

Template Matching with Deformable Diversity Similarity

Itamar Talmi*
Technion

titamar@campus.technion.ac.il

Roey Mechrez*
Technion

roey@tx.technion.ac.il

Lihi Zelnik-Manor
Technion

lihi@ee.technion.ac.il

Abstract

We propose a novel measure for template matching named *Deformable Diversity Similarity* – based on the diversity of feature matches between a target image window and the template. We rely on both local appearance and geometric information that jointly lead to a powerful approach for matching. Our key contribution is a similarity measure, that is robust to complex deformations, significant background clutter, and occlusions. Empirical evaluation on the most up-to-date benchmark shows that our method outperforms the current state-of-the-art in its detection accuracy while improving computational complexity.

1. Introduction

Template Matching is a key component in many computer vision applications such as object detection, tracking, surveillance, medical imaging and image stitching. Our interest is in Template Matching “in the wild” [4], i.e., when no prior information is available on the target image. An example application is to identify the same object in different cameras of a surveillance system [1]. Another use case is in video tracking, where Template Matching is used to detect drifts and relocate the object after losing it [9]. This is a challenging task when the transformation between the template and the target in the image is complex, non-rigid, or contains occlusions, as illustrated in Figure 1.

Traditional template matching approaches, such as Sum-of-Squared-Distances or Normalized Cross-Correlation, do not handle well these complex cases. This is largely because they penalize all pixels of the template, which results in false detections when occlusions or large deformations occur. To overcome this limitation the Best-Buddies-Similarity (BBS) measure was proposed in [4, 18]. BBS is based on properties of the Nearest-Neighbor (NN) matches between features of the target and features of the template. It relies only on a subset of the points in the template, thus hoping to latch on to the relevant features that correspond between the template and the target. This makes BBS more robust than previous methods.

*Authors contributed equally



Figure 1: **Template Matching challenges:** Template Matching results of the proposed Deformable Diversity Similarity (DDIS). (Top) The Miami Heats logo on Lebron James’s shirt, is marked as template (in green). (Bottom) Best matches found by DDIS in four target images (in magenta). Different challenges are evident: the ball occludes the logo, Lebron moves non-rigidly with out-of-plane rotation and complex deformation of the shirt.

In this paper we adopt the feature-based, parameter-free, approach of BBS and propose a novel similarity measure for template matching named DDIS: *Deformable Diversity Similarity*. DDIS is based on two properties of the Nearest Neighbor field of matches between points of a target window and the template. The first is that the *diversity* of NN matches forms a strong cue for template matching. This idea is supported by observations in [11], where patch diversity was used to match objects for texture synthesis. We propose formulas for measuring the NN field diversity and further provide theoretical analysis as well as empirical evaluations that show the strength of these measures.

The second key idea behind DDIS is to explicitly consider the *deformation* implied by the NN field. As was shown by the seminal work of [7] on Deformable Part Models, allowing deformations while accounting for them in the

matching measure is highly advantageous for object detection. DDIS incorporates similar ideas for template matching leading to a significant improvement in template detection accuracy in comparison to the state-of-the-art.

A benefit of DDIS with respect to BBS [4, 18] is reduced computational complexity. Both measures rely on NN matches, however, BBS is formulated in a way that requires heavier computations. DDIS is more efficient while providing statistical properties similar to BBS.

To summarize, in this paper we introduce DDIS, a measure for *template matching in the wild* that relies on two observations: (i) The diversity of NN matches between template points and target points is indicative of the similarity between them. (ii) The deformation implied by the NN field should be explicitly accounted for. DDIS is robust and parameter free, it operates in unconstrained environments and shows improved accuracy compared to previous methods on a real challenging data-set. Our code is available at <https://github.com/roimehrez/DDIS>

2. Related Work

The similarity measure between the template and a sub-window of the target image is the core part of template matching. A good review is given in [19]. The commonly used methods are pixel-wise, e.g., Sum of Squared differences (SSD), Sum of Absolute Differences (SAD) and Normalized Cross-Correlation (NCC), all of which assume only translation between the template and target. They could be combined with tone mapping to handle illumination changes [8] or with asymmetric correlation to handle noise [6]. To increase robustness to noise pixel-wise measures such as M-estimators [2, 23] or Hamming-based distances [22, 20] have been proposed.

More general geometric transformation such as affine are addressed by [27, 12]. In [13] parametric transformations are handled by approximating the global optimum of the parametric model. In [26] non-rigid transformations are addressed via parametric estimation of the distortion. All of these methods work very well when their underlying geometric assumptions hold, however, they fail in the presence of complex deformations, occlusions and clutter.

A second group of methods consider a global probabilistic property of the template. For example in [3, 21] color Histogram Matching is used (for tracking). This does not restrict the geometric transformation, however, in many cases the color histogram is not a good representation, e.g., in the presence of background clutter and occlusions. Other methods combine geometric cues with appearance cues. For example, a probabilistic solution was suggested in [16], where geometric and color cues are used to represent the image in the location-intensity space. Oron et al. [17] extend this idea by measuring one-to-one distance in $xyRGB$ space. These methods all make various assumptions that do

not hold in complex scenarios.

A more robust approach, that can handle complex cases has been recently suggested in [4, 18]. Their approach, named the Best-Buddies-Similarity (BBS) is based on the Bi-Directional Similarity (BDS) concept of [24]. They compute the similarity between a template and a target window by considering matches between their patches. The matches are computed in both directions providing robustness to outliers. A similar idea was suggested in [5] by replacing the max operator of the Hausdorff distance [10] with a sum. The BBS of [4, 18] lead to a significant improvement in template matching accuracy over prior methods. In this paper we propose a different measure, that shares with BBS its robustness properties, while yielding even better detection results.

3. Diversity as a Similarity Measure

To measure similarity between a target window and a template we first find for every target patch its Nearest Neighbor (NN), in terms of appearance, in the template. Our key idea is that the similarity between the target and the template is captured by two properties of the implied NN field. First, as shown in Figure 2d, when the target and template correspond, most target patches have a unique NN match in the template. This implies that the NN field is highly *diverse*, pointing to many different patches in the template. Conversely, as shown in Figure 2e, for arbitrary targets most patches do NOT have a good match, and the NNs converge to a small number of template points that happen to be somewhat similar to the target patches. Second, we note that arbitrary matches typically imply a large *deformation*, indicated by long arrows in Figure 2e.

Next, we propose two ways for quantifying the amount of diversity and deformation of the NN field. The first is more intuitive and allows elegant statistical analysis. The second is slightly more sophisticated and more robust.

3.1. Diversity Similarity (DIS)

Let points $p_i, q_j \in \mathbb{R}^d$ represent patches of the template and target, respectively. Our goal is to measure the similarity between two sets of points, the template points $P = \{p_i\}_{i=1}^N$ and the target points $Q = \{q_j\}_{j=1}^M$. We require finding the NN in P for every point $q \in Q$, s.t., $NN(q, P) = \operatorname{argmin}_{p \in P} d(q, p)$ for some given distance function $d(q, p)$. The first property our measures are based on is the diversity of points $p \in P$ that were found as NNs.

An intuitive way to measure diversity is to count the number of unique NNs. We define the *Diversity Similarity* (DIS) as:

$$\mathbf{DIS}_{Q \rightarrow P} = c \left| \{p_i \in P : \exists q_j \in Q, NN(q_j, P) = p_i\} \right| \quad (1)$$

where $c = 1/\min\{M, N\}$ is a normalization factor and $|\{\cdot\}|$




Figure 2: **Deformable Diversity:** A face template (c), marked in green (a), is searched in a target image (b). The face undergoes a complex transformation: background clutter, out of plain rotation, and non-rigid deformation. We zoom on a bad target window (e) and on the best target window (d), both are also marked by color in (b). The blue arrows are samples of the NN field: their start point marks a target patch while the end point marks the position of its NN in the template. The bad target (e), shows low diversity of NN patches with large deformation (long arrows). Conversely, the best target (d) shows high diversity and small deformation (short arrows). Best viewed on screen.

denotes group size.

To provide further intuition as to why DIS captures the similarity between two sets of points we provide an illustration in 2D in Figure 5. Figure 5a demonstrates that when the distributions of points in P and Q are similar, most of the points $q \in Q$ have a unique NN $p \in P$ implying a high DIS value. Conversely, when P and Q are distributed differently, as illustrated in Figure 5b, DIS is low. This is since in areas where Q is sparse while P is dense most of the points in P are not NN of any q . In addition, in areas where Q is dense while P is sparse most of the points in Q share the same NNs. In both cases, since the number of points is finite, the overall contribution to DIS is low.

3.2. Deformable Diversity Similarity (DDIS)

While capturing well diversity, DIS does not explicitly consider the deformation field. Accounting for the amount of deformation is important since while non-rigid transformations should be allowed, they should also be restricted to give preference to plausible deformations of real objects.

In order to integrate a penalty on large deformations we make two modifications to the way we measure diversity. First, to obtain an explicit representation of the deformation field we distinguish between the appearance and the posi-

tion of each patch and treat them separately. Second, we propose a different way to measure diversity, that enables considering the deformation amount.

Let p^a denote the appearance and p^l the location of patch p (and similarly for q). We find the appearance based NN p_i for every point q_j s.t. $p_i = \text{NN}^a(q_j, P) = \text{argmin}_{p \in P} d(q_j^a, p^a)$ for some given distance $d(q^a, p^a)$. The location distance between a point q_j and its NN^a is denoted by $r_j = d(q_j^l, p_i^l)$. To quantify the amount of diversity as a function of the NN field we define $\kappa(p_i)$ as the number of patches $q \in Q$ whose NN^a is p_i :

$$\kappa(p_i) = |\{q \in Q : \text{NN}^a(q, P) = p_i\}| \quad (2)$$

Finally, we define the *Deformable Diversity Similarity* (DDIS) by aiming for high diversity and small deformation:

$$\text{DDIS} = c \sum_{q_j \in Q} \frac{1}{r_j + 1} \cdot \exp\left(1 - \kappa(\text{NN}^a(q_j, P))\right) \quad (3)$$

where $c = 1/\min\{M, N\}$ is a normalization factor.

This definition can be viewed as a sum of contributions over the points q_j . When a point q_j has a unique NN, then $\kappa(\text{NN}^a(q_j, P)) = 1$ and the exponent reaches its maximum value of 1. Conversely, when the NN of q_j is shared by many other points, then $\kappa(\text{NN}^a(q_j, P))$ is large, the exponent value is low and the overall contribution of q_j to the similarity is low. In addition, the contribution of every point is inversely weighted by the length r_j of its implied deformation vector.

DDIS possesses several properties that make it attractive: (1) it relies mostly on a subset of matches, i.e., points that have distinct NNs. Points that share NNs will have less influence on the score. (2) DDIS does not require any prior knowledge on the data or its underlying deformation. (3) DDIS analyses the NN field, rather than using the actual distance values. These properties allow DDIS to overcome challenges such as background clutter, occlusions, and non-rigid deformations.

3.3. DIS as simplified DDIS

DIS and DDIS capture diversity in two different ways. DIS simply counts unique matches in P , while DDIS measures exponentially the distinctiveness of each NN match of patches in Q . Nonetheless, we next show that DIS and DDIS are highly related.

We start by ignoring the deformations by setting $r_j = 0$ in (3) and simplifying (without loss of generality) by assuming $M = N$. We denote by $1/k$ the fraction of points $p \in P$ that are NNs of at least one point in Q . Both DIS and DDIS reach their maximum value of 1 when $k = 1$, i.e., when $\kappa(p_i) = 1 \forall p_i \in P$. When $k = N$, i.e., all $q \in Q$ share a single NN, both scores reach their minimum value, $\text{DIS} = 1/N$ and $\text{DDIS} = \exp(1 - N)$.




Figure 3: **DIS as simplified DDIS:** Plots of DIS and DDIS (with $r_j=0$, $M=N$) as a function of k , where N/k points in P are NNs of some q , and for all of them $\kappa(p_i) = k$. Both DIS and DDIS are maximal at $k = 1$ and decrease monotonically, reaching their minimal value when $k = N$.

Further intuition can be derived from the case of uniform distribution of NN matches, i.e., when only N/k points in P are NNs of some q , and for all of them $\kappa(p_i) = k$. In this case $\text{DIS} = (N/k)/N = 1/k$, and $\text{DDIS} = 1/N \sum_{q_j \in Q} \exp(1-k) = \exp(1-k)$. Both measures share extrema points between which they drop monotonically as a function of k , with DDIS decreasing faster due to its exponential nature. This is illustrated in Figure 3.

3.4. Statistical Analysis

To further cement our assertions that diversity captures the similarity between two distributions, we provide statistical analysis, similar to that presented in [4, 18]. Our goal is to show that the expectation of DIS and DDIS is maximal when the points in both sets are drawn from the same distribution, and drops sharply as the distance between the two distributions increases. We do that via a simple 1D mathematical model, in which an image window is modeled as a set of points drawn from a general distribution.

Appendix A presents derivations of $E[\text{DIS}]$ (Expected value of DIS) when the points are drawn from two given distributions. The expression for $E[\text{DIS}]$ does not have a closed form solution, but it can be solved numerically for selected underlying distributions. Therefore, we adopt the same setup as [4] where P and Q are assumed to be Gaussian distributions, which are often used as simple statistical models of image patches. We then use Monte-Carlo integration to approximate the Expectation for discrete choices of parameters μ_Q and σ_Q . For BBS and SSD we adopt the derivations in [4], where $E[\text{BBS}]$ was also approximated via Monte-Carlo integration and $E[\text{SSD}]$ is normalized.

Figure 4 presents the resulting approximated expected values. It can be seen that DIS is likely to be maximized when the distributions are the same, and falls rapidly when the distributions differ from each other. In addition it is evident that DIS and BBS present highly similar behaviors. Finally, similar to [4], one can show that this holds also for the multi-dimensional case.

For DDIS we cannot derive nice expressions for its Expectation $E[\text{DDIS}]$. Instead, we use simulations to approximate it. The simulation needs to consider also locations




Figure 4: **Expected behavior in the 1D Gaussian case:** Following [4] two point sets, P and Q , are generated by sampling $N = M = 100$ points from $N(0; 1)$, and $N(\mu_Q; \sigma_Q)$, respectively, with $[\mu_Q, \sigma_Q] \in [0, 10]$. (Top) The approximated expectation of SSD (a), BBS (b) and DIS (c) as a function of μ_Q and σ_Q suggest that BBS and DIS behave similarly and drop much more rapidly than SSD as the distributions move apart. (Bottom) The approximated expectation of DDIS when the mean deformation is small (d), large (e), and ignored (f) ((d) and (e) are color scaled jointly). It can be seen that small deformation fields correspond to a sharper peak in the expectation, while for large deformations the similarity is always low.

to quantify the amount of deformation, r_j , in (3). When $r_j = 0$ the expectation is similar to that of BBS and DIS. For $r_j \neq 0$ we simulate two cases: (i) Small deformation: We sort the points in each set based on their appearance coordinate, and take as position their index in the sorted list. When the distributions are different the diversity is very low anyhow. But when the distributions are similar, the sorting results in points and their NN having a *similar index*, which corresponds to *small deformation*. (ii) Large deformation: We sort the points of one set in descending order and the other set in ascending order, again taking as position their index in the sorted list. When the distributions are similar, the sorting results in points and their NN having a *different index*, which corresponds to *large deformation*. Figure 4 shows that for small deformation $E[\text{DDIS}]$ drops sharply as the distributions become more different. For large deformations it is always low, as desired, since even when the appearances are similar, if the geometric deformation is large the overall similarity between the point sets is low.

4. Comparison to BBS

Our measures bare resemblance to BBS – all rely on NN matches between two sets of points. There are, however,




Figure 5: **Intuition for DIS and BBS:** For each point $q \in Q$ we draw a red arrow pointing at its NN $p \in P$. If q and p are also best-buddies-pair (BBP, see section 4.1), we change the red arrow to a green line. DIS counts blue triangles that are pointed to by either a red arrow or a green line. BBS counts green lines. (a) P and Q are similarly distributed, hence, many p 's are NN of some q and there are many BBPs. Here $DIS = BBS = 8$. (b) P and Q have different distributions. A single q among dense p 's or a single p among dense q 's contribute 1 to both DIS and BBS. Occasionally, there is a unique NN match between Q and P that is not a BBP. Since the distributions of P and Q are different both DIS and BBS are relatively low, $DIS = 3$ and $BBS = 2$.

two key differences: (i) the way in which similarity between the two sets is measured, and, (ii) the penalty on the amount of spatial deformation. We next analyze the implications of these differences.

4.1. You Only Need One Direction

The key idea behind the bi-directional similarity approaches of [4, 18, 24] is that robust matching requires bi-directional feature correspondences. Our unilateral measures contradict this claim. In fact, an observation we make is that, Diversity provides a good approximation to BBS. The analysis we present is for DIS, since it is simpler than DDIS and does not incorporate deformation, making the comparison to BBS more fair and direct.

Recall that BBS counts the number of bi-directional NN matches between the target and template. A pair of points $p_i \in P$ and $q_j \in Q$ are considered a best-buddies-pair (BBP) if p_i is the NN of q_j , and q_j is the NN of p_i . BBS counts the number of BBPs as a measure of similarity between P and Q . Clearly, the bilateral nature of BBS is wasteful in terms of computations, compared to the unilateral DDIS and DIS.

DIS and BBS are defined differently, however, since the number of patches in both target and template is *finite*, DIS provides a good approximation to BBS. As illustrated in Figure 5a when the distributions of points in P and Q are similar, many of the NN relations are bi-directional. This implies that the values of BBS and DIS are very similar. In the extreme case when the template and target are identical, every point q has a unique NN $p = q$ and they form a BBP. In this case $DIS = BBS$ exactly.

DIS and BBS behave similarly also when the distributions are different, as illustrated in Figure 5b. In areas where

P is sparse and Q is dense we get multiple points $q \in Q$ that share the same NN $p \in P$. At most one of them forms a BBP and their joint contribution to both DIS and BBS is 1. Since the number of points in P and Q is finite, this implies that there are other areas where P is dense and Q is sparse. In these areas there are many points in P that are not NN of any $q \in Q$, and have zero contribution to both DIS and BBS.

Our observations are in line with the Expectation analysis of Figure 4. In addition, our experiments (Section 7) show that DIS and BBS achieve comparable template matching accuracy.

4.2. Explicit vs. Implicit Deformation Modeling

The need to penalize large deformations was noted in [4, 18]. This was done implicitly by adding the xy coordinates to the feature vectors when searching for NNs. The distance between a pair of points is taken as a weighted linear combination of their appearance and position difference. This is different from DDIS that considers only appearance for NN matching and explicitly penalizes the deformation in the obtained NN field. Our approach has two benefits: (i) improved runtime, and (ii) higher detection accuracy.

Using only appearance for NN matching significantly reduces runtime since while every image patch is shared by many sub-windows, its xy coordinates are different in each of them. This implies that the NN field needs to be computed for each image sub-window separately. Conversely, working in appearance space allows us to perform a single NN search per image patch. In Section 6 we analyze the benefits in terms of computational complexity.

Separating between appearance and position also leads to more accurate template localization. Overlapping target windows with very similar appearance could lead to very similar similarity scores. DDIS chooses the window implying less deformations. Our experiments indicated that this is important and improves the localization accuracy.

5. Implementation

To utilize DDIS for template matching in images, we follow the traditional raster scan approach. Our algorithm gets as input a target image S and a template T . Its output is a frame placing T within S . We denote the width of T by T_w and its height T_h , similarly for S . Each template sized sub-window $W \in S$ is compared to T . We extract from T and W feature vectors, as described below, yielding sets P and Q respectively. We use the Euclidean distance (L_2) to compare appearance features p^a and q^a . The deformation length r_j is the Euclidean distance between the xy coordinates p^l and q^l . Our implementation consists of 4 phases:

0. Feature extraction: We experimented with two forms of appearance feature, color and deep-features. As color features we set p^a and q^a as vectorized RGB pixel values of 3×3 overlapping patches. To obtain deep-features we

used the popular VGG-Deep-Net [25]. More specifically, we take feature maps from layers *conv1_2*, *conv3_4* and *conv4_4* (akin to the suggestion in [14] for object tracking). We forsake the higher layers since we found the low spatial resolution inaccurate. The features maps were normalized to zero mean and unit standard deviation, and then upsampled, via bi-linear interpolation, to reach the original image size.

1. NN search: We find for each feature vector in S , its approximate NN in the template T . We use TreeCANN [15] with PCA dimensionality reduction to 9 dimensions, kd-tree approximation parameter $\epsilon = 2$, dense search ($g_S = g_T = 1$), and window parameter $w_S = 3$, $w_T = 5$.

2. Similarity map calculation: For each target image pixel (ignoring boundary pixels) we compute the similarity between its surrounding sub-window W and the template T . For each W , we first compute $\kappa(p_i)$, $\forall p_i \in P$ as defined in (2). Since subsequent windows have many overlaps, the computation of κ needs only update the removed and added features with respect to the previous sub-window. We then calculate DDIS as defined in (3).

3. Target localization: Finally, the template location is that with maximum score. Before taking the maximum, we smooth the similarity map with a uniform kernel of size $\frac{T_w}{3} \times \frac{T_h}{3}$, to remove spurious isolated peaks.

6. Complexity

The sets P and Q consist of features from all locations in T and W , receptively. This implies $|P| = |Q| = T_w \cdot T_h \triangleq L$. The number of possible sub-windows¹ is $S_w \cdot S_h \triangleq L$. Recall that d denotes the feature vectors length. For color features d equals the size of the patch $\times 3$ while for deep-features it is determined by the feature map dimension. Next, we analyze the complexity of steps (1-3).

1. NN search: TreeCANN consists of two stages. In the first, the dimension of all template points is reduced from d to d' in $O(dl)$ and a k-d tree is built in $O(d'l \log l)$. The second stage performs the queries. Each query consists of dimensionality reduction $O(d)$, a search in the k-d tree $O(\log l)$ (on average), and a propagation phase which leverages spatial coherency $O(d)$. The overall complexity for finding the Approximate NN for all the features in the target image S is $O(d'l \log l + L(d + \log l))$ on average. The memory consumption is $O(l)$.

2. Similarity map calculation: Assuming for simplicity that $T_w = T_h = \sqrt{L}$, the update of $\kappa(p)$ takes $O(\sqrt{L})$ operations, for any W except for the first one. Next, DDIS is calculated with $O(l)$ operations. Given that the overall number of sub-windows is L this step's complexity is $O(L(l + \sqrt{L})) \equiv O(Ll)$. The memory consumption for this stage is $O(l)$ which is the size of a table holding $\kappa(p)$.

¹In practice, we exclude patches that are not fully inside the template or the sub-window, but these are negligible for our complexity analysis.

3. Target localization: Averaging the similarity map is done efficiently using an integral image in $O(L)$. To find the maxima location another swipe over the image is needed, which takes $O(L)$.

Putting it all together, we get that the overall complexity of Template Matching with DDIS is $O(d'l \log l + Ll)$ where we omitted $O(L(d + \log l))$ since d and l are expected to be of the same order for small T and $d \ll l$ for large T .

Comparison to BBS One of the benefits of DDIS with respect to BBS is that it requires only unilateral matches. The benefit in terms of complexity can now be made clear. According to [18], the BBS complexity using deep features is $O(Ll^4d)$ and for color features it is $O(Ll^2d)$ on average. The latter case uses heavy caching which consumes $O(l^2\sqrt{L})$ memory (assuming $S_w = S_h = \sqrt{L}$).

7. Empirical Evaluation

Our experimental setup follows that of [18] that created a benchmark by sampling frames from video sequences annotated with bounding-boxes for object tracking [28]. The videos present a variety of challenges: complex deformations, luminance changes, scale differences, out-of-plane rotations, occlusion and more. The benchmark consists of three data-sets, generated by sampling $\{270, 270, 254\}$ pairs of frames with a constant frame (time) difference $dFrame = \{25, 50, 100\}$, producing increasingly challenging data-sets, and overall, a challenging benchmark for template matching.

For each pair of frames, one is used to define the template as the annotated ground-truth box, while the second is used as a target image. As commonly done in object tracking, the overlap between the detection result and the ground-truth annotation of the target is taken as a measure of accuracy: $Accuracy = \frac{|R_{est} \cap R_{truth}|}{|R_{est} \cup R_{truth}|}$ where $|\cdot|$ counts the number of pixels in a region and R_{truth} and R_{est} are the ground truth and estimated rectangles, locating T in S .

Quantitative Evaluation: We compare DDIS and DIS to BBS, BDS, SSD, SAD and NCC with both color and deep features. For BBS and DIS we use the exact same setup as in [18], that is, 3×3 non-overlapping patches represented in $xyHSV$ space. In Figure 6 we plot for each data-set and method a success rate curve. It can be seen from the Area-Under-Curve (AUC) scores in Table 6d that DDIS is significantly more successful than all previous methods. Furthermore, DDIS with our simplistic color features outperforms all other methods with either color or deep features. When using Deep features, DDIS improves over BBS with margins of $\approx 11\%$, 15% , 8% for the three data-sets. When using color features the margins are $\approx 10\%$, 16% , 21% .

Qualitative Evaluation: Figure 8 displays several detection results on challenging examples, taken from the

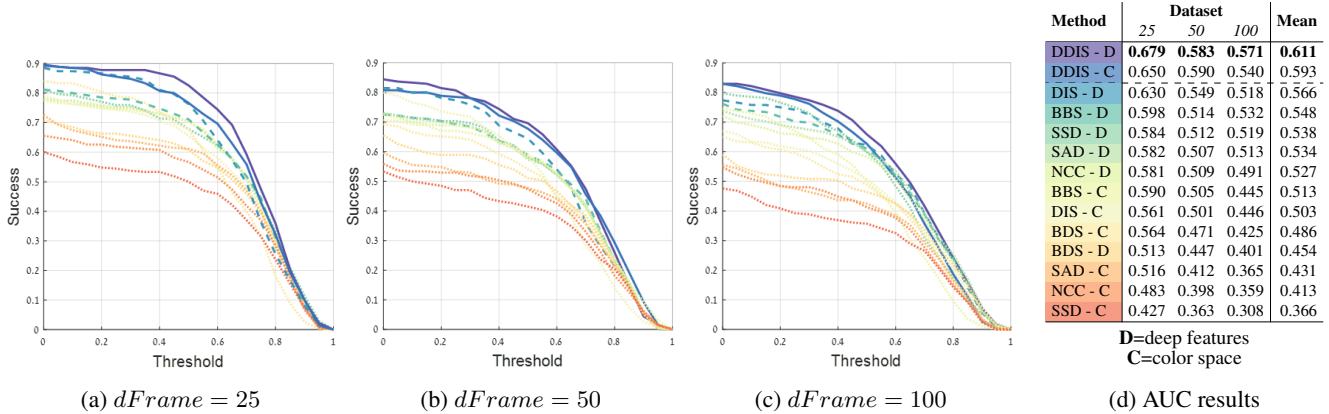


Figure 6: **Template matching accuracy:** Evaluation on the benchmark of [18]: 270 template-image pairs with $dFrame \in \{25, 50, 100\}$. DDIS outperforms competing methods as can be seen in the ROC curves (a)-(c) showing the fraction of image-pairs with $Accuracy > Threshold \in [0, 1]$. The corresponding Area-under-curve (AUC) scores are presented in the table in (d). For all methods D stands for deep features while C stands for color. DDIS with deep features provides the best results. DDIS with color features comes in second, outperforming other methods even when they use deep-features.

web, that include occlusions, significant deformations, background clutter and blur. We compare DDIS and DIS to BBS – the current state-of-the-art. It is evident from the detection likelihood maps that DIS and BBS share a similar behavior, supporting our suggestion that unilateral matching suffices to capture similarity. DDIS, on the other hand, accounts also for deformations, hence, it presents cleaner maps, with fewer distractors.

Runtime: Our implementation is in MATLAB/c++ and all experiments were performed on a 32GB RAM, Intel i7 quad-core machine. The average(std) runtime for an image-pair in the benchmark, using color features is $0.86s(0.59)$, depending on the template size. For comparison, the average(std) time for BBS is orders of magnitude longer: $35.47s(80.36)$. The max and min runtimes of DDIS are $3.44s$ and $0.06s$, respectively, and for BBS are $493s$ and $0.14s$, respectively. Detailed results for $dFrame = 25$ are presented in Figure 7. This matches our complexity analysis that showed that DDIS is less affected by the template

size, while BBS dependence on l is polynomial.

8. Conclusions

We introduced a new approach for template matching in the wild, based on properties of the NN field of matches between target and template features. Our method suggests not only improvement in terms of detection accuracy, but also in terms of computational complexity. A drawback of our algorithm is not dealing with significant scale change of the object. This could possibly be addressed, by computing the likelihood maps over multiple scales. A future research direction is to explore consideration of more than the first NN for each patch. This could be beneficial to handle repetitive textures.

An important observation, our analysis makes, is that one does not necessarily need bi-directional matches to compute similarity. This raises questions regarding the celebrated bi-directional-similarity approach, which provided excellent results, but was heavy to compute.

Acknowledgements This research was supported by the Israel Science Foundation Grant 1089/16 and by the Ollendorf foundation.

A. Appendix: DIS Expectation Term

In this appendix we develop mathematical expressions for the *expectation* of DIS in \mathbb{R}^1 . We start by rewriting DIS in a form convenient for our derivations:

$$\mathbf{DIS}_{Q \rightarrow P} = \frac{1}{N} \cdot \sum_{i=1}^N dis_i(Q, P)$$

$$dis_i(Q, P) = \mathbb{I}[\{q_j \in Q : \text{NN}(q_j, P) = p_i\} \neq \emptyset]$$




Figure 7: **Runtime:** Each point in the scatter plot marks the runtime for one of the 270 image-pairs in the dataset ($dFrame = 25$). For DDIS the runtime is always small while for BBS it becomes very long as the template and image size increase.

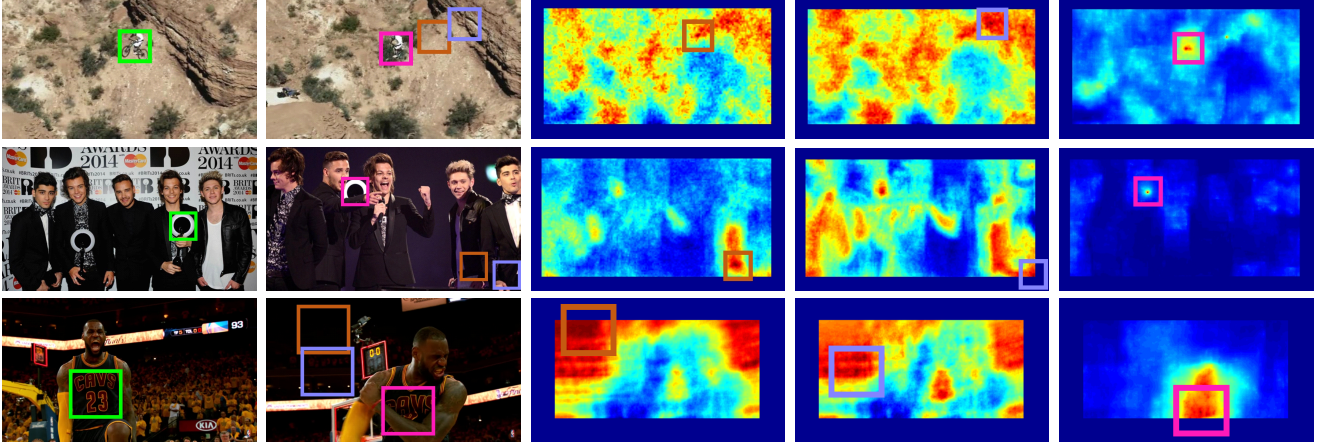


Figure 8: **Qualitative assessment:** The template marked in green (a), is detected in the target image (b) using three Template Matching methods: **BBS**, **DIS** and **DDIS** (all using the RGB features). (c-e) The corresponding detection likelihood maps show that DDIS yields more peaked maps that more robustly identify the template. Going over the rows from top to bottom: (1) BBS prefers a target location where the background matches the template over the location where the motorcycle is at. This happens because the motorcycle deforms and hence there are few bi-directional correspondences between its template appearance and target appearance. DIS and DDIS use more information – they consider all the one-directional correspondences. Therefore, they locate the motorcycle correctly. (2) The trophy won by the band *One Direction* is fully seen in the template, but occluded in the target. Nonetheless, DDIS finds it (as Section 4.1 said, we only need one direction...). (3) Complex deformations together with occlusion confuse both DIS and BBS, but not DDIS.

where \mathbb{I} is an indicator function and $dis_i(Q, P)$ indicates whether p_i is chosen as a NN match at least once.

We proceed with the expectation:

$$E \left[\mathbf{DIS}_{Q \rightarrow P} \right] = \frac{1}{N} \sum_{i=1}^N E [dis_i(Q, P)] = E [dis_k(Q, P)] \quad (4)$$

where the last step is since samples Q and P are drawn independently, so all indexes behave alike and we can choose some arbitrary index k . Continuing with the expectation of the indicator function, we have:

$$\begin{aligned} E [dis_k(Q, P)] &= \Pr \{dis_k(Q, P) = 1\} \\ &= 1 - \Pr \{dis_k(Q, P) = 0\} \end{aligned} \quad (5)$$

Claim:

$$\begin{aligned} \Pr \{dis_i(Q, P) = 0\} &= \\ &= \int_{p_1} \cdots \int_{p_N} (F_Q(p_i^-) + 1 - F_Q(p_i^+))^N \cdot \prod_{k=1}^N f_P(p_k) dp_k \end{aligned} \quad (6)$$

where $F_Q(x) = \Pr \{q \leq x\}$ and $F_P(x) = \Pr \{p \leq x\}$ are the CDF's of Q and P , respectively. p_i^+, p_i^- are defined by:

$$\begin{aligned} p_i^+ &= p_i + \min_{\substack{p_k \in P \cup \{+\infty\} \\ p_k > p_i}} |p_k - p_i|/2 \\ p_i^- &= p_i - \min_{\substack{p_k \in P \cup \{-\infty\} \\ p_k < p_i}} |p_k - p_i|/2 \end{aligned} \quad (7)$$

Proof: Given a known set of samples P , the probability that the NN match for a sampled $q \sim Q$ is NOT p_i is:

$$\begin{aligned} \Pr \{NN(q, P) \neq p_i | P\} &= \int_{-\infty}^{\infty} \mathbb{I} [NN(q, P) \neq p_i] f_Q(q) dq \\ &= \int_{-\infty}^{p_i^-} f_Q(q) dq + \int_{p_i^+}^{\infty} f_Q(q) dq \\ &= F_Q(p_i^-) + 1 - F_Q(p_i^+) \end{aligned} \quad (8)$$

where we split \mathbb{R} into two ranges where the indicator is not zero. Since Q consists of N independently sampled points, the probability that p_i is not a NN match for any $q \in Q$ when Q is sampled and P is known, is:

$$\Pr \{dis_i(Q, P) = 0 | P\} = \left[\Pr \{NN(q, P) \neq p_i | P\} \right]^N \quad (9)$$

Finally, since all of the points are sampled independently, we have:

$$\begin{aligned} \Pr \{dis_i(Q, P) = 0\} &= \\ &= \int_{p_1} \cdots \int_{p_N} \Pr \{dis_i(Q, P) = 0 | P\} \cdot \prod_{k=1}^N f_P(p_k) dp_k \end{aligned} \quad (10)$$

Substituting equations (8) and (9) in (10) results in (6). ■

References

- [1] S. M. Assari, H. Idrees, and M. Shah. Human re-identification in crowd videos using personal, social and environmental constraints. In *European Conference on Computer Vision*, pages 119–136. Springer, 2016. 1
- [2] J.-H. Chen, C.-S. Chen, and Y.-S. Chen. Fast algorithm for robust template matching with m-estimators. *IEEE Transactions on Signal Processing*, 51(1):230–243, 2003. 2
- [3] D. Comaniciu, V. Ramesh, and P. Meer. Real-time tracking of non-rigid objects using mean shift. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 142–149, 2000. 2
- [4] T. Dekel, S. Oron, M. Rubinstein, S. Avidan, and W. T. Freeman. Best-buddies similarity for robust template matching. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2021–2029, 2015. 1, 2, 4, 5
- [5] M. P. Dubuisson and A. K. Jain. A modified hausdorff distance for object matching. In *Proceedings of the 12th IAPR International Conference on Pattern Recognition*, volume 1, pages 566–568, Oct 1994. 2
- [6] E. Elboher and M. Werman. Asymmetric correlation: A noise robust similarity measure for template matching. *IEEE Transactions on Image Processing*, 22(8):3062–3073, 2013. 2
- [7] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(9):1627–1645, 2010. 1
- [8] Y. Hel-Or, H. Hel-Or, and E. David. Matching by tone mapping: Photometric invariant template matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(2):317–330, 2014. 2
- [9] Z. Hong, Z. Chen, C. Wang, X. Mei, D. Prokhorov, and D. Tao. Multi-store tracker (muster): A cognitive psychology inspired approach to object tracking. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 749–758, 2015. 1
- [10] D. P. Huttenlocher, G. A. Klanderman, and W. J. Rucklidge. Comparing images using the hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(9):850–863, 1993. 2
- [11] O. Jamriřka, J. Fiřer, P. Asente, J. Lu, E. Shechtman, and D. Šykora. LazyFluids: appearance transfer for fluid animations. *ACM Transactions on Graphics (TOG)*, 34(4):92, 2015. 1
- [12] H. Y. Kim and S. A. De Araújo. Grayscale template-matching invariant to rotation, scale, translation, brightness and contrast. In *Pacific-Rim Symposium on Image and Video Technology*, pages 100–113. Springer, 2007. 2
- [13] S. Korman, D. Reichman, G. Tsur, and S. Avidan. Fast-match: Fast affine template matching. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2331–2338, 2013. 2
- [14] C. Ma, J.-B. Huang, X. Yang, and M.-H. Yang. Hierarchical convolutional features for visual tracking. In *IEEE International Conference on Computer Vision (ICCV)*, pages 3074–3082, 2015. 6
- [15] I. Olonetsky and S. Avidan. TreeCANN-kd tree coherence approximate nearest neighbor algorithm. In *European Conference on Computer Vision (ECCV)*, pages 602–615. Springer, 2012. 6
- [16] C. F. Olson. Maximum-likelihood image matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(6):853–857, 2002. 2
- [17] S. Oron, A. Bar-Hillel, D. Levi, and S. Avidan. Locally orderless tracking. *International Journal of Computer Vision*, 111(2):213–228, 2015. 2
- [18] S. Oron, T. Dekel, T. Xue, W. T. Freeman, and S. Avidan. Best-buddies similarity-robust template matching using mutual nearest neighbors. *arXiv preprint arXiv:1609.01571*, 2016. 1, 2, 4, 5, 6, 7
- [19] W. Ouyang, F. Tombari, S. Mattoccia, L. Di Stefano, and W.-K. Cham. Performance evaluation of full search equivalent pattern matching algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(1):127–143, 2012. 2
- [20] O. Pele and M. Werman. Robust real-time pattern matching using bayesian sequential hypothesis testing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(8):1427–1443, 2008. 2
- [21] P. Pérez, C. Hue, J. Vermaak, and M. Gangnet. Color-based probabilistic tracking. In *European Conference on Computer Vision (ECCV)*, pages 661–675. Springer, 2002. 2
- [22] B. G. Shin, S.-Y. Park, and J. J. Lee. Fast and robust template matching algorithm in noisy image. In *International Conference on Control, Automation and Systems.*, pages 6–9. IEEE, 2007. 2
- [23] A. Sibiryakov. Fast and high-performance template matching method. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1417–1424, 2011. 2
- [24] D. Simakov, Y. Caspi, E. Shechtman, and M. Irani. Summarizing visual data using bidirectional similarity. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2008. 2, 5
- [25] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015. 6
- [26] Y. Tian and S. G. Narasimhan. Globally optimal estimation of nonrigid image distortion. *International Journal of Computer Vision*, 98(3):279–302, 2012. 2
- [27] D.-M. Tsai and C.-H. Chiang. Rotation-invariant pattern matching using wavelet decomposition. *Pattern Recognition Letters*, 23(1):191–201, 2002. 2
- [28] Y. Wu, J. Lim, and M.-H. Yang. Online object tracking: A benchmark. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2411–2418, 2013. 6