9/67.3
Partial Occlusion	72.2/52.5	52.7/36.6	77.1/52.2	75.4/48.3	56.0/36.4	57.2/38.4	84.0/60.1
Heavy Occlusion	58.3/41.3	37.1/24.3	54.6/34.8	53.1/34.1	36.6/25.9	39.3/27.3	68.4/45.8
Low Illumination	66.6/45.6	47.3/31.1	71.4/46.4	71.6/44.7	52.8/34.5	49.8/33.6	77.1/53.7
Low Resolution	64.1/38.1	46.0/23.1	64.7/37.4	65.8/37.5	54.6/32.5	45.2/27.7	69.2/43.6
Thermal Crossover	82.1/58.8	43.2/29.3	65.8/43.0	64.9/40.7	49.6/33.2	42.8/29.4	76.5/55.2
Deformation	61.2/45.0	44.7/33.0	65.2/45.8	65.3/45.9	44.8/34.4	48.9/35.2	77.7/56.6
Fast Motion	58.2/39.2	42.6/25.0	58.8/34.9	58.0/33.1	37.1/24.1	36.5/23.0	76.7/52.6
Scale Variation	74.5/55.4	53.3/37.5	72.5/49.2	67.4/41.7	50.3/32.6	56.7/40.6	82.2/59.5
Motion Blur	67.8/49.9	34.7/23.8	58.4/40.5	58.6/39.6	30.4/22.0	30.3/22.4	72.5/51.2
Camera Moving	61.7/45.0	38.9/27.4	60.0/41.9	59.0/40.7	36.2/27.0	37.2/27.9	75.3/53.8
Background Clutter	52.9/35.2	38.4/23.7	58.3/35.6	58.6/35.5	42.3/28.4	43.7/28.1	71.5/45.7
ALL	69.0/49.8	49.1/33.0	69.4/46.3	67.5/43.0	49.3/33.1	51.8/36.0	78.6/55.5

Table 3. **Attribute-based Precision Rate and Success Rate (PR/SR %) on RGBT210 dataset with several trackers.** These trackers include popular RGB trackers such as ECO and CSR, recent multi-modal fusion tracker like CMRT and SGT, and also extended RGB-T trackers from KCF and CFnet. Our tracker surpasses almost all the trackers over all the attributes.

5.5. VOT-RGBT2019 dataset

In this section, we evaluate our mfDiMP on the VOT-RGBT2019 dataset in terms of EAO in Table 2. We compare with several high-quality RGB trackers including ECO [10], ATOM [11], DiMP [5], SiameseFC [4], DaSi-amRPN [61]. All of these use only the RGB modality as input. The single modality baseline tracker DiMP [5], which shows dominant performances on various RGB datasets [27, 43, 18, 23], also achieves excellent results on VOT-RGBT2019. By using our multi-modal fusion with end-to-end training, we improve DiMP by an absolute gain of 6.4% in terms of EAO. This significant improvement demonstrates that our selected fusion mechanism is effective for maximally exploiting the multi-modal nature of the given images.

5.6. RGBT210 dataset

We evaluate mfDiMP on the recent RGBT210 dataset [34] using their two evaluation metrics (Figure 3). We compare against the top-10 trackers on this dataset, including CCOT [17], ECO [10], CMRT [35], BACF [26], SRDCF [14], SGT [34], Staple [3], Staple-CA [42], SiameseFC [4]. We can see how our tracker significantly outperforms the second best tracker (DiMP) with an absolute gain of 6.7% and 4.2%, in terms of precision rate and success rate respectively. As a result, mfDiMP achieves a new state-of-the-art also on this dataset, bringing further evidence to the advantages of end-to-end training for multi-modal tracking in terms of accurate object localization.

5.7. Attribute analysis on RGBT210 dataset

There are a total of 12 different attributes in the RGBT210 dataset [34]. We analyze the performance of our method on these attributes in terms of precision rate and success rate (PR/SR %) in Table 3. We compare with some

popular RGB trackers such as ECO [10] and CSR [38]. We also compare with the state-of-the-art RGB-T trackers on this dataset, e.g. CMRT [35] and SGT [34], and some extended RGB-T trackers [35] from RGB modality, e.g. KCF and CFnet. This experiment, which compares trackers on specific scenarios, proves the robustness and generality of our mfDiMP on RGB-T tracking. Our tracker outperforms all other trackers on all attributes but one (thermal crossover). Moreover, in attributes such as partial occlusion, low illumination, deformation, fast motion, camera moving, and background clutter, mfDiMP achieves a significant gain of about 10% in terms of Success Rate (SR) when compared with the second best.

6. Conclusions

Most of the multi-modal trackers are still using hand-crafted features, or simple off-the-shelf deep features. We investigated how to effectively fuse multi-modal data in an end-to-end training manner, which makes optimal use of information from both modalities. We propose three end-to-end multi-modal fusion architectures, consisting of pixel-level fusion, feature-level fusion and response-level fusion. To ensure that the proposed fusion tracker can be trained in an end-to-end manner, we also generated a large-scale paired synthetic RGB-T dataset. We performed extensive experiments on two recent benchmarks: VOT-RGBT2019 [1] and RGBT210 [34]. The results showed that the proposed fusion tracker does significantly improve the performance of the baseline tracker with respect to single modality tracking. As a consequence, our end-to-end multi-modal fusion tracker sets new state-of-the-art results on both datasets.

Acknowledgements. We acknowledge the financial support by the Spanish project TIN2016-79717-R, and mention the Generalitat de Catalunya CERCA Program.

References

- [1] www.votchallenge.net/vot2019. 2, 5, 6, 7, 8
- [2] <https://sites.google.com/view/ahutracking001>. 6
- [3] L. Bertinetto, J. Valmadre, S. Golodetz, O. Miksik, and P. H. Torr. Staple: Complementary learners for real-time tracking. In *CVPR*, pages 1401–1409, 2016. 8
- [4] L. Bertinetto, J. Valmadre, J. F. Henriques, A. Vedaldi, and P. H. Torr. Fully-convolutional siamese networks for object tracking. In *ECCV workshop*, 2016. 1, 2, 6, 7, 8
- [5] G. Bhat, M. Danelljan, L. Van Gool, and R. Timofte. Learning discriminative model prediction for tracking. In *ICCV*, 2019. 2, 3, 5, 6, 7, 8
- [6] S. T. Birchfield and S. Rangarajan. Spatiograms versus histograms for region-based tracking. In *CVPR*, volume 2, pages 1158–1163. IEEE, 2005. 3
- [7] D. S. Bolme, J. R. Beveridge, B. A. Draper, and Y. M. Lui. Visual object tracking using adaptive correlation filters. In *CVPR*, 2010. 1, 2
- [8] C. Ó. Conaire, N. E. OConnor, and A. Smeaton. Thermo-visual feature fusion for object tracking using multiple spatiogram trackers. *Machine Vision and Applications*, 19(5-6):483–494, 2008. 3
- [9] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *CVPR*, 2005. 2
- [10] M. Danelljan, G. Bhat, F. S. Khan, and M. Felsberg. Eco: efficient convolution operators for tracking. In *CVPR*, 2017. 1, 2, 6, 7, 8
- [11] M. Danelljan, G. Bhat, F. S. Khan, and M. Felsberg. Atom: Accurate tracking by overlap maximization. In *CVPR*, 2019. 1, 2, 3, 7, 8
- [12] M. Danelljan, G. Häger, F. Khan, and M. Felsberg. Accurate scale estimation for robust visual tracking. In *BMVC*, 2014. 1, 2
- [13] M. Danelljan, G. Hager, F. Shahbaz Khan, and M. Felsberg. Convolutional features for correlation filter based visual tracking. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2015. 2
- [14] M. Danelljan, G. Hager, F. Shahbaz Khan, and M. Felsberg. Learning spatially regularized correlation filters for visual tracking. In *ICCV*, 2015. 2, 8
- [15] M. Danelljan, G. Hager, F. Shahbaz Khan, and M. Felsberg. Adaptive decontamination of the training set: A unified formulation for discriminative visual tracking. In *CVPR*, 2016. 2
- [16] M. Danelljan, F. S. Khan, M. Felsberg, and J. van de Weijer. Adaptive color attributes for real-time visual tracking. In *CVPR*, 2014. 1, 2
- [17] M. Danelljan, A. Robinson, F. S. Khan, and M. Felsberg. Beyond correlation filters: Learning continuous convolution operators for visual tracking. In *ECCV*, 2016. 1, 2, 8
- [18] 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. *CoRR*, abs/1809.07845, 2018. 2, 3, 6, 8
- [19] M. Felsberg, A. Berg, G. Hager, J. Ahlberg, M. Kristan, J. Matas, A. Leonardis, L. Cehovin, G. Fernandez, T. Vojir, et al. The thermal infrared visual object tracking vot-tir2015 challenge results. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2015. 2
- [20] H. K. Galoogahi, T. Sim, and S. Lucey. Multi-channel correlation filters. In *ICCV*, 2013. 2
- [21] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [22] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista. High-speed tracking with kernelized correlation filters. *TPAMI*, 37(3):583–596, 2015. 1, 2, 6, 8
- [23] L. Huang, X. Zhao, and K. Huang. Got-10k: A large high-diversity benchmark for generic object tracking in the wild. *CoRR*, abs/1810.11981, 2018. 2, 3, 5, 6, 8
- [24] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. 3, 5
- [25] B. Jiang, R. Luo, J. Mao, T. Xiao, and Y. Jiang. Acquisition of localization confidence for accurate object detection. In *ECCV*, pages 784–799, 2018. 3, 4
- [26] H. Kiani Galoogahi, A. Fagg, and S. Lucey. Learning background-aware correlation filters for visual tracking. In *ICCV*, 2017. 2, 8
- [27] M. Kristan, A. Leonardis, J. Matas, M. Felsberg, R. Pflugfelder, L. Č. Zajc, T. Vojir, G. Bhat, A. Lukežič, A. Eldesokey, et al. The sixth visual object tracking vot2018 challenge results. In *ECCV workshop*, 2018. 3, 8
- [28] M. Kristan, J. Matas, A. Leonardis, T. Vojir, R. Pflugfelder, G. Fernandez, G. Nebehay, F. Porikli, and L. Čehovin. A novel performance evaluation methodology for single-target trackers. *TPAMI*, 38(11):2137–2155, Nov 2016. 1
- [29] B. Li, W. Wu, Q. Wang, F. Zhang, J. Xing, and J. Yan. Siamrpn++: Evolution of siamese visual tracking with very deep networks. *CoRR*, abs/1812.11703, 2018. 2
- [30] B. Li, J. Yan, W. Wu, Z. Zhu, and X. Hu. High performance visual tracking with siamese region proposal network. In *CVPR*, 2018. 1, 2
- [31] C. Li, H. Cheng, S. Hu, X. Liu, J. Tang, and L. Lin. Learning collaborative sparse representation for grayscale-thermal tracking. *IEEE Transactions on Image Processing*, 25(12):5743–5756, 2016. 1, 3
- [32] C. Li, X. Liang, Y. Lu, N. Zhao, and J. Tang. Rgb-t object tracking: benchmark and baseline. *CoRR*, abs/1805.08982, 2018. 1, 3
- [33] C. Li, X. Sun, X. Wang, L. Zhang, and J. Tang. Grayscale-thermal object tracking via multitask laplacian sparse representation. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 47(4):673–681, 2017. 1, 3
- [34] C. Li, N. Zhao, Y. Lu, C. Zhu, and J. Tang. Weighted sparse representation regularized graph learning for rgb-t object tracking. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 1856–1864. ACM, 2017. 1, 2, 3, 5, 6, 7, 8
- [35] C. Li, C. Zhu, Y. Huang, J. Tang, and L. Wang. Cross-modal ranking with soft consistency and noisy labels for robust rgb-t tracking. In *ECCV*, pages 808–823, 2018. 1, 2, 3, 4, 8

- [36] 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 *ECCV*, pages 740–755. Springer, 2014. 6
- [37] H. Liu and F. Sun. Fusion tracking in color and infrared images using joint sparse representation. *Science China Information Sciences*, 55(3):590–599, 2012. 1, 3
- [38] A. Lukezic, T. Vojár, L. C. Zajc, J. Matas, and M. Kristan. Discriminative correlation filter with channel and spatial reliability. In *CVPR*, 2017. 1, 2, 8
- [39] C. Ma, J.-B. Huang, X. Yang, and M.-H. Yang. Hierarchical convolutional features for visual tracking. In *ICCV*, 2015. 2
- [40] C. Ma, X. Yang, C. Zhang, and M.-H. Yang. Long-term correlation tracking. In *CVPR*, 2015. 2
- [41] M. Mueller, N. Smith, and B. Ghanem. Context-aware correlation filter tracking. In *CVPR*, 2017. 2
- [42] M. Mueller, N. Smith, and B. Ghanem. Context-aware correlation filter tracking. In *CVPR*, pages 1396–1404, 2017. 8
- [43] M. Muller, A. Bibi, S. Giancola, S. Alsubaihi, and B. Ghanem. Trackingnet: A large-scale dataset and benchmark for object tracking in the wild. In *ECCV*, 2018. 2, 3, 6, 8
- [44] H. Nam and B. Han. Learning multi-domain convolutional neural networks for visual tracking. In *CVPR*, 2016. 2
- [45] E. Park and A. C. Berg. Meta-tracker: Fast and robust online adaptation for visual object trackers. In *ECCV*, 2018. 2
- [46] 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. *IJCV*, 115(3):211–252, 2015. 2
- [47] Y. Song, C. Ma, X. Wu, L. Gong, L. Bao, W. Zuo, C. Shen, R. W. Lau, and M.-H. Yang. Vital: Visual tracking via adversarial learning. In *CVPR*, 2018. 2
- [48] J. Valmadre, L. Bertinetto, J. Henriques, A. Vedaldi, and P. H. Torr. End-to-end representation learning for correlation filter based tracking. In *CVPR*, 2017. 2, 8
- [49] J. Valmadre, L. Bertinetto, J. F. Henriques, R. Tao, A. Vedaldi, A. W. Smeulders, P. H. Torr, and E. Gavves. Long-term tracking in the wild: A benchmark. In *ECCV*, pages 670–685, 2018. 1, 2
- [50] J. Van De Weijer, C. Schmid, J. Verbeek, and D. Larlus. Learning color names for real-world applications. *IEEE Transactions on Image Processing*, 18(7):1512–1523, 2009. 2
- [51] Q. Wang, Z. Teng, J. Xing, J. Gao, W. Hu, and S. Maybank. Learning attentions: residual attentional siamese network for high performance online visual tracking. In *CVPR*, 2018. 2
- [52] Y. Wu, E. Blasch, G. Chen, L. Bai, and H. Ling. Multiple source data fusion via sparse representation for robust visual tracking. In *14th International Conference on Information Fusion*, pages 1–8. IEEE, 2011. 1, 3
- [53] Y. Wu, J. Lim, and M.-H. Yang. Object tracking benchmark. *TPAMI*, 37(9):1834–1848, 2015. 1, 3
- [54] X. Yu and Q. Yu. Online structural learning with dense samples and a weighting kernel. *Pattern Recognition Letters*, 105:59–66, 2018. 2
- [55] X. Yu, Q. Yu, Y. Shang, and H. Zhang. Dense structural learning for infrared object tracking at 200+ frames per second. *Pattern Recognition Letters*, 100:152–159, 2017. 2
- [56] K. Zhang, L. Zhang, Q. Liu, D. Zhang, and M.-H. Yang. Fast visual tracking via dense spatio-temporal context learning. In *ECCV*, 2014. 2
- [57] L. Zhang, D. Bi, Y. Zha, S. Gao, H. Wang, and T. Ku. Robust and fast visual tracking via spatial kernel phase correlation filter. *Neurocomputing*, 204:77–86, 2016. 2
- [58] L. Zhang, A. Gonzalez-Garcia, J. van de Weijer, M. Danelljan, and F. S. Khan. Synthetic data generation for end-to-end thermal infrared tracking. *IEEE Transactions on Image Processing*, 28(4):1837–1850, 2019. 2, 4, 5
- [59] T. Zhang, C. Xu, and M.-H. Yang. Learning multi-task correlation particle filters for visual tracking. *TPAMI*, 41(2):365–378, 2019. 2
- [60] G. Zhu, F. Porikli, and H. Li. Beyond local search: Tracking objects everywhere with instance-specific proposals. In *CVPR*, 2016. 2
- [61] Z. Zhu, Q. Wang, B. Li, W. Wu, J. Yan, and W. Hu. Distractor-aware siamese networks for visual object tracking. In *ECCV*, 2018. 1, 2, 7, 8