Best-Buddies Similarity - Robust Template Matching using Mutual Nearest Neighbors

Shaul Oron, Tali Dekel, Tianfan Xue, William T. Freeman, Shai Avidan

Abstract—We propose a novel method for template matching in unconstrained environments. Its essence is the Best-Buddies Similarity (BBS), a useful, robust, and parameter-free similarity measure between two sets of points. BBS is based on counting the number of Best-Buddies Pairs (BBPs)—pairs of points in source and target sets, where each point is the nearest neighbor of the other. BBS has several key features that make it robust against complex geometric deformations and high levels of outliers, such as those arising from background clutter and occlusions. We study these properties, provide a statistical analysis that justifies them, and demonstrate the consistent success of BBS on a challenging real-world dataset while using different types of features.

1 INTRODUCTION

Finding a template patch in a target image is a core component in a variety of computer vision applications such as object detection, tracking, image stitching and 3D reconstruction. In many real-world scenarios, the template—a bounding box containing a region of interest in the source image —undergoes complex deformations in the target image: the background can change and the object may undergo nonrigid deformations and partial occlusions.

Template matching methods have been used with great success over the years but they still suffer from a number of drawbacks. Typically, all pixels (or features) within the template and a candidate window in the target image are taken into account when measuring their similarity. This is undesirable in some cases, for example, when the background behind the object of interest changes between the template and the target image (see Fig. 1). In such cases, the dissimilarities between pixels from different backgrounds may be arbitrary, and accounting for them may lead to false detections of the template (see Fig. 1(b)).

In addition, many template matching methods assume a specific parametric deformation model between the template and the target image (e.g., rigid, affine transformation, etc.). This limits the type of scenes that can be handled, and may require estimating

- S. Oron, Department of Electrical Engineering, Tel-Aviv University E-mail: shauloro@post.tau.ac.il
- T. Dekel, MIT Computer Science and Artificial Intelligence Lab, Google E-mail: tdekel@google.com
- T. Xue, MIT Computer Science and Artificial Intelligence Lab E-mail: tfxue@mit.edu
- W.T. Freeman, MIT Computer Science and Artificial Intelligence Lab, Google
 - E-mail: billf@mit.edu S Avidan Department of Flect
- S. Avidan, Department of Electrical Engineering, Tel-Aviv University E-mail: avidan@eng.tau.ac.il



Figure 1: **Best-Buddies Similarity (BBS) for Template Matching:** (a), The template, marked in green, contains an object of interest against a background. (b), The object in the target image undergoes complex deformation (background clutter and large geometric deformation); the detection results using different similarity measures are marked on the image (see legend); our result is marked in blue. (c), The Best-Buddies Pairs (BBPs) between the template and the detected region are mostly found the object of interest and not on the background; each BBP is connected by a line and marked in a unique color.

a large number of parameters when complex deformations are considered.

In order to address these challenges, we introduce a novel similarity measure termed *Best-Buddies Similarity (BBS)*, and show that it can be applied successfully to template matching *in the wild*. In order to compute the BBS we first represent both the template patch and candidate query patches as point sets in \mathbb{R}^d . Then, instead of searching for a parametric deformation between template and candidate we directly measure the similarity between these point sets. We analyze key features of BBS, and perform extensive evaluation of its performance compared to a number of commonly used alternatives on challenging datasets.

BBS measures the similarity between two sets of points in \mathbb{R}^d . A key feature of this measure is that it relies only on a subset (usually small) of pairs of points – the *Best-Buddies Pairs* (*BBPs*). A pair of points is considered a BBP if the points are mutual nearest neighbors, i.e. each point is the nearest neighbor of the other in the corresponding point set. BBS is then taken to be the fraction of BBPs out of all the points in the set.

Albeit simple, this measure turns out to have important and

nontrivial properties. Because BBS counts only the pairs of points that are best buddies, it is robust to significant amounts of outliers. Another, less obvious property is that the BBS between two point sets is maximal when the points are drawn from the same distribution, and drops sharply as the distance between the distributions increases. In other words, if two points are BBP, they were likely drawn from the same distribution. We provide a statistical formulation of this observation, and analyze it numerically in the 1D case for point sets drawn from distinct Gaussian distributions (often used as a simplified model for natural images).

Modeling image data as distributions, i.e. using histograms, was successfully applied to many computer vision tasks, due to its simple yet effective non-parametric representation. A prominent distance measure between histograms is the Chi-Square (χ^2) distance, in which contributions of different bins, to the similarity score, are proportional to the overall probability stored in those bins.

In this work we show that for sufficiently large sets, BBS converges to the χ^2 distance between distributions. However, unlike χ^2 computing BBS is done directly on the raw data without the need to construct histograms. This is advantageous as it alleviates the need to choose the histogram bin size. Another benefit is the ability to work with high dimensional representation, such as Deep features, for which constructing histograms is not tractable.

More generally, we show a link between BBS and a well known statistical measure. This provides additional insight into the statistical properties of mutual nearest neighbors, and also sheds light on the ability of BBS to reliably match features coming from the same distribution, in the presence of outliers.

We apply the BBS measure to template matching by representing both the template and each of the candidate image regions as point sets in a joint location-appearance space. To this end, we use normalized coordinates for location and experiment with both color as well as Deep features for appearance (although, BBS is not restricted to these specific choices). BBS is used to measure the similarity between the two sets of points in these spaces. The aforementioned properties of BBS now readily apply to template matching. That is, pixels on the object of interest in both the template and the candidate patch can be thought of as originating from the same underlying distribution. These pixels in the template are likely to find best buddies in the candidate patch, and hence would be considered as inliers. In contrast, pixels that come from different distributions, e.g., pixels from different backgrounds, are less likely to find best buddies, and hence would be considered outliers (see Fig. 1(c)). Given this important property, BBS bypasses the need to explicitly model the underlying object appearance and deformation.

To summarize, the main contributions of this paper are: (a) introducing BBS – a useful, robust, parameter-free measure for template matching in unconstrained environments, (b) analysis providing theoretical justification of its key features and linking BBS with the Chi-Square distance, and (c) extensive evaluation on challenging real data, using different feature representations, and comparing BBS to a number of commonly used template matching methods. A preliminary version of this paper appeared in CVPR 2015 [1].

2 RELATED WORK

Template matching algorithms depend heavily on the similarity measure used to match the template and a candidate window in



Signal F

Signals: First row, Signal P consists of "foreground" points drawn from a normal distribution, $N(\mu_1, \sigma_1)$, marked in blue; and "background" points drawn from $N(\mu_2, \sigma_2)$, marked in red. Similarly, the points in the second signal Q are drawn from the same distribution $N(\mu_1, \sigma_1)$, and a different background distribution $N(\mu_3, \sigma_3)$. The color of points is for illustration only, i.e., BBS does not know which point belongs to which distribution. Second row, only the BBPs between the two signals which are mostly found between foreground points.

the target image. Various similarity measures have been used for this purpose. The most popular are the Sum of Squared Differences (SSD), Sum of Absolute Differences (SAD) and Normalized Cross-Correlation (NCC), mostly due to their computational efficiency [2]. Different variants of these measures have been proposed to deal with illumination changes and noise [3], [4].

Another family of measures is composed of robust error functions such as M-estimators [5], [6] or Hamming-based distance [7], [8], which are less affected by additive noise and 'salt and paper' outliers than cross correlation related methods. However, all the methods mentioned so far assume a strict rigid geometric deformation (only translation) between the template and the target image, as they penalize pixel-wise differences at corresponding positions in the template and the query region.

A number of methods extended template matching to deal with parametric transformations (e.g., [9], [10]). Recently, Korman *et al.* [11] introduced a template matching algorithm under 2D affine transformation that guarantees an approximation to the globally optimal solution. Likewise, Tian and Narasimhan [12] find a globally optimal estimation of nonrigid image distortions. However, these methods assume a one-to-one mapping between the template and the query region for the underlying transformation. Thus, they are prone to errors in the presence of many outliers, such as those caused by occlusions and background clutter. Furthermore, these methods assume a parametric model for the distortion geometry, which is not required in the case of BBS.

Measuring the similarity between color histograms, known as Histogram Matching (HM), offers a non-parametric technique for

Signal Q

dealing with deformations and is commonly used in visual tracking [13], [14]. Yet, HM completely disregards geometry, which is a powerful cue. Further, all pixels are evenly treated. Other tracking methods have been proposed to deal with cluttered environments and partial occlusions [15], [16]. But unlike tracking, we are interested in detection in a single image, which lacks the redundant temporal information given in videos.

Olson [17] formulated template matching in terms of maximum likelihood estimation, where an image is represented in a 3D location-intensity space. Taking this approach one step further, Oron *et al.* [18] use xyRGB space and reduced template matching to measuring the EMD [19] between two point sets. Unlike EMD, BBS does not require 1 : 1 matching. It therefore does not have to account for all the data when matching, making it more robust to outliers.

The BBS is a bi-directional measure. The importance of such two-side agreement has been demonstrated by the Bidirectional similarity (BDS) in [20] for visual summarization. Specifically, the BDS was used as a similarity measure between two images, where an image is represented by a set of patches. The BDS sums over the distances between each patch in one image to its nearest neighbor in the other image, and vice versa.

In the context of image matching, another widely used measure is the Hausdorff distance [21]. To deal with occlusions or degradations, Huttenlocher *et al.* [21] proposed a fractional Hausdorff distance in which the K^{th} farthest point is taken instead of the most farthest one. Yet, this measure highly depends on K that needs to be tuned. Alternatively, Dubuisson and Jain [22] replace the max operator with sum, which is similar to the way BDS is defined.

In contrast, the BBS is based on a *count* of the BBPs, and makes only implicit use of their actual distance. Moreover, the BDS does not distinguish between inliers and outliers. These proprieties makes the BBS a more robust and reliable measure as demonstrated by our experiments.

We show a connection between BBS and the Chi-Square (χ^2) distance used as a distance measure between distributions (or histograms). Chi-Square distance comes from the χ^2 test-statistic [23] where it is used to test the fit between a distribution and observed frequencies. χ^2 was successfully applied to a wide range of computer vision tasks such as texture and shape classification [24], [25], local descriptors matching [26], and boundary detection [27] to name a few.

It is worth mentioning, that the term *Best Buddies* was used by Pomeranz *et al.* [28] in the context of solving jigsaw puzzles. Specifically, they used a metric similar to ours in order to determine if a pair of pieces are compatible with each other.

The power of mutual nearest neighbors was previously leveraged for tasks such as image matching [29], classification of images [30] and natural language data [31], clustering [32] and more. In this work we demonstrate its use for template matching while providing some new statistical analysis.

3 BEST-BUDDIES SIMILARITY

Our goal is to match a template to a given image, in the presence of high levels of outliers (i.e., background clutter, occlusions) and nonrigid deformation of the object of interest. We follow the traditional sliding window approach and compute the Best-Buddies Similarity (BBS) between the template and every window (of the size of the template) in the image. In the following, we give



Figure 3: **BBS template matching results.** Three toys examples are shown: (A) cluttered background, (B) occlusions, (C) nonrigid deformation. The template (first column) is detected in the target image (second column) using the BBS; the results using BBS are marked in a blue. The likelihood maps (third column) show well-localized distinct modes. The BBPs are shown in last column. See text for more details.

a general definition of BBS and demonstrate its key features via simple intuitive toy examples. We then statistically analyze these features in Sec. 4.

General Definition: BBS measures the similarity between two sets of points $P = \{p_i\}_{i=1}^{N_P}$ and $Q = \{q_i\}_{i=1}^{N_Q}$, where $p_i, q_i \in \mathbb{R}^d$. The BBS is the fraction of *Best-Buddies Pairs* (BBPs) between the two sets. Specifically, a pair of points $\{p_i \in P, q_j \in Q\}$ is a BBP if p_i is the nearest neighbor of q_j in the set Q, and vice versa. Formally,

$$bb(p_i, q_j, P, Q) = \begin{cases} 1 & \text{NN}(p_i, Q) = q_j \land \text{NN}(q_j, P) = p_i \\ 0 & \text{otherwise} \end{cases}$$
(1)

where, NN $(p_i, Q) = \operatorname*{argmin}_{q \in Q} d(p_i, q)$, and $d(p_i, q)$ is some distance measure. The BBS between the point sets P and Q is given by:

$$BBS(P,Q) = \frac{1}{\min\{N_P, N_Q\}} \cdot \sum_{i=1}^{N_P} \sum_{j=1}^{N_Q} bb(p_i, q_j, P, Q).$$
(2)

The key properties of the BBS are: 1) it relies only on a (usually small) subset of matches i.e., pairs of points that are BBPs, whereas the rest are considered as outliers. 2) BBS finds the bidirectional inliers in the data without any prior knowledge on the data or its underlying deformation. 3) BBS uses *rank*, i.e., it counts the number of BBPs, rather than using the actual distance values.

To understand why these properties are useful, let us consider a simple 2D case of two point sets P and Q. The set P consist of 2D points drawn from two different normal distributions, $N(\mu_1, \Sigma_1)$, and $N(\mu_2, \Sigma_2)$. Similarly, the points in Q are drawn from the same distribution $N(\mu_1, \Sigma_1)$, and a different distribution $N(\mu_3, \Sigma_3)$ (see first row in Fig. 2). The distribution $N(\mu_1, \Sigma_1)$ can be treated as a *foreground* model, whereas $N(\mu_2, \Sigma_2)$ and $N(\mu_3, \Sigma_3)$ are two different *background* models. As can be seen in Fig. 2, the BBPs are mostly found between the foreground points in P and Q. For set P, where the foreground and background points are well separated, 95% of the BBPs are foreground points. For set Q, despite the significant overlap between foreground and background, 60% of the BBPs are foreground points.

This example demonstrates the robustness of BBS to high



Figure 4: The expectation of BBS in the 1D Gaussian case: Two point sets, P and Q, are generated by sampling points from N(0, 1), and $N(\mu, \sigma)$, respectively. (a), the approximated expectation of BBS(P,Q) as a function of σ (x-axis), and μ (y-axis).(b)-(c), the expectation of SSD(P,Q), and SAD(P,Q), respectively. (d), the expectation of BBS as a function of μ plotted for different σ .

levels of outliers in the data. BBS captures the foreground points and does not force the background points to match. In doing so, BBS sidesteps the need to model the background/foreground parametrically or have a prior knowledge of their underlying distributions. This shows that a pair of points $\{p, q\}$ is more likely to be BBP if p and q are drawn from the same distribution. We formally prove this general argument for the 1D case in Sec. 4. With this observations in hand, we continue with the use of BBS for template matching.

3.1 BBS for Template Matching

To apply BBS to template matching, one needs to convert each image patch to a point set in \mathbb{R}^d . Following [18], we use a joint spatial-appearance space which was shown to be useful for template matching. BBS, as formulated in equation (2), can be computed for any arbitrary feature space and for any distance measure between point pairs. In this paper we focus on two specific appearance representations: (i) using color features, and (ii) using Deep features taken from a pretrained neural net. Using such Deep features is motivated by recent success in applying features taken from deep neural nets to different applications [33], [34]. A detailed description of each of these feature spaces is given in Section 5.1.

Following the intuition presented in the 2D Gaussian example (see Fig. 2), the use of BBS for template matching allows us to overcome several significant challenges such as background clutter, occlusions, and nonrigid deformation of the object. This is demonstrated in three synthetic examples shown in Fig. 3. The templates A and B include the object of interest in a cluttered background, and under occlusions, respectively. In both cases the templates are successfully matched to the image despite the high level of outliers. As can be seen, the BBPs are found only on the object of interest, and the BBS likelihood maps have a distinct mode around the true location of the template. In the third example, the template C is taken to be a bounding box around the forth duck in the original image, which is removed from the searched image using inpating techniques. In this case, BBS matches the template to the fifth duck, which can be seen as a nonrigid deformed version of the template. Note that the BBS does not aim to solve the pixel correspondence. In fact, the BBPs are not necessarily semantically correct (see third row in Fig. 3), but rather pairs of points that likely originated from the same distribution. This property, which we next formally analyze, helps us deal with complex visual and geometric deformations in the presence of outliers.

4 ANALYSIS

So far, we have empirically demonstrated that the BBS is robust to outliers, and results in well-localized modes. In what follows, we give a statistical analysis that justifies these properties, and explains why using the count of the BBP is a good similarity measure. Additionally, we show that for sufficiently large sets BBS converges to the well known Chi-Square. This connection with χ^2 provides additional insight into the way BBS handles outliers.

4.1 Expected value of BBS

We begin with a simple mathematical model in 1D, in which an "image" patch is modeled as a set of points drawn from a general distribution. Using this model, we derive the expectation of BBS between two sets of points, drawn from two given distributions $f_P(p)$ and $f_Q(q)$, respectively. We then analyze numerically the case in which $f_P(p)$, and $f_Q(q)$ are two different normal distributions. Finally, we relate these results to the multi-dimentional case. We show that the BBS distinctively captures points that are drawn from similar distributions. That is, we prove that the likelihood of a pair of points being BBP, and hence the expectation of the BBS, is maximal when the points in both sets are drawn from the same distribution, and drops sharply as the distance between the two normal distributions increases.

One-dimentional Case: Following Eq. 2, the expectation BBS(P,Q), over all possible samples of P and Q is given by:

$$E[BBS(P,Q)] = \frac{1}{\min\{N_P, N_Q\}} \sum_{i=1}^{N_P} \sum_{j=1}^{N_Q} E[bb_{i,j}(P,Q)],$$
(3)

where $bb_{i,j}(P,Q)$ is defined in Eq. 1. We continue with computing the expectation of a pair of points to be BBP, over all possible samples of P and Q, denoted by E_{BBP} . That is,

$$E_{\text{BBP}} = \iint_{P,Q} bb_{i,j}(P,Q) \operatorname{Pr}\{P\} \operatorname{Pr}\{Q\} dP dQ, \qquad (4)$$

This is a multivariate integral over all points in P and Q. However, assuming each point is independent of the others this integral can be simplified as follows.

Claim:

$$E_{\text{BBP}} = \int_{-\infty}^{\infty} (F_Q(p^-) + 1 - F_Q(p^+))^{N_Q - 1} \cdot (F_P(q^-) + 1 - F_P(q^+))^{N_P - 1} f_P(p) f_Q(q) dp dq,$$
(5)

where, $F_P(x)$, and $F_Q(x)$ denote the CDFs of P and Q, respectively. That is, $F_P(x) = \Pr\{p \le x\}$. And, $p^- = p - d(p,q)$, $p^+ = p + d(p,q)$, and q^+, q^- are similarly defined.

Proof: Due to the independence between the points, the integral in Eq.4 can be decoupled as follows:

$$E_{\text{BBP}} = \int_{p_1} \cdots \int_{p_{N_P}} \int_{q_1} \cdots \int_{q_{N_Q}} bb_{i,j}(P,Q) \prod_{k=1}^{N_P} f_P(p_k) \prod_{l=1}^{N_Q} f_Q(q_l) dP dQ$$
(6)

With abuse of notation, we use $dP = dp_1 \cdot dp_2 \cdots dp_N$, and $dQ = dq_1 \cdot dq_2 \cdots dq_M$. Let us consider the function $bb_{i,j}(P,Q)$ for a given realization of P and Q. By definition, this indicator function equals 1 when p_i and q_j are nearest neighbors of each

other, and zero otherwise. This can be expressed in terms of the distance between the points as follows:

$$bb_{i,j}(P,Q) = \prod_{\substack{N_P \\ k \neq i,k=1}}^{N_P} \mathbb{I}[d(p_k,q_j) > d(p_i,q_j)] \prod_{\substack{l \neq j,l=1}}^{N_Q} \mathbb{I}[d(q_l,p_i) > d(p_i,q_j)]$$
(7)

where \mathbb{I} is an indicator function. It follows that for a given value of p_i and q_j , the contribution of p_k to the integral in Eq. 6 can be decoupled. Specifically, we define:

$$Cp_k = \int_{-\infty}^{\infty} \mathbb{I}[d(p_k, q_j) > d(p_i, q_j)] f_P(p_k) dp_k \tag{8}$$

Assuming $d(p,q)=\sqrt{(p-q)^2}=|p-q|,$ the latter can be written as:

$$Cp_k = \int_{-\infty}^{\infty} \mathbb{I}[p_k < q_j^- \lor p_k > q_j^+] f_P(p_k) dp_k \tag{9}$$

where $q_j^- = q_j - d(p_i, q_j)$, $q_j^+ = q_j + d(p_i, q_j)$. Since $q_j^- < q_j^+$, it can be easily shown that Cp_k can be expressed in terms of $F_P(x)$, the CDF of P:

$$Cp_k = F_P(q_j^-) + 1 - F_P(q_j^+)$$
(10)

The same derivation hold for computing Cq_l , the contribution of q_l to the integral in Eq. 6, given p_i , and q_j . That is,

$$Cq_l = F_Q(p_i^-) + 1 - F_Q(p_i^+)$$
(11)

where p_i^-, p_i^+ are similarly defined and $F_Q(x)$ is the CDF of Q. Note that Cp_k and Cq_l depends only on p_i and q_j and on the underlying distributions. Therefore, Eq. 6 results in:

$$E_{\text{BBP}} = \iint_{p_i, q_j} dp_i dq_j f_P(p_i) f_Q(q_j) \prod_{\substack{k=1, k \neq i \\ p_i, q_j}}^{N_P} Cp_k \prod_{\substack{l=1, l \neq j \\ l=1, l \neq j}}^{N_Q} Cq_l$$
$$= \iint_{p_i, q_j} dp_i dq_j f_P(p_i) f_Q(q_j) Cp_k^{N_P-1} Cq_l^{N_Q-1}$$
(12)

Substituting the expressions for Cp_k and Cq_l in Eq. 12, and omitting the subscripts i, j for simplicity, result in Eq. 5, which completes the proof.

In general, the integral in Eq. 5 does not have a closed form solution, but it can be solved numerically for selected underlying distributions. To this end, we proceed with Gaussian distributions, which are often used as simple statistical models of image patches. We then use Monte-Carlo integration to approximate $E_{\rm BBP}$ for discrete choices of parameters μ and σ of Q in the range of [0, 10] while fixing the distribution of P to have $\mu = 0, \sigma = 1$. We also fixed the number of points to $N_P = N_Q = 100$. The resulting approximation for $E_{\rm BBP}$ as a function of the parameters μ, σ is shown in Fig. 4, on the left. As can be seen, $E_{\rm BBP}$ is the highest at $\mu = 0, \sigma = 1$, i.e., when the points are drawn from the same distribution, and drops rapidly as the the underlying distribution of Q deviates from N(0, 1).

Note that E_{BBP} does not depends on p and q (because of the integration, see Eq. 5. Hence, the expected value of the BBS between the sets (Eq. 3) is given by:

$$E[BBS(P,Q)] = c \cdot E_{BBP} \tag{13}$$

where $c = \frac{N_P N_Q}{\min\{N_P, N_Q\}}$ is constant.

We can compare the BBS to the expectation of SSD, and SAD. The expectation of the SSD has a closed form solution given by:

$$E[SSD(P,Q)] = \iint_{-\infty}^{\infty} (p-q)^2 f_P(p) f_Q(q|k) dp dq = 1 + \mu^2 + \sigma^2.$$
(14)

Replacing $(p-q)^2$ with |p-q| results in the expression of the SAD. In this case, the expected value reduces to the expectation of the Half-Normal distribution and is given by:

$$E[\text{SAD}(\mathbf{P},\mathbf{Q})] = \frac{1}{\sqrt{2\pi}} \sigma_K \exp^{-\mu^2/(2\sigma^2)} + \mu (1 - 2f_P(-\mu/\sigma))$$
(15)

Fig. 4(b)-(c) shows the maps of the expected values for $1 - \text{SSD}_n(P, Q)$, and $1 - \text{SAD}_n(P, Q)$, where SSD_n , SAD_n are the expectation of SSD and SAD, normalized to the range of [0,1]. As can be seen, the SSD and SAD results in a much wider spread around their mode. Thus, we have shown that the likelihood of a pair of points to be a BBP (and hence the expectation of the BBS) is the highest when P and Q are drawn from the same distribution and drops sharply as the distance between the distributions increases. This makes the BBS a robust and distinctive measure that results in well-localized modes.

Multi-dimensional Case: With the result of the 1D case in hand, we can bound the expectation of BBS when P and Q are sets of multi-dimensional points, i.e., $p_i, q_j \in \mathbb{R}^d$.

If the d-dimensions are uncorrelated (i.e., the covariance matrices are diagonals in the Gaussian case), a sufficient (but not necessary) condition for a pair of points to be BBP is that the point would be BBP in each of the dimensions. In this case, the analysis can be done for each dimension independently similar to what was done in Eq. 5. The expectation of the BBS in the multi-dimensional case is then bounded by the product of the expectations in each of the dimensions. That is,

$$E_{\rm BBS} \ge \prod_{i=1}^{a} E_{\rm BBS}^{i},\tag{16}$$

where E_{BBS}^i denote the expectation of BBS in the *i*th dimension. This means that the BBS is expected to be more distinctive, i.e., to drop faster as *d* increases. Note that if a pair of points is not a BBP in one of the dimensions, it does not necessarily imply that the multi-dimentional pair is not BBP. Thus, this condition is sufficient but not necessary.

4.2 BBS and Chi-Square

Chi-Square is often used to measure the distance between histograms of two sets of features. For example, in face recognition, χ^2 is used to measure the similarity between local binary patterns (LBP) of two faces [35], and it achieves superior performance relative to other distance measures.

In this section, we will discuss the connection between this well known statistical distance measure and BBS. Showing that, for sufficiently large point sets, BBS converges to the χ^2 distance.

We assume, as before, that point sets P and Q are drawn i.i.d. from 1D distribution functions $f_P(p)$ and $f_Q(q)$ respectively. We begin by considering the following lemma:

Lemma 1. Given a point $p_i = p$ in P, let $Pr[bb(p_i = p; P, Q)]$ be the probability that p_i has a best buddy in Q. Then we have:

$$\lim_{N \to +\infty} \Pr[bb(p_i = p; P, Q)] = \frac{f_Q(p)}{f_P(p) + f_Q(p)},$$
 (17)

For the proof of the lemma see appendix A. Intuitively, if there are many points from P in the vicinity of point p, but only few points from Q, i.e. $f_P(p)$ is large but $f_Q(p)$ is small. It is then



Figure 5: Finding a Best-Buddy: We illustrate how the underlying density functions affect the probability that a point p (bold red circle) has a best buddy. (a) Points from set P (red circles) are dense but points from set Q (blue cross) are sparse. Although q is the nearest neighbor of p in Q, p is not the nearest neighbor of q in P (p' is closer). (b) Points from set Q are dense and points from set P are sparse. In this case, p and q are best buddies, as p is the closest point to q.

hard to find a best buddy in Q for p, as illustrated in Figure 5(a). Conversely, if there are few points from P in the vicinity of p but many points from Q, i.e. $f_P(p)$ is small and $f_Q(p)$ is large. In that case, it is easy for p to find a best buddy, as illustrated in Figure 5(b).

A synthetic experiment illustrating the lemma is shown in Figure 6. Two Gaussian mixtures, each consisting of two 1D-Gaussian distributions are used (Figure 6(a)). Sets P and Q are sampled from these distributions (each set from a different mixture model). We then empirically calculate the probability that a certain point p_i from set P has a best buddy in set Q for different set sizes, ranging from 10 to 10000 points, Figure 6(b) . As the sets size increases, the empirical probability converges to the analytical value given by Lemma 1, marked by the dashed black line. Note how the results agree with our intuition. For example, at p = 0, $f_P(p)$ is very large but $f_Q(p)$ is almost 0, such that $Pr[bbs(p_i; P, Q)]$ is almost 0. At p = 5, however, $f_P(p)$ is very small and $f_Q(p)$ is almost 0, so $Pr[bbs(p_i; P, Q)]$ is almost 1.

Lemma 1 assumes the value of the point p_i is fixed. However, we need to consider that p_i itself is also sampled from the distribution $f_P(p)$, in which case the probability this point has a best buddy is:

$$Pr[bb(p_i; P, Q)] = \int_{p=-M}^{M} f_P(p) \cdot Pr(p_i = p; P, Q) dp. = \int_{p=-M}^{M} \frac{f_Q(p)f_P(p)}{f_P(p) + f_Q(p)} dp$$
(18)

Where we assume both density functions are defined on the closed interval [-M, M].

We are now ready to show that BBS converges to Chi-Square,

Theorem 1. Suppose both density functions are defined on a close interval [-M, M], non-zero and Lipschitz continuous ¹. That is,

1)
$$\forall p, q, f_P(p) \neq 0, f_Q(q) \neq 0$$

2) $\exists A > 0, \forall p, q, h, s.t. |f_P(p+h) - f_P(p)| < A|h| and |f_Q(q+h) - f_Q(q)| < A|h|,$

1. Note that most of density functions, like the density function of a Gaussian distribution, are non-zero and Lipschitz continuous in their domain.

then we have,

$$\lim_{N \to +\infty} E[BBS(P,Q)] = \int_{p=-M}^{M} \frac{f_P(p)f_Q(p)}{f_P(p) + f_Q(p)} dp$$

$$= \frac{1}{2} - \frac{1}{4}\chi^2(f_P, f_q),$$
(19)

where $\chi^2(f_p, f_q)$ is the Chi-Square distance between two distributions.

To see why this theorem holds, consider the BBS measure between two sets, P and Q. When the two sets have the same size, the BBS measure equals to the fraction of points in P that have a best buddy, that is $BBS(P,Q) = \frac{1}{N} \sum_{i=1}^{N} bbs(p_i; P, Q)$. Taking expectation on both sides of the equation, we get:

$$E[BBS(P,Q)] = \frac{1}{N} \sum_{i=1}^{N} E[bbs(p_i; P,Q)]$$

$$= \frac{1}{N} \cdot N \cdot E[bbs(p_i; P,Q)]$$

$$= \int_{p=-M}^{M} \frac{f_Q(p)f_P(p)}{f_P(p) + f_Q(p)} dp.$$
(20)

Where for the last equality we used lemma 1. This completes the proof of Theorem 1.

The theorem helps illustrate why BBS is robust to outliers. To see this, consider the signals in Figure 6(a). As can be seen f_P and f_Q are both Gaussian mixtures. Let us assume that the Gaussian with mean -5 represents the foreground (in both signals), i.e. $\mu_{fg} = -5$, and that the second Gaussian in each mixture represents the background, i.e. $\mu_{bg1} = 0$ and $\mu_{bg2} = 5$. Note how, $f_P(p)$ is very close to zero around μ_{bg2} and similarly $f_Q(q)$ is very close to zero around μ_{bg1} . This means that the background distributions will make very little contribution to the χ^2 distance, as the numerator $f_P(p)f_Q(q)$ of Eq. 19 is very close to 0 in both cases.

We note that using BBS has several advantages compared to using χ^2 . One such advantage is that BBS does not require binning data into histograms. It is not trivial to set the bin size, as it depends on the distribution of the features. A second advantage is the ability to use high dimensional feature spaces. The computational complexity and amount of data needed for generating histograms quickly explodes when the feature dimension goes higher. On the contrary, the nearest neighbor algorithm used by BBS can easily scale to high-dimensional features, like Deep features.

5 IMPLEMENTATION DETAILS

In this section we provide information on the specific feature spaces used in our experiments. Additionally, we analyze the computational complexity of BBS and propose a caching scheme allowing for more efficient computation.

5.1 Feature Spaces

In order to perform template matching BBS is computed, exhaustively, in a sliding window. A joint spatial-appearance representation is used in order to convert both template and candidate windows into point sets. For the spatial component normalized xy coordinates within the windows are used. For the appearance descriptor we experiment with both color features as well as Deep features.



Figure 6: Illustrating Lemma 1: Point sets P and Q are sampled iid from the two Gaussian mixtures shown in (a). The probability that a point in set P has a best buddy in set Q is empirically computed for different set sizes (b). When the size of the sets increase, the empirical probability converges to the analytical solution in Lemma 1 (dashed black line).

Color features: When using color features, we break the template and candidate windows into $k \times k$ distinct patches. Each such $k \times k$ patch is represented by its $3 \cdot k^2$ color channel values and xy location of the central pixel, relative to the patch coordinate system. For our toy examples and qualitative experiments RGB color space is used. However, for our quantitative evaluation HSV was used as it was found to produced better results. Both spatial and appearance channels were normalized to the range [0,1]. The point-wise distance measure used with our color features is:

$$d(p_i, q_j) = ||p_i^{(A)} - q_j^{(A)}||_2^2 + \lambda ||p_i^{(L)} - q_j^{(L)}||_2^2$$
(21)

where superscript A denotes a points appearance and superscript L denotes a points location. The parameter $\lambda = 0.25$ was chosen empirically and was fixed in all of our experiments.

Deep features: For our Deep feature we use features taken from the VGG-Deep-Net [36]. Specifically, we take features from two layers of the network, $conv1_2$ (64 features) and $conv3_4$ (256 features). The feature maps from $conv1_2$ are down-sampled twice, using max-pooling, to reach the size of the $conv3_4$ which is down-sampled by a factor of 1/4 with respect to the original image. In this case we treat every pixel in the down-sampled feature maps as a point. Each such point is represented by its xy location in the down-sampled window and its appearance is given by the 320 feature channels. Prior to computing the pointwise distances each feature channel is independently normalized to have zero mean and unit variance over the window. The pointwise distance in this case is:

$$d(p_i, q_j) = \langle p_i^{(A)}, q_j^{(A)} \rangle + exp(-\lambda ||p_i^{(L)} - q_j^{(L)}||_2^2)$$
(22)

where $\langle \cdot, \cdot \rangle$ denotes the inner product operator between feature vectors. Unlike the color features we now want to maximize d rather then minimize it (we can always minimize -d). The parameter $\lambda = 1$ was chosen empirically and was fixed in all of our experiments.

5.2 Complexity

Computing BBS between two point sets $P, Q \in \mathbb{R}^d$, requires computing the distance between each pair of points. That is, constructing a distance matrix D where $[D]_{i,j} = d(p_i, q_j)$. Given D, the nearest neighbor of $p_i \in P$, i.e. $NN(p_i, Q)$, is the minimal element in the i^{th} row of D. Similarly, $NN(q_j, P)$ is the minimal element in the j^{th} column of D. BBS is then computed by counting the number of mutual nearest neighbors (divided by a constant).

In this section we analyze the computational complexity of computing BBS exhaustively for every window in a query image. We then propose a caching scheme, allowing extensive computation reuse which dramatically reduces the computational complexity, trading it off with increased memory complexity.

Naive implementation: For our analysis we consider a target window of size $w \times h$ and some query image I of size $W \times H$. Both represented using a feature space with d feature channels. Let us begin by considering each pixel in our target window as a point in our target point set P and similarly every pixel in some query window is considered as a point in the query point set Q. In this case, $|P| = |Q| = w \cdot h \triangleq l$ and our distance matrices D are of size $l \times l$. Assuming some arbitrary image padding, we have $W \cdot H \triangleq L$ query windows for which BBS has to be computed. Computing all the L distance matrices requires $O(Ll^2d)$. For each such distance matrix we need to find the minimal element in every row and column. The minimum computation for a single row or column is done in O(l) and for the entire matrix in $O(l^2)$. Therefore, the complexity of computing BBS naively for all query windows of image I is,

$$O(Ll^4d) \tag{23}$$

This is a high computational load compared to simpler methods such as sum-of-square-difference (SSD) that require only O(Lld).

Distance computation reuse: When carefully examining the naive scheme above we notice that many pairwise distance computations are performed multiple times. This observation is key to



Figure 7: **BBS results on Real Data:** (a), the templates are marked in green over the input images. (b) the target images marked with the detection results of 6 different methods (see text for more details). BBS results are marked in blue. (c)-(e), the resulting likelihood maps using BBS, EMD and NCC, respectively; each map is marked with the detection result, i.e., its global maxima.

our proposed caching scheme.

Assuming our sliding window works column by column, the first distance matrix in the image has to be fully computed. The second distance matrix, after sliding the query window down by one pixel, has many overlapping distance computations with the previously computed matrix. Specifically, we only have to compute the distances between pixels in the new row added to the query window and the target window. This means we have to recompute only w columns of D and not the entire matrix. Taking this one step further, if we cache all the distance matrices computed along the first image column, then staring from the second matrix in the second column, we would only have to compute the distance between one new candidate pixel and the target window, which means we only have to recompute one column of D which requires only O(l). Assuming W, H >> w, h the majority of distance matrices can be computed in O(l), instead of $O(l^2)$. This means that computing BBS for the entire image I would now require:

$$O(Ll^3 f) \tag{24}$$

Minimum operator load reduction: So far we have shown how caching can be used to reduce the load of building the distance matrices. We now show how additional caching can reduce the computational load of the minimum operator applied to each row and column of D in order to find the BBP.

As discussed earlier, for the majority of query windows we only have to recompute *one* column of D. This means that for all other l - 1 columns we have already computed the minimum. Therefore, we actually obtain the minimum over all columns in just O(l). For the minimum computation along the rows there are two cases to consider. First, that the minimal value, for a certain row, was in the column that was pushed out of D. In this case we would have to find the minimum value for that row, which would require O(l). The second option is that the minimal value of the row was not pushed out and we know where it is from previous computations. In such a case we only have to compare the new element added to the row (by the new column introduced into D) relative to the previous minimum value, this operation requires O(1). Assuming the position of the minimal value along a row is uniformly distributed, on average, there will be only one row where the minimum value needs to be recomputed. To see this consider a set of random variables $\{X_i\}_{i=1}^l$ such that $X_i = 1$ if and only if the minimal value in the *i*'th row of D was pushed out of the matrix when a new column was introduced. Assuming a uniform distribution $X_i \sim Bernoulli(1/l)$. The number of rows for which the minimum has to be recomputed is given by $m = \sum_{i=1}^l X_i$, and the expected number of such rows is,

$$E[m] = E\left[\sum_{i=1}^{l} X_i\right] = \sum_{i=1}^{l} E[X_i] = \sum_{i=1}^{l} \frac{1}{l} = 1$$
(25)

This means, that on average, there will be only one row for which the minimum has to be computed in O(l) time (for other rows only O(1) is required). Therefore, on average, we are able to find the minimum of all rows and columns in D, in O(l) instead on $O(l^2)$. By combining the efficient minimum computation scheme, along with the reuse of distance computations for building D, we reduce the overall BBS complexity over the entire image to,

$$O(Ll^2d) \tag{26}$$

Additional load reduction: When using color features, we note that the actual complexity of computing BBS for the entire image I is even lower due to the use of non-overlapping $k \times k$ patches instead of individual pixels. This means that both image and target windows are sampled on a grid with spacing k which in turn leads to an overall complexity of:

$$O\left(\frac{Ll^2d}{k^4}\right) \tag{27}$$

We note that the reuse schemes presented above cannot be used with our Deep features due to the fact that we normalize the features differently, with respect to each query window. Also the above analysis does not consider the complexity of extracting the Deep features themselves.

6 RESULTS

We perform qualitative as well as extensive quantitative evaluation of our method on real world data. We compare BBS with several measures commonly used for template matching. 1) Sum-of-Square-Difference (SSD), 2) Sum-of-Absolute-Difference (SAD), 3) Normalized-Cross-Correlation (NCC), 4) color Histogram Matching (HM) using the χ^2 distance, 5) Bidirectional Similarity [20] (BDS) computed in the same appearance-location space as BBS.

6.1 Qualitative Evaluation

Four template-image pairs taken from the Web are used for qualitative evaluation. The templates, which were manually chosen, and the target images are shown in Figure 1(a)-(b), and in Figure 7. In all examples, the template drastically changes its appearance due to large geometric deformation, partial occlusions, and change of background.

Detection results, using color features with RGB color space, are presented in Figure 1(a)-(b), and in Figure 7(b), and compared to the above mentioned methods as well as to the Earth Movers Distance [19] (EMD). The BBS is the only method successfully matching the template in all these challenging examples. The confidence maps of BBS, presented in Figure 7(c), show distinct and well-localized modes compared to other methods². The BBPs for the first example are shown in Figure 1(c). As discussed in Sec. 3, BBS captures the bidirectional inliers, which are mostly found on the object of interest. Note that the BBPs, as discussed, are not necessarily true physical corresponding points.

6.2 Quantitative Evaluation

We now turn to the quantitative evaluation. The data for our experiments was generated from a dataset of 100 annotated video sequences³ [37], both color and gray-scale. These videos capture a wide range of challenging scenes in which the objects of interest are diverse and typically undergo nonrigid deformations, photometric changes, motion blur, in/out-of-plane rotation, and occlusions.

Three template matching datasets were randomly sampled from the annotated videos. Each dataset is comprised of templateimage pairs, where each such pair consists of frames f and f+df, where f was randomly chosen. For each dataset a different value of df was used (25, 50 or 100). The ground-truth annotated bounding box in frame f is used as the template, while frame f + df is used as the query image. This random choice of frames creates a challenging benchmark with a wide baseline in both time and space (see examples in Figure 9 and Figure 10). For df = 25, 50 the data sets consist of 270 pairs and for df = 100there are 254 pairs.

BBS using both color (with HSV color space) and Deep features was compared with the 5 similarity measures mentioned above. The ground-truth annotations were used for quantitative evaluation. Specifically, we measure the accuracy of both the top match as well as the top k ranked matches, as follows.

9

Accuracy: was measured using the common bounding box overlap measure: $Acc. = \frac{\operatorname{area}(B_e \cap B_g)}{\operatorname{area}(B_e \cup B_g)}$ where B_e and B_g are the estimated and ground truth bounding boxes, respectively. The ROC curves show the fraction of examples with overlap larger than a threshold $(TH \in [0, 1])$. Mean average precision (mAP) is taken as the area-under-curve (AUC). The success rates, of all methods, were evaluated considering only the global maximum (best mode) prediction as well as considering the best out of the top 3 modes (using non-maximum suppression, NMS).

Results for both color feature and Deep features for the dataset with df = 25 are shown in Figure 8. Overall it can be seen that BBS outperforms competing methods using both color and Deep features. Using color features and considering only the top mode Figure 8(a), BBS outperforms competing methods with a margin ranging from 4.6% compared to BDS to over 30% compared to SSD. When considering the top 3 modes, Figure 8(c), the performance of all methods improves. However, we can clearly see the dominance of BBS, increasing its margin over competing methods. BBS reaches mAP of 0.648 (compared to 0.589 with only the top mode). For example the margin between BBS and BDS, which is the runner up, increases to 5.9%. The increase in performance when considering the top 3 modes suggests that there are cases where BBS is able to produce a mode at the correct target position however this mode might not be the global maximum of the entire map.

Some successful template matching examples, along with the likelihood maps produced by BBS, using the color features, are shown in Figure 9. Notice how BBS can overcome non-rigid deformations of the target.

Typical failure cases are presented in Figure 12. Most of the failure cases using the color features can be attributed to either illumination variations (c), distracting objects with a similar appearance to the target (a)-(b), or cases were BBS matches the background or occluding object rather than the target (d). This usually happens when the target is heavily occluded or when the background region in the target window is very large.

Results using our Deep feature and considering only the top mode are shown in figures Figure 8(b). We note that HM was not evaluated in this case due to the high dimensionality of the feature space. We observe that BBS outperforms the second best methods by only a small margin of 2.4%. Considering the top 3 modes allows BBS to reach mAP of 0.684 increasing its margin relative to competing methods. For example the margin relative to the second best method (SSD) is now 5.2%.

Some template matching examples, along with their associated likelihood maps, using the Deep features, are shown in Figure 10. The Deep features are not sensitive to illumination variations and can capture both low level information as well as higher level object semantics. As can be seen the combination of using Deep features and BBS can deliver superior results due to its ability to explain non-rigid deformations. Note how when using the Deep feature, we can correctly match the bike rider in Figure 10(c) for which color features failed (Figure 12(d)). BBS with Deep features produce very well localized and compact modes compared to when color features are used.

Some typical failure cases when using the Deep features are presented in Figure 13. As for the color features, many failure cases are due to distracting objects with a similar appearance (a)-(b) or cases were BBS matches the background or occluding object (d).

^{2.} Our data and code are publicly available at: http://people.csail.mit.edu/ talidekel/Best-BuddiesSimilarity.html

^{3.} https://sites.google.com/site/benchmarkpami/



10



Figure 8: **Template matching accuracy:** Evaluation of method performance using 270 template-image pairs with df = 25. BBS outperforms competing methods as can be seen in ROC curves showing fraction of examples with overlap greater than threshold values in [0,1]. Top: only best mode is considered. Bottom: best out of top 3 modes is taken. Left: Color features. Right: Deep features. Mean-average-precision (mAP) values taken as area-under-curve are shown in the legend. Best viewed in color.

It is interesting to see that BDS which was the runner up when color features were used come in last when using Deep features switching places with SSD which was worst previously and is now second in line. This also demonstrates the robustness of BBS which is able to successfully use different features. Additionally, we see that overall BBS with Deep features outperforms BBS with color features (a margin of 5.5% with top 3 modes). However, this performance gain requires a significant increased in computational load both since the features have to be extracted and also since the proposed efficient computation scheme cannot be used in this case. It is interesting to see that BBS with color features is able perform as well as SSD with Deep features.

Finally, we note that, when using the color features BBS outperforms HM which uses the χ^2 distance. Although BBS converges to χ^2 for large sets there are clear benefits for using BBS over χ^2 . Computing BBS does not require modeling the distributions (i.e. building normalized histograms) and can be performed on the raw data itself. This alleviates the need to

choose the histogram bin size which is known to be a delicate issue. Moreover, BBS can be performed on high dimensional data, such as our Deep features, for which modeling the underlying distribution is not practical.

The space time baseline: effect on performance was examined using data-sets with different df values (25, 50, 100). Figure 11 shows mAP of competing methods for different values of df. Results using color features are shown on the left and using Deep features on the right. All results were analyzed taking the best out of the top 3 modes. It can be seen that BBS outperforms competing methods for the different df values with the only exception being Deep feature with df = 100 in which case BBS and SSD produce similar results reaching mAP of 0.6.

7 CONCLUSIONS

We have presented a novel similarity measure between sets of objects called the Best-Buddies Similarity (BBS). BBS leverages



Figure 9: **Example results using color features**. Top, input images with annotated template marked in green. Middle, target images and detected bounding boxes (see legend); ground-truth (GT) marked in green (our results in blue). Bottom, BBS likelihood maps. BBS successfully match the template in all these examples.



Figure 10: **Example results using Deep features**. Top, input images with annotated template marked in green. Middle, target images and detected bounding boxes (see legend); ground-truth (GT) marked in green (our results in blue). Bottom, BBS likelihood maps. BBS successfully match the template in all these examples.



Figure 11: Effect of space time baseline: Methods performance evaluated for data sets with different space-time baseline, df = 25, 50 and 100. Left: Color features, Right: Deep features. BBS outperforms competing methods for both feature choices and for all df values. Best viewed in color.



Figure 12: **Example of failure cases using color features**. Top, input images with annotated template marked in green. Bottom, target images and detected bounding boxes (see legend); ground-truth (GT) marked in green (our results in blue). As can be seen, some common failure causes are illumination changes, similar distracting targets or locking onto the background.



Figure 13: **Example of failure cases using Deep features**. Top, input images with annotated template marked in green. Bottom, target images and detected bounding boxes (see legend); ground-truth (GT) marked in green (our results in blue). Some common failure causes are similar distracting targets or locking onto the background.

statistical properties of mutual nearest neighbors and was shown to be useful for template matching in the wild. Key features of BBS were identified and analyzed demonstrating its ability to overcome several challenges that are common in real life template matching scenarios. It was also shown, that for sufficiently large point sets, BBS converges to the Chi-Square distance. This result provides interesting insights into the statistical properties of mutual nearest neighbors, and the advantages of using BBS over χ^2 where discussed.

Extensive qualitative and quantitative experiments on challenging data were performed and a caching scheme allowing for an efficient computation of BBS was proposed. BBS was shown to outperform commonly used template matching methods such as normalized cross correlation, histogram matching and bidirectional similarity. Different types of features can be used with BBS, as was demonstrated in our experiments, where superior performance was obtained using both color features as well as Deep features.

Our method may fail when the template is very small compared to the target image, when similar targets are present in the scene or when the outliers (occluding object or background clutter) cover most of the template. In some of these cases it was shown that BBS can predict the correct position (produce a mode) but non necessarily give it the highest score.

Finally, we note that since BBS is generally defined between sets of objects it might have additional applications in computervision or other fields that could benefit from its properties. A natural future direction of research is to explore the use of BBS as an image similarity measure, for object localization or even for document matching.

APPENDIX A PROOF OF LEMMA 1

Because of independent sampling, all points in Q have equal probability being the best buddy of p. From this we have:

$$Pr[bb(p_i = p; P, Q)] = = \sum_{i=1}^{N} Pr(bb(p, q_i; P, Q) = 1) = N \cdot Pr(bb(p, q; P, Q)),$$
(28)

where q is a point from Q and subscript is dropped for ease of description.

The probability that two points are best buddies is given by:

$$Pr(bb(p_i = p, q; P, Q)) = (F_Q(p^-) + 1 - F_Q(p^+))^{N-1} (F_P(q^-) + 1 - F_P(q^+))^{N-1}.$$
(29)

where $F_P(x)$ and $F_Q(x)$ denote CDFs of these two distributions, that is, $F_P(x) = \Pr\{p \le x\}$. And, $p^- = p - |p - q|, p^+ = p + |p - q|,$ and q^+, q^- are similarly defined. Combining Eq.28 and Eq. 29, the probability that p_i has a best buddy equals to

$$\lim_{N \to +\infty} N \int_{q=-M}^{M} (F_Q(p^-) + 1 - F_Q(p^+))^{N-1} \\ \cdot (F_P(q^-) + 1 - F_P(q^+))^{N-1} f_Q(q) dq.$$
(30)

We denote the signed distance between two points by m = p - q. Intuitively, because the density function are non-zero at any place, when N goes to infinity, the probability that two points $p \in P, q \in Q$ are BBP decreases rapidly as m increases. Therefore, we only need to consider the case when the distance between p and q is very small. Formally, for any positive \overline{m} , changing the integration limits in Eq. 30 from $\int_{p=-M}^{M}$ to $\int_{q=p-\overline{m}}^{p+\overline{m}}$ does not change the result (see Claim 2 in the supplementary material). Then let us break down $F_P(\cdot)$ and $F_Q(\cdot)$ in Eq. 30. Given that the density functions $f_P(p)$ and $f_Q(q)$ are Lipschitz continuous (Condition 2 in Theorem 1), we can assume that they take a constant value in the interval $[p^-, p^+]$, and $[q^-, q^+]$. That is,

$$f_P(p^-) \approx f_P(p^+) \approx f_P(p)$$

$$f_Q(q^-) \approx f_Q(q^+) \approx f_Q(q)$$
(31)

And thus, the expression $F_Q(p^+) - F_Q(p^-)$ can be approximated as follows:

$$F_Q(p^+) - F_Q(p^-) = \int_{p^-}^{p^+} f_Q(q) dq \approx f_Q(q) \cdot (p^+ - p^-) = 2|m| \cdot f_Q(p).$$
(32)

Similarly, $F_p(q^+) - F_P(q^-) \approx 2|m| \cdot f_P(q)$. Note that this approximation can also be obtained using Taylor expansion on $F_p(q^+)$ and $F_p(q^-)$. At last, since p and q are very close to each other, we assume:

$$f_Q(q) \approx f_Q(p).$$
 (33)

Plugging all these approximations (Eq. 32 and Eq. 33) to Eq. 30 and replacing q by m, we get:

Eq. 30 =

$$= \lim_{N \to +\infty} N \int_{m=-\overline{m}}^{\overline{m}} (1 - 2|m| f_Q(p))^{N-1} \\ \cdot (1 - 2|m| f_P(p))^{N-1} f_Q(p) dq$$
(34)

$$= f_Q(p) \lim_{N \to +\infty} N \int_{m = -\overline{m}} \left(1 - 2(f_P(p) + f_Q(p)) |m| + 4f_P(p) f_Q(p) m^2 \right)_{-1}^{N-1} dm$$
(35)

$$=f_Q(p)\lim_{N \to +\infty} N \int_{m=-\overline{m}}^{\overline{m}} \left(1 - 2(f_P(p) + f_Q(p))m\right)^{N-1} dm$$
(36)

It is worth mentioning that the approximated equality in Eq. 32 and Eq. 33 becomes restrict equality when N goes to infinity (for the proof see Claim 3 in the supplementary material). Also, since the distance between two points m is very small, the second order term $4f_P(p)f_Q(p)m^2$ in Eq. 35 is negligible and is dropped in Eq. 36 (for full justification see Claim 4 in the supplementary material).

full justification see Claim 4 in the supplementary material). At last, $\lim_{N \to +\infty} N \int_{m=-\overline{m}}^{\overline{m}} (1 - a|m|)^{N-1} dm = \frac{2}{a}$ (see Claim 1 in supplementary material). Thus Eq. 36 equals to:

$$\frac{f_Q(p)}{f_P(p) + f_Q(p)} \tag{37}$$

which completes the proof of Lemma 1.

Acknowledgments.

This work was supported in part by an Israel Science Foun- dation grant 1556/10, National Science Foundation Robust Intelligence 1212849 Reconstructive Recognition, and a grant from Shell Research and National Science Foundation 1447476 BigData.

REFERENCES

- T. Dekel, S. Oron, S. Avidan, M. Rubinstein, and W. Freeman, "Best buddies similarity for robust template matching," in *Computer Vision* and Pattern Recognition (CVPR), 2015 IEEE Conference on. IEEE, 2015. 2
- [2] W. Ouyang, F. Tombari, S. Mattoccia, L. Di Stefano, and W.-K. Cham, "Performance evaluation of full search equivalent pattern matching algorithms," *PAMI*, 2012. 2
- [3] Y. Hel-Or, H. Hel-Or, and E. David, "Matching by tone mapping: Photometric invariant template matching," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 36, no. 2, pp. 317–330, 2014. [Online]. Available: http://doi.ieeecomputersociety.org/10.1109/TPAMI.2013.138 2
- [4] E. Elboher and M. Werman, "Asymmetric correlation: a noise robust similarity measure for template matching," *Image Processing, IEEE Transactions on*, 2013. 2

- [5] J.-H. Chen, C.-S. Chen, and Y.-S. Chen, "Fast algorithm for robust template matching with m-estimators," *Signal Processing, IEEE Transactions on*, 2003. 2
- [6] A. Sibiryakov, "Fast and high-performance template matching method," in CVPR, 2011. 2
- [7] B. G. Shin, S.-Y. Park, and J. J. Lee, "Fast and robust template matching algorithm in noisy image," in *Control, Automation and Systems, 2007. ICCAS'07. International Conference on*, 2007. 2
- [8] O. Pele and M. Werman, "Robust real-time pattern matching using bayesian sequential hypothesis testing," *PAMI*, 2008. 2
- [9] D.-M. Tsai and C.-H. Chiang, "Rotation-invariant pattern matching using wavelet decomposition," *Pattern Recognition Letters*, 2002. 2
- [10] H. Y. Kim and S. A. De Araújo, "Grayscale template-matching invariant to rotation, scale, translation, brightness and contrast," in *AIVT*. Springer, 2007. 2
- [11] S. Korman, D. Reichman, G. Tsur, and S. Avidan, "Fast-match: Fast affine template matching," in CVPR, 2013. 2
- [12] Y. Tian and S. G. Narasimhan, "Globally optimal estimation of nonrigid image distortion," *IJCV*, 2012. 2
- [13] D. Comaniciu, V. Ramesh, and P. Meer, "Real-time tracking of non-rigid objects using mean shift," in CVPR, 2000. 3
- [14] P. Pérez, C. Hue, J. Vermaak, and M. Gangnet, "Color-based probabilistic tracking," in ECCV 2002, 2002. 3
- [15] C. Bao, Y. Wu, H. Ling, and H. Ji, "Real time robust 11 tracker using accelerated proximal gradient approach," CVPR, 2012. 3
- [16] X. Jia, H. Lu, and M. Yang, "Visual tracking via adaptive structural local sparse appearance model," CVPR, 2012. 3
- [17] C. F. Olson, "Maximum-likelihood image matching," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 6, pp. 853–857, 2002.
 [Online]. Available: http://doi.ieeecomputersociety.org/10.1109/TPAMI. 2002.1008392 3
- [18] S. Oron, A. Bar-Hillel, D. Levi, and S. Avidan, "Locally orderless tracking," *IJCV*, 2014. 3, 4
- [19] Y. Rubner, C. Tomasi, and L. Guibas, "The earth mover's distance as a metric for image retrieval," *IJCV*, 2000. 3, 9
- [20] D. Simakov, Y. Caspi, E. Shechtman, and M. Irani, "Summarizing visual data using bidirectional similarity," in CVPR, 2008. 3, 9
- [21] D. P. Huttenlocher, G. A. Klanderman, and W. J. Rucklidge, "Comparing images using the hausdorff distance," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 15, no. 9, pp. 850–863, 1993. 3
- [22] M.-P. Dubuisson and A. Jain, "A modified hausdorff distance for object matching," in *Pattern Recognition*, 1994. Vol. 1 Conference A: Computer Vision amp; Image Processing., Proceedings of the 12th IAPR International Conference on, vol. 1, Oct 1994, pp. 566–568 vol.1. 3
- [23] G. Snedegor, W. G. Cochran et al., "Statistical methods." Statistical methods., no. 6th ed, 1967. 3
- [24] M. Varma and A. Zisserman, "A statistical approach to material classification using image patch exemplars," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 11, pp. 2032–2047, 2009.
- [25] S. Belongie, J. Malik, and J. Puzicha, "Shape matching and object recognition using shape contexts," *Pattern Analysis and Machine Intelligence*, *IEEE Transactions on*, vol. 24, no. 4, pp. 509–522, 2002. 3
- [26] P.-E. Forssén and D. G. Lowe, "Shape descriptors for maximally stable extremal regions," in *Computer Vision*, 2007. ICCV 2007. IEEE 11th International Conference on. IEEE, 2007, pp. 1–8. 3
- [27] D. R. Martin, C. C. Fowlkes, and J. Malik, "Learning to detect natural image boundaries using local brightness, color, and texture cues," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 26, no. 5, pp. 530–549, 2004. 3
- [28] D. Pomeranz, M. Shemesh, and O. Ben-Shahar, "A fully automated greedy square jigsaw puzzle solver," in *The 24th IEEE Conference* on Computer Vision and Pattern Recognition, CVPR 2011, Colorado Springs, CO, USA, 20-25 June 2011, 2011, pp. 9–16. [Online]. Available: http://dx.doi.org/10.1109/CVPR.2011.5995331 3
- [29] T.-t. Li, B. Jiang, Z.-z. Tu, B. Luo, and J. Tang, *Intelligent Computation in Big Data Era*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2015, ch. Image Matching Using Mutual k-Nearest Neighbor Graph, pp. 276–283. [Online]. Available: http://dx.doi.org/10.1007/978-3-662-46248-5_34 3
- [30] H. Liu, S. Zhang, J. Zhao, X. Zhao, and Y. Mo, "A new classification algorithm using mutual nearest neighbors," in 2010 Ninth International Conference on Grid and Cloud Computing, Nov 2010, pp. 52–57. 3
- [31] K. Ozaki, M. Shimbo, M. Komachi, and Y. Matsumoto, "Using the mutual k-nearest neighbor graphs for semi-supervised classification of natural language data," in *Proceedings of the Fifteenth Conference* on Computational Natural Language Learning, ser. CoNLL '11. Stroudsburg, PA, USA: Association for Computational Linguistics,

2011, pp. 154–162. [Online]. Available: http://dl.acm.org/citation.cfm? id=2018936.2018954_3

- [32] Z. Hu and R. Bhatnagar, "Clustering algorithm based on mutual k-nearest neighbor relationships," *Statistical Analy Data Mining*, vol. 5, no. 2, pp. 110–113, 2012. 3
- [33] C. Ma, J.-B. Huang, X. Yang, and M.-H. Yang, "Hierarchical convolutional features for visual tracking," in *Proceedings of the IEEE International Conference on Computer Vision*), 2015. 4
- [34] L. Wang, W. Ouyang, X. Wang, and H. Lu, "Visual tracking with fully convolutional networks," in *The IEEE International Conference on Computer Vision (ICCV)*, December 2015. 4
- [35] T. Ahonen, A. Hadid, and M. Pietikainen, "Face description with local binary patterns: Application to face recognition," in *PAMI*, vol. 28, no. 12. IEEE, 2006, pp. 2037–2041. 5
- [36] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1556, 2014. 7
- [37] Y. Wu, J. Lim, and M. Yang, "Online object tracking: A benchmark," in CVPR, 2013. 9