Part-based Tracking by Sampling

George De Ath   Richard M. Everson
University of Exeter, United Kingdom
{g.de.ath,r.m.everson}@exeter.ac.uk

Abstract

We propose a novel part-based method for tracking an arbitrary object in challenging video sequences. The colour distribution of tracked image patches on the target object are represented by pairs of RGB samples and counts of how many pixels in the patch are similar to them. Patches are placed by segmenting the object in the given bounding box and placing patches in homogeneous regions of the object. These are located in subsequent image frames by applying non-shearing affine transformations to the patches' previous locations, locally optimising the best of these, and evaluating their quality using a modified Bhattacharyya distance. In experiments carried out on VOT2018 and OTB100 benchmarks, the tracker achieves higher performance than all other part-based trackers. An ablation study is used to reveal the effectiveness of each tracking component, with largest performance gains found when using the patch placement scheme.

1. Introduction

The tracking of arbitrary objects, also know as model-free tracking, has applications across many different fields including video surveillance, activity analysis, robot vision, and human-computer interfaces. The goal of model-free tracking is to determine the location of an unknown object, specified only by a bounding box in the first frame of a sequence, in all subsequent frames. Although the topic has attracted much recent research interest and great strides have been made in tracking performance in recent years [31], visual object tracking still remains a challenging problem because of the many difficult real-world tracking scenarios that can be encountered. These include camera motion, illumination change, object motion and size change, as well as occlusion and self-occlusion [33].

The way an object is represented, *i.e.* its model, is arguably the most important component in a tracking algorithm. Objects can either be represented in a holistic manner [59, 5, 54, 24, 27, 46, 15, 16], with a part-based representation, [7, 3, 63, 20, 21, 42, 10, 66, 4, 34, 44, 43], or

as a combination of the two [11, 58, 12]. Holistic methods represent the object using one global model that characterises the entire region that the object resides in, typically its bounding box. However, if large amounts of deformation or occlusion occur then global models can fail to robustly track the target [21]. A promising way of countering these types of problems is through part-based methods, which use smaller, localised models to represent sub-regions of the object, known as *parts* or *patches*, which together can then be used to estimate the object's overall location.

Here we construct a novel part model of the tracked object, drawing inspiration from ViBe [6], a non-parametric pixel-based algorithm that is one of the simplest and most effective background subtraction techniques [60]. It models each pixel in a image by storing samples of their values over time, and using the set of these samples to model the pixel's color distribution. We reverse the background subtraction problem and model the image's foreground (*i.e.* the object) rather than its background. The pixel-based representation is extended to a part-based one by using samples of features (usually colours) taken from image patches to characterise the distribution of a part's features. We also develop an object localisation scheme to find the modelled object's parts in subsequent frames, along with an update scheme to update the models once their locations have been predicted.

A commonly overlooked areas of model-free tracking is the problem of determining which pixels, within the given bounding box, actually correspond to the object [18]. For trackers that use a global model, such as correlation filters, this is less of a concern because the majority of the region given to the tracker (typically over 70% [56]) belongs to the object, resulting in a small proportion of spurious pixels in their model(s). However, for part-based trackers, the placement of their parts is of paramount importance because if a part is initialised on the background of the object then this part will continue to match a stationary region in the image, rather than the moving object. This can lead to some parts trying to follow the object, while others remain static, "tracking" part of the background, resulting in a reduction in tracking performance. We tackle this using a novel object part placement scheme.

In summary, we propose a part-based visual tracking framework which has the following novel contributions:

- A sparse, part-based visual model that empirically characterises the underlying colour distribution of the image patches with clusters in colour space.
- An object part placement mechanism that, given a bounding box containing an object, segments the object and selects part locations in homogeneous regions on the object.
- An object localisation method that models an object's motion using global non-shearing affine transformations followed by localised patch location optimisation, mirroring typical object movement.

The remainder of the paper is organized as follows: §2 provides a review of related methods, §3 describes our tracking method in detail, and in §4 we present experimental results on the VOT2018 and OTB100 datasets, including an ablation study to highlight the contribution of each component of the tracker. Conclusions are drawn in §5.

2. Related Work

We review the main components of part-based tracking methods, focussing on how parts are placed on objects, the way in which parts are represented and inter-part constraints. General surveys of visual object tracking methods may be found in [32, 13, 47, 52, 39, 65, 51, 61, 45, 64, 48].

2.1. Object Part Placement

Part-based methods usually select the locations of the patches representing parts in an *ad hoc* fashion, either distributing them uniformly across the object, or trying to select regions that have good properties in relation to their part model. Some place their constituent parts uniformly on a grid across the bounding box [11, 29, 63], while others place patches so that they touch each other and cover the entire bounding box [3, 62, 42, 22, 2]. Those that place their patches in this fashion make no assumptions as to which pixels within the bounding box belong to the object.

Other techniques place their parts based on the type of part-model they have. The ANT tracker [12], for example, chooses patch locations that maximise the chance of having good optical flow characteristics. Methods that represent object parts as superpixels [10, 20, 21] select superpixels that lie wholly within the bounding box (*i.e.* have no pixels overlapping it). However, [58] also selecting superpixels that overlap the bounding box, because since superpixeling segments an image into homogeneous regions, any superpixels that contain background are likely to contain only background pixels, meaning that they can quickly prune the poorly matching superpixels later in tracking.

PBTS uses superpixels to segment the initial bounding box, but we first label pixels as object or background, using only the initial image and the bounding box [18]. We then superpixel the pixels in the bounding box that have a high likelihood of belonging to the object, thereby drastically reducing the chance of placing patches on pixels that do not belong to the object.

2.2. Object/Part Representation and Matching

Parts are generally represented using either keypoints, correlation filters, or histograms of features extracted from the parts. Keypoint methods represent an object at a set of distinguished feature points. Keypoints are characterised by the local image texture using, *e.g.*, ORB [50], SIFT [41], or SURF [8]. The approximate invariance to image scale and rotation of these descriptors allows them to be matched to the corresponding points in the subsequent frame.

Correlation filters match regions in successive frames by maximising the correlation between an image patch in two frames. They are effective and efficient in object tracking [31], but the range over which the correlation is maximised is limited to half the width or height of the template being matched [42], limiting their use for matching parts, which are generally small compared to the object movement.

Histograms capture the feature distribution of the patch corresponding to a part. Commonly used features are the RGB pixel intensities, but they have been used widely with a variety of features, *e.g.* [11, 42, 34]. Patch similarity, expressing the overlap between the probability distributions estimated by the histograms, can be computed in a variety of ways, methods based on the Bhattacharyya distance [14, 9] being popular. Here we use a novel object patch representation and compare the similarity of patches using a modified version of the Bhattacharyya distance.

2.3. Part Constraints and Localisation

The way in which part-based methods relate their parts to one another geometrically plays an important role in their tracking performance. Some enforce no explicit geometric relationships [34, 58], others a star-based topology [10, 37, 43, 66], and some have a fully-connected sets of parts [4, 42]. Intermediate between these are local constraints between neighbouring parts, *e.g.* [20, 11, 21].

Implicit constraints on a patch's location may also result from limiting the range of potential patch locations searched in subsequent frames, often by modelling part motion between frames by global affine or rigid transformation, after which part locations are individually optimised in a restricted region around the global transformation [63, 11].

In PBTS we also implicitly constrain the potential locations patches to be evaluated, based on our observation that an object's motion between frames is largely rigid and only varies slightly from this rigidity locally. We therefore generate a set of candidate global non-shearing affine transformations and apply them to each part's location in the previous
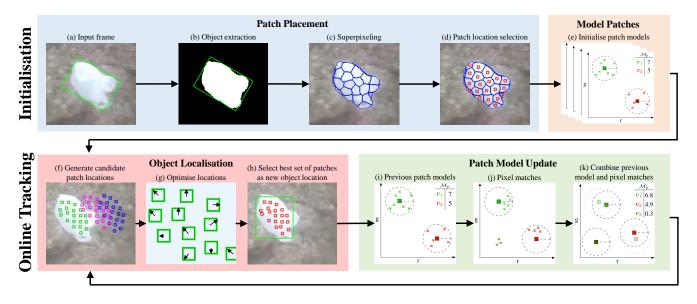
Figure 1. Schematic diagram of PBTS. For ease of illustration in (e) and (i-k) pixels are represented as (r,g) pairs rather than (r,g,b) triples.

frame. The best of these is then used as the base transformation for individual optimisations of each part's location in a small window around the best globally transformed location, thus allowing a localised non-rigidity.

3. Part-based Tracking by Sampling

Fig. 1 gives an overview of PBTS, which we first describe in outline before giving details about each stage.

The tracker is initialised in the first frame of a video sequence from the supplied bounding box containing the object to be tracked (Fig. 1a). As described in §3.5, pixels likely to belong to the object in the bounding box are identified (Fig. 1b) and superpixelled (Fig. 1c), and a patch/part chosen at the centre of each superpixel (red squares in Fig. 1d). Patches are represented by their centres $\mathbf{x}_p$ and a model $\mathcal{M}_p$ of the colour distribution of pixels in the patch. As described in §3.1, $\mathcal{M}_p$ comprises of samples from the patch's pixels together with counts describing the number of similar pixels in the patch, Fig. 1e. In subsequent frames the patches are transformed from their locations in the previous frame by randomly chosen global non-shearing linear transformations to candidate locations; represented by green, magenta and blue squares in Fig. 1f. The most likely of these is determined by matching the colour models of patches in the previous frame to their transformed locations, after which the individual patch locations are locally optimised (Fig. 1g). Finally, the patch colour models are updated (§3.4). In outline, the sparse colour models are updated by moving the features in the old colour model towards the matching pixels in the new patch, and introducing new features to represent previously unseen pixel values or pruning now redundant features (Figs. 1i – 1k).

3.1. Object Model

We represent an object $\mathcal{O}$ as an ordered set of P patches $\mathcal{O} = \{\mathcal{P}_p\}_{p=1}^P$. Patches are rectangular with width and height of w_x and w_y pixels respectively, and each patch $\mathcal{P}_p = (\mathbf{x}_p, \mathcal{M}_p)$ is characterised by its location $\mathbf{x}_p$ and a colour model $\mathcal{M}_p = \{(\mathbf{c}_1, h_1), ..., (\mathbf{c}_S, h_S)\}$, containing S pairs of centres $\mathbf{c}_s$ and their corresponding match counts h_s; see Fig. 1e. Centres are (r, g, b) samples from pixels in the patch and the counts h_s indicate how many pixels within the patch match the centre $\mathbf{c}_s$. A pixel $\mathbf{I}_i = (r_i, g_i, b_i)$ is said to match the centre if $\|\mathbf{I}_i - \mathbf{c}_s\| < R$ for radius R. Thus the colour model can be thought of a sparse empirical histogram of the frequency of colours in the patch. Although we describe the colour model in terms of R, G, B colour descriptors the method naturally generalises to other colour spaces or more sophisticated colour or texture descriptors.

The colour model is initialised from the pixels $\mathbf{I}_i$ in the patch region, which we denote by Ω_p. We start by randomly selecting a pixel $\mathbf{I}_i$ as the first sample in the model $\mathbf{c}_1 = \mathbf{I}_i$, with match count $h_1 = 1$. Successive pixels are randomly selected, without replacement, and compared to the centres in the model. If the pixel matches, the count of the centre it is closest to is incremented by one: $h_{s'} := h_{s'} + 1$ where $s' = \operatorname{argmin}_{s'} \|\mathbf{I}_i - \mathbf{c}_{s'}\|$. If it does not match any centre (i.e. $\|\mathbf{I}_i - \mathbf{c}_s\| > R \ \forall s = 1, \ldots, S$) then that pixel is added to the model as a new centre: $\mathbf{c}_{S+1} = \mathbf{I}_i$ with $h_{S+1} = 1$. This process is repeated until all pixels in the patch have been evaluated once. Finally, if there are more than a predetermined number centre-count pairs S_{max}, then the pairs with the lowest counts are removed until the model only contains S_{max} pairs of centres and counts.

This initialisation yields a kind of pseudo-clustering,

without the need for an explicit clustering algorithm, tailored to producing clusters of a fixed radius (R), centred on pixel samples. This stochastic sampling method is much less computationally intensive than traditional clustering techniques. Since the models are initialised from relatively homogeneous patches and samples are therefore close in colour space, it yields similar cluster centres regardless of the sampling sequence. Limiting the number of centres to S_{max} in each model effectively reduces model over-fitting by excluding rare unrepresentative samples.

Unlike histograms that partition colour space into rectangular bins with predetermined centres on a rectangular grid, this model makes no assumptions as to how the colour space should be partitioned. Instead, the use of colour samples themselves as bins yields a sparse representation as, unlike conventional histograms, there are no empty bins.

3.2. Patch Similarity

Fundamental to PBTS is evaluating how well a candidate patch $\tilde{\mathcal{P}}$ in the current frame matches $\mathcal{P}$ without building a new model for the candidate patch. The quality of the match is evaluated by comparing how well pixels in $\tilde{\mathcal{P}}$ match $\mathcal{M}_p$. This is achieved by counting the number of pixels in $\tilde{\mathcal{P}}$ that are within radius R of each centre $\mathbf{c}_s$ in $\mathcal{M}_p$.[1] Define $\tilde{\mathbf{h}}$ to be the vector of match counts for pixels in $\tilde{\mathcal{P}}$ to the S centres of $\mathcal{P}_p$ and let $\mathbf{h}$ be the normalised vector of counts for $\mathcal{P}$ itself. Both vectors are normalised by dividing by the number of pixels in the patch. Then the similarity between $\mathcal{P}$ and $\tilde{\mathcal{P}}$ can be quantified using a modification of the Bhattacharyya distance [14]. The Bhattacharyya coefficient [9] between two distributions, $\mathbf{p}$ and $\mathbf{q}$, where $p_j \geq 0$, $q_j \geq 0$, $\sum p_j = \sum q_j = 1$ is given by:

$$BC\left(\mathbf{p}, \mathbf{q}\right) = \sum_{j=1}^{S} \sqrt{p_j q_j}. \qquad (1)$$

Note that $0 \leq BC\left(\mathbf{p}, \mathbf{q}\right) \leq 1$ measures the overlap between two distributions, $\mathbf{p}$ and $\mathbf{q}$, with $BC\left(\mathbf{p}, \mathbf{q}\right) = 0$ iff there is no overlap. We define the modified version of the Bhattacharyya Distance (MBD) to be

$$MBD\left(\mathbf{p}, \mathbf{q}\right) = [1 - BC\left(\mathbf{p}, \mathbf{q}\right)]^b \qquad (2)$$

where $b \geq 0$ controls the weight given to good matches. A larger value of b favours good matches, whereas they are down-weighted with smaller values. Setting $b = \frac{1}{2}$ recovers the original Bhattacharyya distance [14], which is also equal to the Hellinger distance [25].

Using the modified Bhattacharyya distance we define the quality of the match between $\mathcal{M}$ and $\tilde{\mathcal{P}}$ as

$$Q(\mathcal{P}, \tilde{\mathcal{P}}) = 1 - MBD(\mathbf{h}, \tilde{\mathbf{h}}). \qquad (3)$$

The overall quality of a set of a candidate parts $\tilde{\mathcal{O}} = \left\{ \tilde{\mathcal{P}}_p \right\}_{p=1}^{P}$ representing an object is thus the average similarity of the individual patches to the object being matched:

$$L(\mathcal{O}, \tilde{\mathcal{O}}) = \frac{1}{P} \sum_{p=1}^{P} Q(\mathcal{P}_p, \tilde{\mathcal{P}}_p). \qquad (4)$$

Note that after normalization by the number of pixels in a patch, the vector of match counts $\sum \tilde{h}_j \leq 1$ because only the pixels that match each centre in $\mathcal{M}_p$ are counted. This is a desirable property as it will implicitly down-weight candidate patches that have fewer pixels matching the model.

3.3. Localisation

Empirically we observe that object movement between consecutive frames is largely rigid and it only has minor local deviations from rigidity [17]. More precisely, we find that object movement is well represented by a rotation, an isotropic scaling and a translation. We therefore match patches in a frame to those in the previous frame in two steps: we first generate non-shearing affine transformations of the previous frame's patch locations to match the rigidity assumption, and then optimise each patch's location within a small window around the best rigid transformation to allow for local non-rigidity; see Figs. 1f and 1g.

G non-shearing affine transformations $\mathbf{A}_g = \mathbf{T}(x,y)\mathbf{S}(s)\mathbf{R}(r)\mathbf{O}$ are generated as the products of a translation to the origin $\mathbf{O}$, a rotation $\mathbf{R}(r)$ by r radians, an isotropic scaling $\mathbf{S}(s)$ by a factor s and a translation $\mathbf{T}(x,y)$ by $[x, y]^T$. The parameters are drawn from random distributions centred on the patch location in the previous frame, thus $r \sim \mathcal{N}(0, \sigma_r)$ and $s \sim \mathcal{N}(1, \sigma_s)$ where $\sigma_r = \pi/16$ and $\sigma_s = 0.02$ were chosen to match the typical rotations and scale changes found in an extensive survey of the VOT2016 data [56, 17]. We observe that the distributions of inter-frame vertical and horizontal translations are heavy-tailed and therefore draw translation parameters from Laplace distributions centred on the object location in the previous frame and with scale parameters $w'\sigma_x$ and $h'\sigma_y$ where w' and h' are the predicted width and height of the object in the previous frame and $\sigma_x = 0.15$ and $\sigma_y = 0.1$ were determined empirically [17]. Scaling the translation distributions' length-scales by the object's predicted width and height allows for the predicted movement of the object to be adjusted relative to its size, because generally, larger objects move further than smaller objects, but are comparatively similar when this movement is considered as a proportion of their own size.

The patch locations resulting from the best L candidates of the G randomly generated transformations are each locally optimised by an exhaustive search of potential patch locations within in a square window with side length W pixels, centred on the transformed patch at $\tilde{\mathbf{x}}_p = \mathbf{A}_g \mathbf{x}_p$. The

---

[1]If a pixel matches more than one centre in $\mathcal{M}_p$, (because two centres lie within $2R$ of one another) the centre closest to the pixel is matched.

match quality of the patch centred at each location within the window is evaluated and the location with the highest quality is selected as the patch's new predicted location; Fig. 1g. If there are multiple locations with the same match quality then the one closest to the globally transformed location $\tilde{\mathbf{x}}_p$ is selected, with equidistant ties broken randomly.

Following preliminary experiments, we report the predicted bounding box of the object as the axis-aligned bounding box (AABB) 20% wider and taller than the AABB that minimally encloses the patches after matching; Fig. 1h.

We note that while the methods of [63] and [11] share some similarities with our object localisation scheme, there are several important differences. In contrast to [63], who, after applying an affine transformation to the set of patches, randomly move each patch and evaluate its quality, we locally optimise within a small region around the affine transformed position. [63] also limit the class of transformations to only include translation, whereas we also include both isotropic scaling and rotation. The localisation scheme of [11] uses the cross-entropy method to first find the optimal affine transformation and then locally optimises each patch. Our methods differs from this as we locally optimise the L sets of patch locations with the highest quality, each set of which can have different patch locations to other sets.

[11] start a search for the optimal affine transform by sampling from a Gaussian distribution with a covariance of $20\mathbf{I}$ for all sizes of objects, similar to [63] who also sample from a Gaussian distribution with a scale parameter of 8 pixels for both horizontal and vertical translation. In contrast, we sample translations drawn from distributions whose scale parameter is relative to the respective width or height of the object, as we have observed that the size of inter-frame object motion is approximately proportional to the object's size.

3.4. Model Update

In common with other models that represent the colour distribution with histograms such as [2, 12, 34], we linearly interpolate the probability of finding a colour in a patch between the original model and the newly matched image region. Our sparse representation of the colour histogram means that in addition to updating the centre-counts pairs $(\mathbf{c}_i, h_i)$, we need to be able to introduce new and remove redundant centre-count pairs.

Let Ω_p be the set of pixels comprising the image region that matches model $\mathcal{M}_p$ and let $\omega_{p,s} = \{ \mathbf{I} \in \Omega_p \mid \|\mathbf{I} - \mathbf{c}_s\| < R \}$ be the set of pixels in Ω_p which match the s-th centre in the model (*i.e.* pixels inside the dashed circles in Fig. 1j).

Then the model counts are updated by linearly interpolating:

$$h_{p,s} := \beta_c |\omega_{p,s}| + (1 - \beta_c) h_{p,s} \qquad (5)$$

where β_c is an update rate. Similarly, the locations of the centres are updated to move them towards the mean of the matched pixels:

$$\mathbf{c}_s := \beta_s \frac{1}{|\omega_{p,s}|} \sum_{\mathbf{I}_j \in \omega_{p,s}} \mathbf{I}_j + (1 - \beta_s)\mathbf{c}_s \qquad (6)$$

The use of two update rates, β_c and β_s, allows for the centres and counts to be adapted at different speeds. This is an important property as when an image region changes colour due to illumination intensity the brightness of each colour in the region increases, rather than the relative proportion of colours. Updating $\mathbf{c}_j$ has the effect of moving the centre of the sphere containing matching pixels towards (or past if $\beta_s > 1$) the matches' centre of mass. It allows the model to *follow* the region of colour space that it characterises as it changes over time. In the standard histogram-based approach, when the centres change over time they move from matching one bin to matching another. When this occurs there will be a sudden loss of probability mass within the histogram, resulting in poorer matching. Following extensive experimentation we found that $\beta_s = 1.7$ yields the best tracking performance, meaning that the update is predicting where the colour distribution of the patch will lie in the subsequent frame.

New colour centres corresponding to pixels Ω_p^- which do not match the original colour model (*i.e.* pixels outside the dashed circles in Fig. 1j) are incorporated in $\mathcal{M}_p$ by first creating a new patch model $\mathcal{M}_p^-$ from them. Each of the counts h_s^- is scaled by β_c and the centre-count pairs in $\mathcal{M}_p^-$ are added to $\mathcal{M}_p$; *e.g.* $(\mathbf{c}_3, h_3)$ in Fig. 1k. The scaling by β_c is done in order to match the behaviour of count updating, as this introduces the centre-count pair as though they were already in the model with a previous count of 0; cf. Eq. (5).

Lastly, the centre-count pairs with $h_s < \beta_c$ are removed from $\mathcal{M}_p$. This allows for centres that have not been seen recently to be removed in order to limit the model's size and computational complexity. A threshold of β_c was chosen as it is the smallest value that the count for a new centre-count pair can have if it were added to the model at the same time-step, *i.e.* it had a value of $h_s = 1$ before scaling.

3.5. Patch Placement

Since our colour model provides a compact representation of the colours in a patch, we seek to select spatially compact regions of homogeneous colour as patches. This is achieved by superpixelling [49] the supplied bounding box; see Fig. 1g. Superpixelling over-segments images into perceptually meaningful regions that are generally uniform in colour and texture. As superpixels tend to adhere to colour and shape boundaries, they also retain the image's structure [53]. As illustrated in Fig. 1d, object parts are placed at the centre of superpixels that reside within the object's predicted location as determined by the alpha-matting segmentation scheme described in [18].

A parameterless version (SLICO[2]) of the SLIC superpixelling algorithm [1] is used to segment the region surrounding the predicted location of the object into approximately P superpixels. The zero-parameter version of the state-of-the-art SLIC [53] allows for each superpixel to have its own compactness parameter, so that regular shaped superpixels are generated for both smooth and rugged image regions.

The location of the first patch is taken as the centroid of the largest superpixel. Successive patches are then greedily initialised to the centroid of the largest unassigned superpixel if doing so would cause the patch to overlap with all other patches with a proportion of its area less than γ. This process is repeated until either P patch locations have been selected or there are no more superpixels left to consider.

The degree of permitted overlap γ between patches controls the density of the patches across the object. For larger objects, where patches do not reach the boundaries of superpixels, its value is immaterial. However, for smaller objects, patches may overlap multiple superpixels, if the area of the superpixels is less than that of the patches' area, or when the superpixels are roughly rectangular with dimensions shorter or longer than the patches. Limiting the patch overlap in these cases reduces the amount redundant information in neighbouring patches.

An alternative would be to initialise patches on the boundaries between superpixels where there tends to be high contrast so that patches would be centred on image features such as corners or "keypoints". However, experiments show that initialisation in homogeneous regions results in better tracking performance.

4. Experiments

The tracker was evaluated on the VOT2018 benchmark [31], which provides a tracking dataset with fully annotated frames, and reports the performance of a large number of state-of-the-art trackers. The dataset comprises of 60 sequences, containing difficult tracking scenarios such as occlusion, scale variation, camera motion, object motion change, and illumination changes. We followed the VOT challenge protocol, with the tracker initialised on the first frame of a sequence using the ground-truth bounding box provided, and reinitialised if the tracker drifted away from the target. Trackers were evaluated in terms of their accuracy (target localisation), robustness (failure frequency), and Expected Average Overlap (EAO). For full details see [31]. PBTS was compared to the 10 best part-based trackers in VOT2018 as well as the top 3 state-of-the-art trackers.

The tracker was also evaluated on the OTB100 benchmark [57] which can be viewed as complementary to the VOT evaluation as its main focus is on unsupervised tracking. However, trackers which explicitly include re-detection

mechanisms are at an advantage when compared to the those trackers evaluated on the VOT benchmark, which generally do not include any long-term tracking components. We compare the performance of PBTS to the top-performing trackers on 75 colours sequence in OTB100 since PBTS uses only colour information (RGB features). PBTS is compared with other trackers using their published results[3] on the same 75 colour videos. Performance is measured using success and precision plots, which are measure-threshold plots of the proportion of tracked frames that have an intersection-over-union (IOU) overlap greater than some threshold and the proportion of frames that have a centre error (the distance between the centroids of the predicted and ground-truth bounding boxes) of less than a threshold.

We also report an ablation study which highlights the contribution of each component of the tracker, together with a qualitative analysis of success and failure modes.

4.1. Implementation Details

PBTS was implemented in Python 3.6 on an Intel i5-4690 CPU with 32Gb RAM, and runs at ≈ 15 frames per second, calculated on the VOT2018 sequences with no GPU or multi-threading. Its parameters were chosen via extensive cross-validation experimentation [17].

Object Representation: Objects were modelled using $P = 35$ patches of size $w_x = w_y = 5$ pixels, with the colour model using RGB pixel features. A matching distance of $R = 20$ [6] was used, with a MBD coefficient value of $b = 1.4$ used for up-weighting better matching patches. Patch update parameters were $\beta_c = 0.05$ and $\beta_s = 1.7$.

Patch Placement and Model Initialisation: Using the notation of [18], we set the object segmentation parameters as $\rho^- = 0.8$, $\rho^+ = 1.2$, $\tau = 0.85$, and $\lambda = 10^{-2}$. Patches were placed on objects with a maximum patch area overlap of $\gamma = 0.25$, and the number of samples limited in each model to $S_{max} = 10$ during initialisation. Performance was found to be roughly constant for $S_{max} \in [8, 20]$.

Object Localisation: Parameters for the random generation of non-shearing affine transformations are given in §3.3. $G = 1000$ transforms were sampled, with the best $L = 100$ locally optimised. We found no performance gain by increasing L or G. During local optimisation a patch may move in any direction 2 pixels per time-step ($W = 5$).

4.2. Evaluation on the VOT2018 Dataset

Table 1 shows the results of PBTS compared with the top 10 part-based trackers and top 3 state-of-the-art trackers on the VOT2018 benchmark.[4] Results for other trackers are taken from [31]. The average overlap (AO) and robustness (R) scores are averaged over each video, and so

---

[2]https://ivrl.epfl.ch/research-2/
research-current/research-superpixels/#SLICO

[3]http://cvlab.hanyang.ac.kr/tracker_benchmark/

[4]Results previously reported for PBTS in [31] did not include any centre-count updating ($\beta_c = 0$) and used a MBD weighting of $b = 1/2$.

Tracker	EAO	AO	R	Tracker	EAO	AO	R
PBTS	0.196	0.427	1.723	BDF [44]	0.093	0.336	4.200
ANT [12]	0.168	0.439	2.250	Matflow†	0.092	0.362	4.550
DPT [42]	0.158	0.463	2.567	FragTrack [2]	0.068	0.325	6.650
LGT [11]	0.144	0.403	2.641	Matrioska [43]	0.065	0.365	6.900
DFPReco†	0.138	0.467	2.983	LADCF [59]	0.389	0.507	0.567
FoT [55]	0.130	0.385	3.667	MFT [5]	0.385	0.496	0.500
BST [7]	0.116	0.244	3.136	SiamRPN [38]	0.383	0.567	0.900

Table 1. Results of PBTS on the VOT2018 benchmark compared to part-based trackers and state-of-the-art trackers (lower-right). Note that † indicates that details of the tracker are only published in the VOT2018 benchmark [31].

PBTS	Default MBD	No local optimisation	No model update	No Alpha matting	Uniform Patch placement
EAO	0.196	0.179	0.169	0.162	0.127	0.118
AO	0.427	0.353	0.396	0.386	0.245	0.234
R	1.723	1.658	2.346	2.194	1.500	1.591

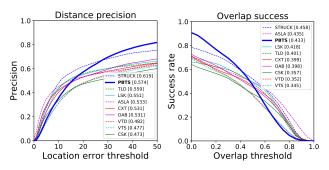Table 2. Ablation study using the VOT2018 benchmark.

Figure 2. Precision and success plots using one-pass evaluation on the OTB100 dataset. The precision of the tracker with a 20 pixel threshold and the area under the success plot are shown in brackets in the left and right-hand plots respectively. Better trackers have distance precision curves towards the top-left of the plots and overlap success curves towards the top-right of the plots.

are per-video scores. As can be seen, PBTS outperforms all other part-based trackers in terms of both EAO and robustness, and places fourth in terms of accuracy. Given the relative simplicity of the tracker, compared to those using global models (*e.g.* ANT [12], DPT [42], LGT [11]) to guide tracking in addition to their part-based, local models, the PBTS methodology performs well. However, compared to the current state-of-the-art methods (those in the lower-right portion of the table) our tracker performs considerably worse with respect to all three attributes. We suggest that this is due to both the tracker architecture and the types of features used for object representation. The 10 top-performing trackers on the benchmark [31] all use deep neural network features (combined with various other handcrafted features) in either correlation filter-based or CNN-based frameworks.

4.3. Evaluation on the OTB100 Dataset

We compare PBTS to the top-performing trackers on the OTB100 OPE benchmark, namely STRUCK [24], CXT [19], ASLA [28], TLD [30], CSK [26], LSK [40], VTD [35], OAB [23], and VTS [36]. The performance of PBTS over all colour sequences is shown in Fig. 2. Overall, PBTS performs well on the benchmark, achieving second

and third place in terms of precision and overlap respectively. The general shape of the overlap success plots (Fig. 2 (right)) of PBTS tends to be more *S*-shaped than the majority, which we suspect is due to the underestimation of the bounding box; this is discussed further in §4.5. PBTS generally achieves a higher level of precision than the other trackers once the location error threshold greater than approximately 25 pixels. We believe that this is due to the robustness of the part-based formulation of the tracker, as demonstrated by its performance on the VOT2018 benchmark. Even if several object parts have drifted, the majority will continue to track the object, despite the fact that the drifted parts may be a significant distance from the main group of parts tracking the object, leading to a good level of a precision at the expense of overlap. This means that the predicted bounding box's centre may be away from the centre of the object, but closer than typical holistic trackers (which model the entire bounding box using one appearance model). In the holistic case, if the tracker begins to model a region containing a large amount of background, the entire tracker will tend to drift off the object, resulting in failure.

4.4. Ablation Study

We performed a component ablation study using the VOT2018 benchmark to evaluate the contribution of each key component of PBTS. Table 2 shows the performance of varying one component, keeping all others fixed. The label *No alpha matting* indicates that the procedure of [18] is not used and the entire bounding box is predicted as containing the object (the default assumption of most trackers). *Uniform patch placement* indicates that the entire patch placement scheme was not used; instead patches were placed uniformly over the bounding box, similarly to [11, 29, 63]. *Default MBD* refers to using $b = \frac{1}{2}$ in Equation 1, relatively up-weighting the contribution of poorer performing patches compared to the optimised value of $b = 1.4$. *No local optimisation* denotes PBTS with $L = 0$, restricting transformations to only search non-shearing, affine transformations. *No model update* denotes setting $\beta_s = \beta_c = 0$ and using the patch models initialised in the first frame throughout.

Table 2 shows that each component is important to the overall performance. In particular, the initialisation via object segmentation and placing patches at the centre of superpixels has the most influence on tracking performance, because this avoids placing patches on relatively static background pixels which are subsequently tracked. Using the Bhattacharyya distance parameter $b = \frac{1}{2}$, gives the least
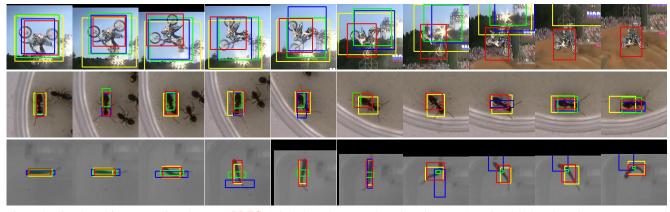
Figure 3. Visual tracking comparison between **PBTS** and the best three part-based trackers on the VOT2018 benchmark, ANT, DPT, and LGT, on the sequences *motocross1*, *ants1*, and *zebrafish1* (first, second, and third rows respectively). Frames in which there is no bounding box for a tracker indicate that it failed recently and will be reinitialised.

reduction in performance, but shows that the modified distance is a valuable addition.

4.5. Qualitative Analysis

Fig. 3 shows qualitative tracking results for PBTS compared to ANT [12], DPT [42], and LGT [11]. In the *motocross1* sequence, only PBTS is able to cope with in-plane and out-of-plane rotation, as well as large flashes of light from the floodlights (right of images). This is largely due to the enhanced PBTS patch placement procedure, which avoids placing patches on background pixels (*e.g.* sky), enabling PBTS to track the object rather than the background. This appears to be the reason for the large expansions of the other trackers' bounding boxes in the initial frames.

The *ants1* sequence shows PBTS coping well with rotation and a non-convex object, whereas both ANT and DPT fail to track the ant as it quickly rotates $90°$. LGT also tracks the ant without failure but starts to drift as the ant rotates. Objects with large aspect ratios, such as the fish in the *zebrafish1*, highlight problems with trackers that fail to estimate object rotation correctly. ANT and DPT incorrectly estimate the object's rotation, resulting in failure for DPT and a severe miscalculation of the object's scale by ANT.

Three failure cases of the PBTS tracker are illustrated in Fig. 4. Due to a poor segmentation in the *bolt1* video, some patches are initialised partially on the background, which leads to those patches tracking the background, while other patches, initialised on the runner's torso when it was in shade, match poorly as he moves into the light and they too start to drift. In the *rabbit* sequence the rabbit is very similar in appearance to its background, leading to patches drifting off the rabbit as it moves across the snow. As PBTS relies solely on a colour model it struggles to track objects that are very similar to their background. In *gymnastics2* the gymnast rotates out of plane, undergoes large amounts of deformation, has large amounts of motion blur, and is filmed
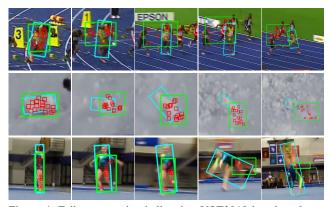

Figure 4. Failure cases in challenging VOT2018 benchmark sequences (*bolt1*, *rabbit*, and *gymnastics2*). Tracked object parts are shown in red, with the predicted and ground-truth bounding boxes shown in green and cyan respectively.

on a hand-held, moving camera. These challenging visual attributes combined to rapidly change the object's appearance, resulting in the tracker underestimating the object's scale and eventually losing track of the object completely.

5. Conclusion

PBTS is a novel part-based tracking framework for short-term, model-free tracking. Essential to its performance is the patch placement mechanism that attempts to place patches on the object being tracked, avoiding tracking the background, and the observation that inter-frame movement is on the whole rigid with small local deviations, which permits efficient global followed by local search for the optimum transformation. The sparse colour model, characterising the patch, shows surprisingly good performance for such a simple representation and we anticipate that performance can be improved by employing *e.g.* texture or deep convolutional features.

References

[1] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk. SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 34(11):2274–2282, 2012. 6

[2] A. Adam, E. Rivlin, and I. Shimshoni. Robust fragments-based tracking using the integral histogram. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 798–805, 2006. 2, 5, 7

[3] O. Akin, E. Erdem, A. Erdem, and K. Mikolajczyk. Deformable part-based tracking by coupled global and local correlation filters. *Journal of Visual Communication and Image Representation*, 38:763–774, 2016. 1, 2

[4] N. M. Artner, A. Ion, and W. G. Kropatsch. Multi-scale 2d tracking of articulated objects using hierarchical spring systems. *Pattern Recognition*, 44(4):800–810, 2011. 1, 2

[5] S. Bai, Z. He, T. Xu, Z. Zhu, Y. Dong, and H. Bai. Multi-hierarchical independent correlation filters for visual tracking. *CoRR*, abs/1811.10302, 2018. 1, 7

[6] O. Barnich and M. Van Droogenbroeck. ViBe: A universal background subtraction algorithm for video sequences. *IEEE Transactions on Image Processing (TIP)*, 20(6):1709–1724, 2011. 1, 6

[7] F. Battistone, A. Petrosino, and V. Santopietro. Watch Out: Embedded Video Tracking with BST for Unmanned Aerial Vehicles. *Journal of Signal Processing Systems*, 2017. 1, 7

[8] H. Bay, A. Ess, T. Tuytelaars, and L. V. Gool. Speeded-up robust features (SURF). *Computer Vision and Image Understanding (CVIU)*, 110(3):346–359, 2008. 2

[9] A. Bhattacharyya. On a measure of divergence between two multinomial populations. *Sankhyā: The Indian Journal of Statistics (1933-1960)*, 7(4):401–406, 1946. 2, 4

[10] Z. Cai, L. Wen, Z. Lei, N. Vasconcelos, and S. Z. Li. Robust deformable and occluded object tracking with dynamic graph. *IEEE Transactions on Image Processing (TIP)*, 23(12):5497–5509, 2014. 1, 2

[11] L. Čehovin, M. Kristan, and A. Leonardis. Robust visual tracking using an adaptive coupled-layer visual model. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 35(4):941–953, 2013. 1, 2, 5, 7, 8

[12] L. Čehovin, A. Leonardis, and M. Kristan. Robust visual tracking using template anchors. In *IEEE Applications of Computer Vision (WACV)*, pages 1–8, 2016. 1, 2, 5, 7, 8

[13] Z. Chen, Z. Hong, and D. Tao. An experimental survey on correlation filter-based tracking. *CoRR*, abs/1509.05520, 2015. 2

[14] D. Comaniciu, V. Ramesh, and P. Meer. Real-time tracking of non-rigid objects using mean shift. In *IEEE Conference on Computer Vision and Pattern Recognition (CPVR)*, volume 2, pages 142–149, 2000. 2, 4

[15] M. Danelljan, G. Häger, F. S. Khan, and M. Felsberg. Learning spatially regularized correlation filters for visual tracking. In *IEEE International Conference on Computer Vision (ICCV)*, pages 4310–4318, 2015. 1

[16] M. Danelljan, A. Robinson, F. S. Khan, and M. Felsberg. Beyond correlation filters: Learning continuous convolution operators for visual tracking. In *IEEE European Conference on Computer Vision (ECCV)*, pages 472–488, 2016. 1

[17] G. De Ath. *Object Tracking in Video with Part-Based Tracking by Feature Sampling*. PhD thesis, University of Exeter, 2019. 4, 6

[18] G. De Ath and R. Everson. Visual object tracking: The initialisation problem. In *Conference on Computer and Robot Vision (CRV)*, pages 142–149, 2018. 1, 2, 5, 6, 7

[19] T. B. Dinh, N. Vo, and G. Medioni. Context tracker: Exploring supporters and distracters in unconstrained environments. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, pages 1177–1184, 2011. 7

[20] D. Du, H. Qi, W. Li, L. Wen, Q. Huang, and S. Lyu. Online deformable object tracking based on structure-aware hyper-graph. *IEEE Transactions on Image Processing (TIP)*, 25(8):3572–3584, 2016. 1, 2

[21] D. Du, H. Qi, L. Wen, Q. Tian, Q. Huang, and S. Lyu. Geometric hypergraph learning for visual tracking. *IEEE Transactions on Cybernetics (TCYB)*, PP(99):1–14, 2017. 1, 2

[22] H. Fan and J. Xiang. Robust visual tracking via local-global correlation filter. In *AAAI Conference on Artificial Intelligence*, pages 4025–4031, 2017. 2

[23] H. Grabner, M. Grabner, and H. Bischof. Real-time tracking via on-line boosting. In *British Machine Vision Conference (BMVC)*, 2006. 7

[24] S. Hare, S. Golodetz, A. Saffari, V. Vineet, M. M. Cheng, S. L. Hicks, and P. H. S. Torr. Struck: Structured output tracking with kernels. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 38(10):2096–2109, 2016. 1, 7

[25] E. Hellinger. Neue begründung der theorie quadratischer formen von unendlichvielen veränderlichen. *Journal für die Reine und Angewandte Mathematik*, 136:210–271, 1909. 4

[26] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista. Exploiting the circulant structure of tracking-by-detection with kernels. In *IEEE European Conference on Computer Vision (ECCV)*, pages 702–715, 2012. 7

[27] 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 (TPAMI)*, 37(3):583–596, 2015. 1

[28] X. Jia, H. Lu, and M. Yang. Visual tracking via adaptive structural local sparse appearance model. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, pages 1822–1829, 2012. 7

[29] J. Johnander, M. Danelljan, F. S. Khan, and M. Felsberg. Dcco: Towards deformable continuous convolution operators for visual tracking. In *International Conference on Computer Analysis of Images and Patterns*, pages 55–67, 2017. 2, 7

[30] Z. Kalal, K. Mikolajczyk, and J. Matas. Tracking-learning-detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 34(7):1409–1422, 2012. 7

[31] M. Kristan, A. Leonardis, J. Matas, M. Felsberg, R. Pflugfelder, L. Č. Zajc, T. Vojír, G. Bhat, A. Lukežič, A. Eldesokey, G. Fernández, et al. The visual object tracking vot2018 challenge results. In *IEEE European Conference on*

Computer Vision Workshop (ECCVW), pages 3–53, 2019. 1, 2, 6, 7

[32] M. Kristan, J. Matas, A. Leonardis, T. Vojír, R. Pflugfelder, G. Fernández, G. Nebehay, F. Porikli, and L. Č. Zajc. A novel performance evaluation methodology for single-target trackers. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 38(11):2137–2155, 2016. 2

[33] M. Kristan, R. Pflugfelder, A. Leonardis, J. Matas, F. Porikli, Z. Luka Čehovin, G. Nebehay, G. Fernandez Dominguez, T. Vojír, et al. The visual object tracking vot2013 challenge results. In *IEEE International Conference on Computer Vision Workshop (ICCVW)*, pages 98–111, 2013. 1

[34] J. Kwon and K. M. Lee. Tracking of a non-rigid object via patch-based dynamic appearance modeling and adaptive Basin Hopping Monte Carlo sampling. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, pages 1208–1215, 2009. 1, 2, 5

[35] J. Kwon and K. M. Lee. Visual tracking decomposition. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, pages 1269–1276, 2010. 7

[36] J. Kwon and K. M. Lee. Tracking by sampling trackers. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1195–1202, 2011. 7

[37] J. Kwon and K. M. Lee. Tracking by sampling and integrating multiple trackers. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 36(7):1428–1441, 2014. 2

[38] B. Li, J. Yan, W. Wu, Z. Zhu, and X. Hu. High performance visual tracking with siamese region proposal network. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, 2018. 7

[39] X. Li, W. Hu, C. Shen, Z. Zhang, A. Dick, and A. V. D. Hengel. A survey of appearance models in visual object tracking. *IEEE Transactions on Intelligent Transportation Systems (TITS)*, 4(4:58):1–48, 2013. 2

[40] B. Liu, J. Huang, C. Kulikowski, and L. Yang. Robust visual tracking using local sparse appearance model and k-selection. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 35(12):2968–2981, 2013. 7

[41] D. G. Lowe. Object recognition from local scale-invariant features. In *IEEE International Conference on Computer Vision (ICCV)*, volume 2, pages 1150–1157, 1999. 2

[42] A. Lukežič, L. Č. Zajc, and M. Kristan. Deformable parts correlation filters for robust visual tracking. *IEEE Transactions on Cybernetics (TCYB)*, pages 1–13, 2018. 1, 2, 7, 8

[43] M. E. Maresca and A. Petrosino. Matrioska: A multi-level approach to fast tracking by learning. In *Image Analysis and Processing (ICIAP)*, pages 419–428. Springer Berlin Heidelberg, 2013. 1, 2, 7

[44] M. E. Maresca and A. Petrosino. Clustering local motion estimates for robust and efficient object tracking. In *IEEE European Conference on Computer Vision (ECCV)*, pages 244–253, 2015. 1, 7

[45] T. B. Moeslund, A. Hilton, and V. Krüger. A survey of advances in vision-based human motion capture and analysis. *Computer Vision and Image Understanding (CVIU)*, 104(2):90–126, 2006. 2

[46] H. Nam and B. Han. Learning multi-domain convolutional neural networks for visual tracking. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, pages 4293–4302, 2016. 1

[47] S. Ojha and S. Sakhare. Image processing techniques for object tracking in video surveillance - a survey. In *International Conference on Pervasive Computing (ICPC)*, pages 1–6, 2015. 2

[48] F. Porikli. Achieving real-time object detection and tracking under extreme conditions. *Journal of Real-time Image Processing*, 1(1):33–40, 2006. 2

[49] X. Ren and J. Malik. Learning a classification model for segmentation. In *IEEE International Conference on Computer Vision (ICCV)*, volume 1, pages 10–17, 2003. 5

[50] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski. ORB: An efficient alternative to SIFT or SURF. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2564–2571, 2011. 2

[51] S. Salti, A. Cavallaro, and L. D. Stefano. Adaptive appearance modeling for video tracking: Survey and evaluation. *IEEE Transactions on Image Processing (TIP)*, 21(10):4334–4348, 2012. 2

[52] A. W. M. Smeulders, D. M. Chu, R. Cucchiara, S. Calderara, A. Dehghan, and M. Shah. Visual tracking: An experimental survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 36(7):1442–1468, 2014. 2

[53] D. Stutz, A. Hermans, and B. Leibe. Superpixels: An evaluation of the state-of-the-art. *Computer Vision and Image Understanding (CVIU)*, 166:1–27, 2018. 5, 6

[54] C. Sun, D. Wang, H. Lu, and M.-H. Yang. Learning spatial-aware regressions for visual tracking. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, 2018. 1

[55] T. Vojír and J. Matas. *The Enhanced Flock of Trackers*, pages 113–136. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014. 7

[56] T. Vojír and J. Matas. Pixel-wise object segmentations for the VOT 2016 dataset. Research Report CTU-CMP-2017-01, Center for Machine Perception, Czech Technical University, Prague, Czech Republic, 2017. 1, 4

[57] Y. Wu, J. Lim, and M.-H. Yang. Object tracking benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 37(9):1834–1848, 2015. 6

[58] J. Xiao, R. Stolkin, and A. Leonardis. Single target tracking using adaptive clustered decision trees and dynamic multi-level appearance models. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, pages 4978–4987, 2015. 1, 2

[59] T. Xu, Z. Feng, X. Wu, and J. Kittler. Learning adaptive discriminative correlation filters via temporal consistency preserving spatial feature selection for robust visual tracking. *CoRR*, abs/1807.11348, 2018. 1, 7

[60] Y. Xu, J. Dong, B. Zhang, and D. Xu. Background modelling methods in video analysis: A review and comparative evaluation. *CAAI Transactions on Intelligence Technology*, 1(1):43–60, 2016. 1

[61] H. Yang, L. Shao, F. Zheng, L. Wang, and Z. Song. Recent advances and trends in visual tracking: A review. *Neurocomputing*, 74(18):3823–3831, 2011. 2

[62] R. Yao, S. Xia, Z. Zhang, and Y. Zhang. Real-time correlation filter tracking by efficient dense belief propagation with structure preserving. *IEEE Transactions on Multimedia (TMM)*, 19(4):772–784, 2017. 2

[63] K. M. Yi, H. Jeong, S. W. Kim, S. Yin, S. Oh, and J. Y. Choi. Visual tracking of non-rigid objects with partial occlusion through elastic structure of local patches and hierarchical diffusion. *Image and Vision Computing (IVC)*, 39:23–37, 2015. 1, 2, 5, 7

[64] A. Yilmaz, O. Javed, and M. Shah. Object tracking: A survey. *ACM Computing Surveys*, 38(4), 2006. 2

[65] S. Zhang, H. Yao, X. Sun, and X. Lu. Sparse coding based visual tracking: Review and experimental comparison. *Pattern Recognition*, 46(7):1772–1788, 2013. 2

[66] G. Zhu, J. Wang, C. Zhao, and H. Lu. Weighted part context learning for visual tracking. *IEEE Transactions on Image Processing (TIP)*, 24(12):5140–5151, 2015. 1, 2