

Best Buddies Registration for Point Clouds

Amnon Drory, Tal Shomer, Shai Avidan, and Raja Giryes

Tel Aviv University, Israel

Abstract. We propose new, and robust, loss functions for the point cloud registration problem. Our loss functions are inspired by the Best Buddies Similarity (BBS) measure that counts the number of mutual nearest neighbors between two point sets. This measure has been shown to be robust to outliers and missing data in the case of template matching for images. We present several algorithms, collectively named *Best Buddy Registration (BBR)*, where each algorithm consists of optimizing one of these loss functions with Adam gradient descent. The loss functions differ in several ways, including the distance function used (point-to-point vs. point-to-plane), and how the BBS measure is combined with the actual distances between pairs of points. Experiments on various data sets, both synthetic and real, demonstrate the effectiveness of the BBR algorithms, showing that they are quite robust to noise, outliers, and distractors, and cope well with extremely sparse point clouds. One variant, BBR-F, achieves state-of-the-art accuracy in the registration of automotive lidar scans taken up to several seconds apart, from the KITTI and Apollo-Southbay datasets.

1 Introduction

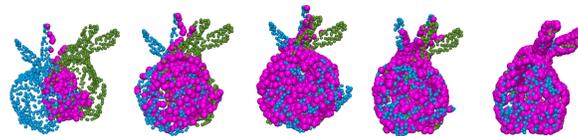


Fig. 1: **Best-Buddies Pairs:** The BBR methods perform point cloud registration by iteratively optimizing (with gradient descent) a loss defined by a set of soft or hard *best buddy pairs*. Blue and Green points denote the original point clouds to register. Best buddy pairs are marked in purple. From left to right: iteration 1, 2, 4, 80 and the last iteration (120).

Point clouds registration is an important task in 3D computer vision. The same object or scene is scanned from two view points, e.g. with a laser scanner, and the goal is to recover the rigid transformation (rotation and translation) that aligns the two scans to each other. Realistic scenarios add complications: measurement noise, occlusions due to the change in view point, and outliers due

to independent motions of free moving objects in the scene (*distractors*). This makes robustness a central issue for point cloud registration algorithms.

Probably the most popular approach to solve the problem is using some variant of the Iterative Closest Point (ICP) algorithm [1]. This method works by iterating between two stages: first match pairs of points between the two clouds, and then apply the transformation that minimizes a loss defined by the distance between the two points in each pair. The simplest version of ICP uses the Euclidean distance between points, but later versions make use of more complex distance measures to achieve faster and more accurate convergence. Some of the most popular and successful variants use local normals to define point-to-plane distance measures. ICP-like methods are typically sensitive to noise, requiring the use of steps such as explicit outlier removal to improve their robustness.

Recently, Oron *et al.* introduced the Best-Buddies Similarity (BBS) measure [2]. BBS counts the number of mutual nearest-neighbors between two point sets. This simple measure was used for template matching between images and proved to be resilient to outliers and occluders. This success motivated us to study how the BBS measure could be adapted to the task of point cloud registration. We suggest several differentiable loss functions inspired by BBS. Our registration algorithms consist of optimizing over these losses with a variant of gradient descent (Adam [3]), to recover the parameters of the aligning transformation. We collectively name the resulting algorithms *Best Buddy Registration (BBR)*, and demonstrate their high level of robustness to noise, occlusions and distractors, as well as an ability to cope with extremely sparse point clouds. Some of the algorithms are able to achieve very high accuracy in noisy settings where robustness is essential.

Deep neural networks (DNN) have increasingly been used for the processing of point clouds lately. BBR can easily be integrated into such DNNs as a registration stage, and be optimized as part of the overall gradient descent optimization of the network.¹ To facilitate this, we implemented BBR in Pytorch², which also makes it possible to run the algorithms on widely available neural network infrastructure, such as GPUs. (See diagram in supplementary material).

The main contributions of this paper are:

1. A robust and accurate point cloud registration algorithm that is especially useful in realistic scenarios with a large time offset between the pair of point clouds, meaning large occlusions and outlier motions.
2. The algorithm naturally fits into the deep learning settings as a component: it can be implemented using operations that already exist in deep learning frameworks, and optimized using Adam gradient descent, which is commonly used for neural network optimization

¹ For example, the DeepMapping network [4] includes a registration stage based on the non-robust Chamfer distance, which could be replaced by BBR.

² <https://github.com/AmnonDrory/BestBuddiesRegistration>

2 Related Work

There are various approaches to the problem of point cloud registration. These algorithms can be divided into classic (i.e., non-deep) and deep methods.

Classic methods. ICP was introduced by Besl and Mckay [1], and Chen and Medioni [5]. See the survey of Rusinkiewicz and Levoy [6] or the recent review of the topic by Pomerelo *et al.* [7]. The basic ICP algorithm deals with point-to-point registration, but already [5] considered point-to-plane registration to improve accuracy. This, however, requires the use of normals as an additional source of information.

Segal *et al.* [8] later extended ICP to a full plane-to-plane formulation and gave it a probabilistic interpretation. Jian and Vemuri [9] proposed a robust point set registration. Their approach reformulated ICP as the problem of aligning two Gaussian mixtures such that a statistical discrepancy measure between the two corresponding mixtures is minimized. It was recently accelerated by Eckart *et al.* [10] who introduced a Hierarchical multi-scale Gaussian Mixture Representation (HGMR) of the point clouds. Similarly, FilterReg [11] is a probabilistic point-set registration method that is considerably faster than alternative methods due to its computationally-efficient probabilistic model. Their key idea is to treat registration as a maximum likelihood estimation, which can be solved using the EM algorithm. With a simple augmentation, they formulate the E step as a filtering problem and solve it using advances in efficient Gaussian filters.

ICP is prone to errors due to outliers and missing data. Thus, a variety of heuristics, as well as more principled methods, were introduced to deal with it. Chetverikov *et al.* [12] proposed a robust version of ICP, termed Trimmed ICP, which is based on Least Trimmed Squares that is designed to robustify the minimization of the error function. Bouaziz *et al.* [13] used sparse inducing norms to cope with missing points and outliers.

Rusinkiewicz [14] recently introduced a symmetric objective function for ICP that approximated a locally-second-order surface centered around the corresponding points. The proposed objective function achieved a larger basin of convergence, compared to regular ICP, while providing state-of-the-art accuracy.

Fitzgibbon [15] replaces ICP with a general-purpose nonlinear optimization (the Levenberg-Marquardt algorithm) that minimizes the registration error directly. His surprising finding is that his technique is comparable in speed to the special-purpose ICP algorithm.

Another line of research gives the correspondence problem a probabilistic interpretation. Instead of assuming a one-to-one correspondence, assignments are assumed to be probabilistic. Similar to us, these methods, described next, use gradient descent to find the optimal registration between two point clouds.

The differentiable approximation we take resembles that taken in SoftAssign [16]. There, they solve the correspondence problem, as an intermediate step, using a permutation matrix. Because that matrix is non-differentiable it is replaced with a Doubly-Stochastic Matrix.

EM-ICP [17] treats point matches as hidden variables and suggests a method that corresponds to an ICP with multiple matches weighted by normalized Gaussian weights, giving birth to the EM-ICP acronym of the method.

In KReg [18], the authors take an information theoretic approach to the problem. First, they define a kernel correlation that measures affinity between every pair of points. Then, they use that to measure the compactness of a point set and then show that registering two point sets minimizes the overall compactness. In addition, they show that this is equivalent to minimizing Renyi’s Quadratic Entropy. In fact, the only difference between the gradients of KReg [18] and EM-ICP [17] is the normalization term.

Deep methods. The introduction of PointNet [19] for processing unordered point clouds led to the development of PointNet-based registration algorithms.

PointNetLK [20] maps the two point clouds to some latent space in which it applies the Lucas-Kanade registration [21]. To do that, they define a supervised learning problem that takes two rotated versions of the same point cloud and produces the rotation between the two. The method is implemented using a Recurrent Neural Network, and avoids the costly step of point correspondence. On the downside, it requires a training phase to learn the embedding space, unlike our work that requires no training at all.

Deep Closest Point [22] consists of three parts: a point cloud embedding network, an attention-based module combined with a pointer generation layer [23] to approximate combinatorial matching, and a differentiable singular value decomposition (SVD) layer to extract the final rigid transformation. PointGMM [24] represents the data via a hierarchical Gaussian mixture and learns to perform registration by transforming shapes to their canonical position. DeepVCP [25], for *Virtual Corresponding Points*, trains a network to detect keypoints, match them probabilistically, and recover the registration using them.

A major drawback of deep learning based registration methods is that they strongly depend on the data that they have been trained on. A registration network that is trained for a given dataset does not necessarily generalize well to other datasets [26]. As we do not have a training step, our approach does not suffer from this problem.

Our approach builds on the work of Oron *et al.* [2] and that of Plötz and Roth [27]. Oron *et al.* introduced the concept of the best-buddies similarity measure as a robust method for template matching in images. The idea was to map image patches to points in some high dimensional space and count the number of mutual nearest neighbor matches between the two point sets. This was shown to converge to the χ^2 error measure when the number of points tends to infinity.

Plötz and Roth [27] proposed an approximation scheme to the nearest neighbor problem. Instead of selecting a particular element to be the nearest neighbor to a query point, they use a soft approximation that is governed by a temperature parameter. When the temperature goes to zero, the approximation converges to the deterministic nearest neighbor. Similarly to ICP and its variants, the best-buddies similarity relies on nearest neighbor search, and we use this nearest neighbor approximation in our work.

3 Method

We consider two point clouds $P = \{p_i\}_{i=1}^n$, $Q = \{q_i\}_{i=1}^m$, where $p_i, q_j \in \mathbb{R}^3$. We wish to find the transformation that aligns them, and in this work we assume this is a rigid transformation with 6 degrees of freedom (6DOF). We define several differentiable loss functions inspired by the *Best Buddy Similarity measure (BBS)*. We collectively name our registration algorithms *Best Buddy Registration (BBR)*. For each loss function L , the algorithm *BBR-L* works by optimizing over this loss function to find the aligning transformation:

$$\arg \min_{R,t,\alpha} \mathcal{L}(P, RQ + t), \quad (1)$$

where R is a 3D rotation matrix, t a 3D translation vector, and α a temperature parameter (discussed ahead). We parameterize the rotation using Euler angles: $R = R(\theta, \phi, \psi)$. We next describe the four variants of our algorithm: *BBR-softBBS*, *BBR-softBD*, *BBR-N* and *BBR-F*.

We start by defining the BBS measure: Let $D \in \mathbb{R}^{n \times m}$ denote the distance matrix between points in P and points in Q . A best buddies matrix B determines if a pair of points p_i and q_j are mutual nearest neighbors:

$$B_{ij} = \llbracket i = \arg \min_{i'} D_{i'j} \rrbracket \cdot \llbracket j = \arg \min_{j'} D_{ij'} \rrbracket, \quad (2)$$

where $\llbracket \cdot \rrbracket$ equals 1 if the term in the brackets is true and zero otherwise.

The Best-Buddies Similarity (BBS) loss $\mathcal{L}_{BBS}(P, Q)$ is negative the number of best buddies pairs³:

$$\mathcal{L}_{BBS}(P, Q) = - \sum_{i,j} B_{ij}. \quad (3)$$

The best-buddies similarity measure was shown to be very robust to outliers and missing data [2], in the context of template matching for images. We bring it to 3D point clouds. \mathcal{L}_{BBS} is a robust measure for the quality of the matching between two point clouds P and Q . However, it cannot be used for gradient-descent optimization, because it uses the non-differentiable argmin operator. To overcome this, we use the soft argmin approximation introduced by [27] for Neural Nearest Neighbors Networks. Specifically, \bar{B} approximates B as follows:

$$\bar{B}_{ij} = \frac{e^{-D_{ij}/\alpha}}{\epsilon + \sum_j e^{-D_{ij}/\alpha}} \cdot \frac{e^{-D_{ij}/\alpha}}{\epsilon + \sum_i e^{-D_{ij}/\alpha}}, \quad (4)$$

where α is a temperature parameter (see ahead), and ϵ is a small constant used for numerical stability. The matrix \bar{B} is the element-wise multiplication of row-wise and column-wise soft argmin of the distance matrix D . Observe how it corresponds to the brackets in Equation (2). The *softBBS* loss is now given by:

$$\mathcal{L}_{softBBS}(P, Q) = - \sum_i \sum_j \bar{B}_{i,j}. \quad (5)$$

³ in the original definition [2], BBS is normalized by nm . We omit that here.

$\mathcal{L}_{softBBS}$ is not only a differentiable approximation to \mathcal{L}_{BBS} , but also a generalization. While B_{ij} is only non-zero if p_i and q_j are mutual nearest neighbors, \bar{B}_{ij} can also be non-zero when, for example, p_i is q_j 's 3rd nearest neighbor, while q_j is p_i 's 4th nearest neighbor. The value of the temperature parameter, α , controls this behaviour. The smaller it is, the more strict \bar{B}_{ij} becomes, meaning more similar to B_{ij} . α is also learned during the optimization (together with the other optimized parameters, R and t). However, we find it important to initialize it with a reasonable value. A bad choice of α can result in a flat loss, unsuitable for gradient descent. In all our experiments we set $\alpha_{init} = 1e-2$ as it generally provides a smooth approximation to \mathcal{L}_{BBS} , but with a good slope near the minimum. For numerical stability, we allow it to decrease down to $\alpha_{final} = 1e-8$.

The next loss we suggest is the *soft buddy distance* loss, or \mathcal{L}_{softBD} , which makes use of the distance between the two points in each pair. This is in contrast to the BBS measure, which only counts the *number* of best-buddies, and softBBS loss, which is a soft approximation to that. We define *softBD* as follows:

$$\mathcal{L}_{softBD}(P, Q) = \frac{\sum_i \sum_j \bar{B}_{i,j} D_{ij}}{\sum_i \sum_j \bar{B}_{i,j}}. \quad (6)$$

The next loss is *softBD with normals* loss, or \mathcal{L}_N , which uses a point-to-plane distance measure calculated from local normals. Such a distance measure is used in some of the most popular and successful ICP variants, such as generalized ICP [8], and symmetric ICP [14]. In the previous methods we suggested, we used the Euclidean point-to-point distance to create the distance matrix D . In *BBR-N* we replace that with the following symmetric point-to-plane distance, based on Rusinkiewicz [14]:

$$D_{i,j}^n = \text{dist}(Rq_i + t, p_j) = |\langle Rq_i + t - p_j, Rn_{q_i} + n_{p_j} \rangle|. \quad (7)$$

where n_{q_i} is the normal at point q_i , n_{p_j} is the normal at point p_j , and D^n denotes the version of the distance matrix calculated using this distance. The loss function is then calculated as in *softBD*, except using D^n instead D . This distance is symmetric in the sense that it uses the normals from both points, unlike algorithms such as point-to-plane ICP [5] that only use normals from one of the point clouds.⁴

The final loss we present is the *best-buddy filtering* loss, or \mathcal{L}_F . At the heart of the BBS measure lies the robustness achieved by using mutual nearest neighbors. *BBR-F* translates this idea into *best buddy filtering*: using only pairs that are mutual nearest neighbors. In addition, it follows the trend in ICP-like algorithms, in that it uses both point-to-point and point-to-plane distance measures: the Euclidean point-to-point distance is used for the stage of matching pairs between the two point clouds. Then the symmetric point-to-plane distance between these pairs is used to define the following loss:

$$\mathcal{L}_F(P, Q) = - \sum_i \sum_j B_{i,j} \cdot D_{i,j}^n. \quad (8)$$

⁴ This is not to be confused with the symmetric nature of the best-buddies similarity measure that we introduce here.

Notice that in this variant of BBR we have a hard selection of pairs, which isn't optimized during gradient descent. Instead, the *distances* between the points in each pair are minimized.

The central difference between *BBR-F* and symmetric ICP [14] is that best buddy filtering replaces explicit outlier rejection. This removes the necessity to calibrate outlier-rejection parameters, while resulting in better accuracy in settings where robustness is important, as Section 4 shows. Another difference is that *BBR-F* uses Adam gradient descent for optimization, while symmetric ICP uses a closed form solution to an approximate linearized version of the symmetric point-to-plane distance measure.

BBR-F is especially useful for very large point clouds, where memory and running time become a constraint, because it does not require the full distance matrix D or D^n . For the pair matching step, we use the KD-tree method [28], and then calculate the point-to-plane distances only for best-buddy pairs. See supplementary material for run time measurements.

4 Experiments

We present experiments that are designed to analyze the behaviour of the BBR methods, and evaluate them on several datasets including Stanford [29], TUM RGBD [30], KITTI Odometry [31], and Apollo-Southbay [32]. We compare our approach to several established alternatives, focusing on classic approaches (e.g., ICP) as opposed to learned methods, as the latter do not necessarily generalize well across datasets [26].

Figure 1 shows the best-buddies pairs during a typical run of *BBR-softBBS* on the Stanford Bunny model [29]. The algorithm converges after 120 iterations. At first, there are few best-buddies pairs, but as time progresses their number grows until convergence.

4.1 Performance Evaluation Setup

We conduct a set of experiments to evaluate accuracy and robustness. To do that, we apply different rotations, translations, sub-sampling, and noise, to each point cloud. We compare ourselves to the following popular point cloud registration algorithms: (i) HGMR [10], a GMM based method; (ii) Coherent Point Drift (CPD) [33], a probabilistic algorithm based on GMM; (iii) Generalized ICP (G-ICP) [8], a very popular and accurate ICP variant that uses local normals, and (iv) Symmetric ICP (Sym-ICP) [14], which provides state-of-the-art performance on several point cloud registration challenges. Additional experiments with different ranges of noise and motions appear in the supplementary material, as do the definitions we use for translational and angular error.

ICP algorithms are sensitive to noise, and therefore commonly employ a set of standard practices for outlier rejection. Our BBR methods require no such processes.

Setting. We test the 4 variants of our algorithm: *BBR-softBBS*, *BBR-softBD*, *BBR-N* and *BBR-F*. *BBR-softBBS* tends to be less accurate than the others, and therefore we omit it from most experiments.

The local normals that are used in *BBR-N* and *BBR-F* are estimated by calculating the principal axis of a neighborhood of $k = 13$ neighbors around each point in the full cloud (before subsampling). For consistency, Sym-ICP was given the same normals used for *BBR-N*. G-ICP calculates its own normals on-the-fly, from the subsampled point cloud that it takes as input. For all BBR methods, the optimization is performed by running Adam for a pre-defined number of iterations.

We use the Probreg⁵ library’s implementation of *HGMR* and *CPD*, and the original *SymICP* implementation, for all of which we use the default parameters. We use the Point Cloud Library’s [34] implementation of *G-ICP*, setting the parameters as in [25].

4.2 Comparing Accuracy and Robustness between BBR variants

We start by comparing the different variants of BBR in a simple experiment, testing their accuracy and ability to converge with different initial rotations: 5, 10, 30, 60 and 90 degrees. For each angle, we repeat 20 times: select two random subsets of 500 points from the Stanford Bunny point cloud, rotate the target point cloud around a random axis, perform registration and measure the angular error. We consider registration to have failed if the final error is over 5° . Results are shown in figure 2. *BBR-softBBS* is the clear leader in robustness to large initial error. Unlike the others, it is able to handle initial rotations of 90° with hardly any failures. Notice that *BBR-N* is especially sensitive to large initial rotations, due to its reliance on normals to recognize best-buddies. When the initial rotation is large, the same object will have very different normals in each of the two scans. For all algorithms, large initial rotations do not cause degradation in accuracy - for the attempts that did succeed. When looking at the accuracy of the successful attempts, it’s clear that the methods that make use of local normals (*BBR-N* and *BBR-F*) are significantly more accurate.

4.3 Accuracy

The experiments shown in Figure 3 demonstrate our ability to register point clouds with random rotation and translation, using the Bunny, Horse, and Dragon point clouds from [29]. We randomly select a source and target subset from the original point cloud, each containing M points. The target point cloud is rotated around a randomly selected axis by an angle of θ_{rot} . It is then translated along a random axis to a distance of Δ_{trans} . We then run the registration algorithms on the point-cloud pair and record their translational and angular error. We repeat the experiment T times and report the median error for each algorithm and each cardinality M of the point clouds.

⁵ <https://github.com/neka-nat/probreg>

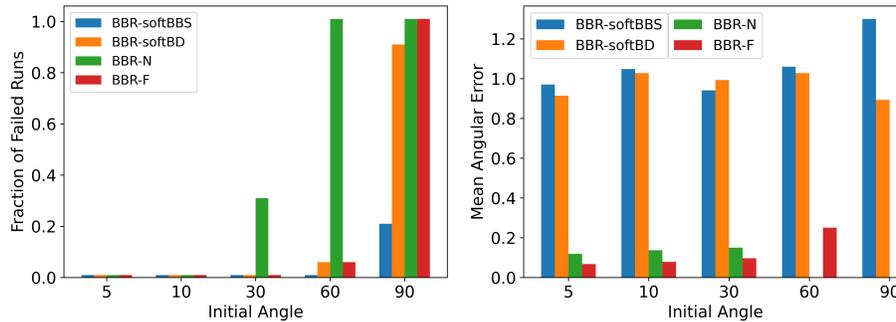


Fig. 2: **Accuracy and Robustness to Initial Error** Comparison among the BBR methods using the Stanford Bunny point cloud. Left - the fraction out of 20 registration attempts that failed, Right - accuracy as measured by the mean angular error of the successful attempts. *BBR-softBBS* is the best at converging from a large initial error. *BBR-N* and *BBR-F*, that make use of local normals, are the most accurate.

For all experiments in Figure 3 we set $\Delta_{trans} = 0.005m$ and $T = 20$. The other parameters are: Bunny: $M \in \{200, 300, \dots, 1000\}$, $\theta_{rot} = 8^\circ$. Horse: $150 \leq M \leq 1000$, $\theta_{rot} = 10^\circ$. Dragon: M up to 2000, $\theta_{rot} = 5^\circ$. *BBR-F* achieves the lowest error rate, across almost all point cloud sizes. It is followed closely by *BBR-N*. Only when the point density is high, Sym-ICP performs on-par with *BBR-N*. This demonstrate BBR’s ability to work well with very sparse point clouds. It should also be pointed out that *BBR-softBD* outperforms the comparable registration method that do not use normals (as well as G-ICP).

4.4 Robustness

The next experiments demonstrate the resistance of the algorithm to a variety of challenges, including occlusions, the presence of a distractor, measurement noise, and a large initial error.

Partial Overlap and Occlusion. The experiment shown in Figure 4 evaluates the resistance to partial overlap and occlusion. We perform registration between two partial scans of the Stanford Bunny, *bun000* and *bun090*, each captured from a different view point. Following Rusinkiewicz’s [14] experimental setup, we first align the scans according to the ground truth motion. Then we follow the same experimental method as in Section 4.3, with $M \in \{200, 300, \dots, 1000\}$, $\theta_{rot} = 5^\circ$ and $\Delta_{trans} = 0.005m$. *BBR-F* achieves the most accurate results, followed by *BBR-N* and *Sym-ICP*. However, *Sym-ICP* deteriorates considerably when given a very sparse point cloud, while both of our algorithms cope with it very well.

Resistance to Distractors. This experiment evaluates the effects of distractor noise. This is the case where in addition to the main object of interest, the

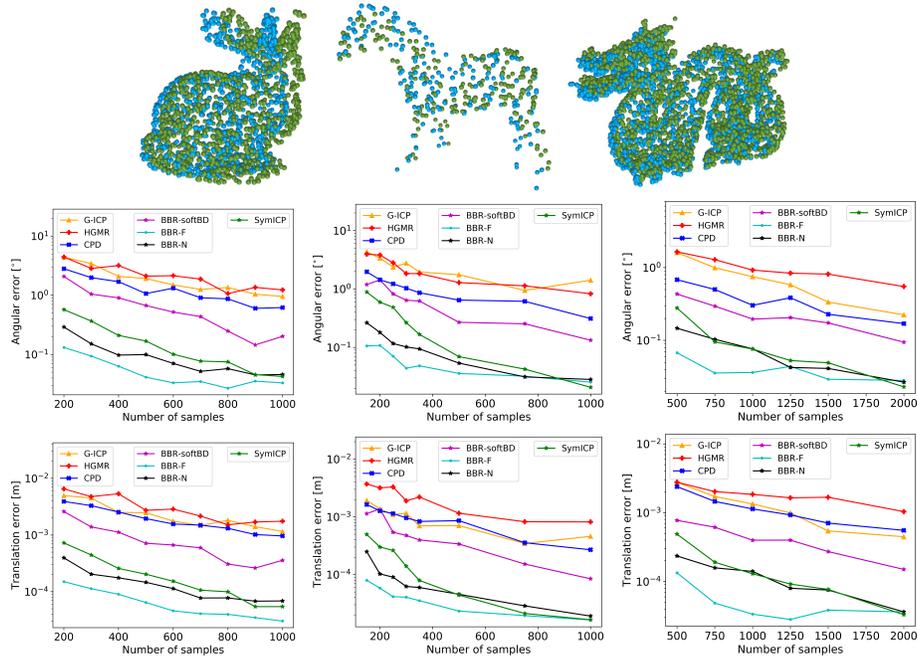


Fig. 3: **Accuracy test:** Top Row - Point clouds examples as used in the accuracy test. Middle and bottom rows - angular and translation error as a function of the number of points. *BBR-F* achieves the most accurate results in all scenarios, followed closely by *BBR-N*. *SymICP* becomes competitive only when the number of points increases.

scene also contains a second object with a different motion. In this synthetic experiment, the main object, a large horse, was randomly translated and rotated as in Section 4.3, with $T = 20$, $\theta_{rot} = 10^\circ$ and $\Delta_{trans} = 0.005m$. The distractor object, a small horse, underwent a different motion. Such a situation may occur when attempting to estimate the ego-motion of a vehicle, using scans that include other independently moving vehicles. The main object contains 1000 points, and we vary the number of points in the distractor object from 200 up to 900. In figure 5 we show the median error as a function of the number of points in the distractor. *BBR-F* shows a strong resistance to distractor noise in this experiment, while *Sym-ICP* is quite susceptible to it. Among methods that do not make use of normals, *BBR-sofBD* is the most accurate.

Measurement noise. The TUM RGB-D dataset [30] contains point clouds of indoor scenes captured with the Kinect sensor. It contains natural measurement noise, due to the warp and scanning noise from the Kinect sensor. It has been noted by Rusinkiewicz [14] that this dataset poses a qualitatively different challenge than the bunny point cloud. He demonstrates his algorithm on a spe-

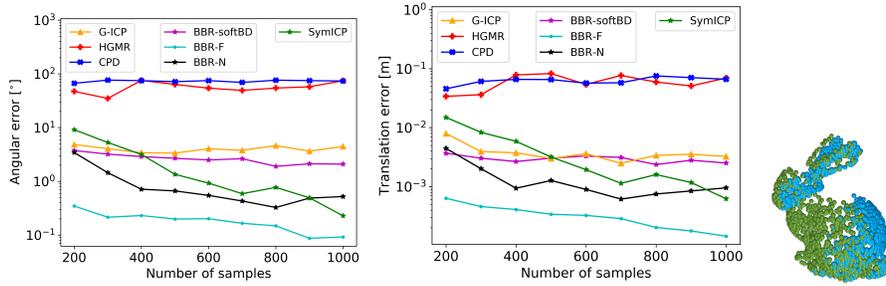


Fig. 4: **Robustness to point of view changes:** Right - aligned *bun000* and *bun090* with 1000 points, with overlap of less than 30%. Left - angular and translation error for varying number of points. *BBR-sofBD*, *BBR-N* and *BBR-F* all outperform Sym-ICP [14] when the number of samples is low.

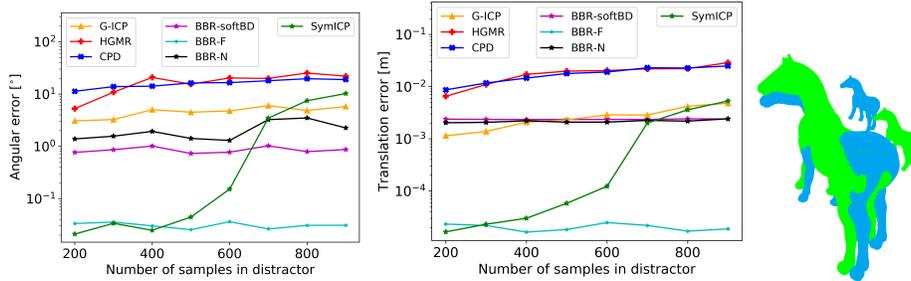


Fig. 5: **Distractor noise:** Right - an example of our experimental setting: The main object (large horse) has moved differently than the distractor (small horse). Left and center - the angular and translation errors for varying distractor strengths. The BBR algorithms are able to ignore the distractor and focus on the motion of the main object.

cific pair of partially-overlapping scans from this dataset. We use the same pair, 1305031104.030279 and 1305031108.503548 of the freiburg1_xyz sequence from the TUM RGB-D (See image in supplementary material), sample 1000 points from each, and experiment with adding a random rotation of up to 5° around a random axis. We repeat this 50 times and perform registration for each, showing the cumulative distribution of the final errors in Figure 6. The BBR algorithms perform better than all competing methods.

Large initial error with partial overlap. In this experiment we evaluate the ability to converge when the source and target point clouds are only partially overlapping, and starting from a large initial error. We use two scans of the bunny point cloud, *bun180* and *bun270*, that were captured from significantly different view points. As before, we first align them, sample 1000 points from each, and then apply a random motion to one of them. In this experiment the motion is

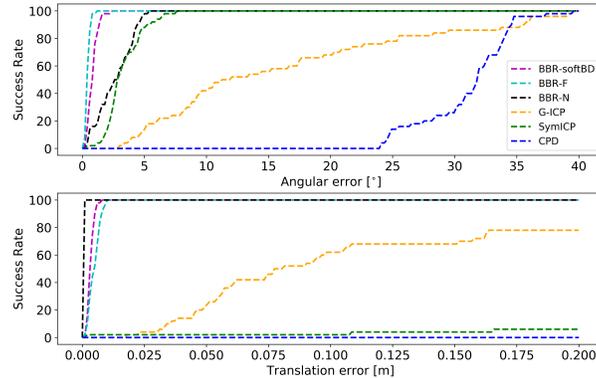


Fig. 6: **Convergence Analysis (TUM)**: The cumulative distribution of errors over 50 repeats. x -axis: the error threshold; y axis: the fraction of results that achieved an error smaller than this threshold. For example, G -ICP has 20% of its registration results with an angular error lower than 5° , $SymICP$ has 90% of its results under 5° , and the BBR algorithms have 100% of their errors under 5° .

large: a random rotation in the range $[30, 50]$ degrees, around a random axis, and a translation of half the extent of the point cloud ($0.05m$) along a random direction. Figure 7 shows the cumulative distribution of errors over 50 repeats. BBR -softBD performs considerably better than all other algorithms.

4.5 Odometry

In this section we present experiments in the realistic setting of vehicle navigation, specifically focusing on a difficult setting where the separation between the two clouds is relatively large, leading to significant occlusions and outliers caused by independently moving objects in the scene. We use two datasets of high-resolution Lidar scans: KITTI Odometry [31], and Apollo-Southbay [32], both of which consist of large point clouds of over 100K points. We follow the experimental setup of [25], where the test set consists of pairs of clouds scanned by the same vehicle at two times, during which the vehicle has travelled up to 5 meters, for up to 2 seconds (KITTI) or even 5 seconds (Apollo-Southbay). An initial estimate for the motion is assumed to be available, which is inaccurate by up to 1 meter in each of x, y, z , and up to 1 degree in each of θ, ϕ, ψ . We report the mean translation and angular errors, as well as the maximum (worst case).

For this test, we use the BBR -F variant of our BBR algorithm. Normals are estimated from neighborhoods of $k = 94$ points in the full cloud. To achieve high accuracy, we only moderately subsample the point clouds to 30K points, uniformly at random.

Tables 1 and 2 compare our results to those reported in [25], and to those of Sym-ICP [14]. In both experiments, BBR -F achieves state of the art accuracy

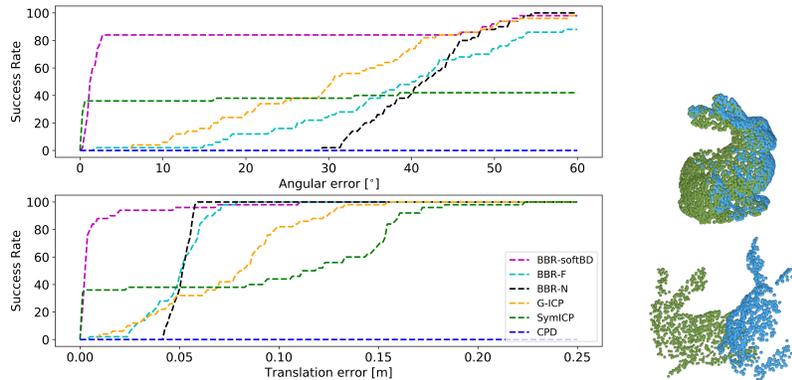


Fig.7: **Convergence Analysis (partially-overlapping Bunny):** Right: *bun180* and *bun270*, aligned (top), and in their initial position: rotated in 50° and translated by a factor of $0.05m$ (bottom). Left: the cumulative distribution of errors over 50 repeats. The x -axis is the error threshold and the y axis is fraction of registration results whose error is smaller than this threshold.

by three of the four accuracy measures. This demonstrates BBR’s capabilities of achieving high accuracy in realistic scenarios that include occlusions, outlier motions and measurement noise.

5 Conclusions

We proposed *Best Buddy Registration (BBR)* algorithms for point cloud registration inspired by the Best Buddy Similarity (BBS) measure. First we show that registration can be performed by running gradient descent on a differential approximation to the negated BBS measure. This results in an algorithm that is quite robust to noise, occlusions and distractors, and able to cope with very sparse point clouds. We then present additional algorithms that achieve higher accuracy while maintaining robustness, by incorporating point-to-point and point-to-plane distances into the loss function. Finally, we present the *BBR-F* algorithm that uses *best buddy filtering* to achieves state of the art accuracy on challenges that include significant noise, occlusions and distractors, including registration of automotive lidar scans that are relatively widely separated in time. Our algorithms are implemented in Pytorch and optimized with Adam gradient descent, allowing them to be incorporated as a registration stage in Deep Neural Networks for processing point clouds.

6 Acknowledgements

This research was supported by ERC-StG grant no. 757497 (SPADE) and by ISF grant number 1549/19.

Method	Angular Error($^{\circ}$)		Translation Error(m)	
	Mean	Max	Mean	Max
ICP-Po2Po [1]	0.139	1.176	0.089	2.017
ICP-Po2P1 [1]	0.084	1.693	0.065	2.050
G-ICP [8]	0.067	0.375	0.065	2.045
AA-ICP [35]	0.145	1.406	0.088	2.020
NDT-P2D [36]	0.101	4.369	0.071	2.000
CPD [33]	0.461	5.076	0.804	7.301
3DFeat-Net [37]	0.199	2.428	0.116	4.972
DeepVCP-Base [25]	0.195	1.700	0.073	0.482
DeepVCP-Duplication [25]	0.164	1.212	0.071	0.482
Sym-ICP [14]	0.066	0.422	0.058	0.863
BBR-F (ours)	0.065	0.356	0.058	0.730

Table 1: The KITTI Odometry dataset. Our algorithm achieves state of the art results for almost all accuracy measures.

Method	Angular Error($^{\circ}$)		Translation Error(m)	
	Mean	Max	Mean	Max
ICP-Po2Po [1]	0.051	0.678	0.089	3.298
ICP-Po2P1 [1]	0.026	0.543	0.024	4.448
G-ICP [8]	0.025	0.562	0.014	1.540
AA-ICP [35]	0.054	1.087	0.109	5.243
NDT-P2D [36]	0.045	1.762	0.045	1.778
CPD [33]	0.054	1.177	0.210	5.578
3DFeat-Net [37]	0.076	1.180	0.061	6.492
DeepVCP-Base [25]	0.135	1.882	0.024	0.875
DeepVCP-Duplication [25]	0.056	0.875	0.018	0.932
Sym-ICP [14]	0.018	2.589	0.010	6.625
BBR-F (ours)	0.015	0.308	0.007	2.517

Table 2: The Apollo-SouthBay odometry dataset. Our algorithm achieves state of the art results for almost all accuracy measures.

References

1. Besl, P.J., McKay, N.D.: A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell.* **14** (1992) 239–256
2. Oron, S., Dekel, T., Xue, T., Freeman, W.T., Avidan, S.: Best-buddies similarity—robust template matching using mutual nearest neighbors. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **40** (2018) 1799–1813
3. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. *International Conference on Learning Representations* (2014)
4. Ding, L., Feng, C.: Deepmapping: Unsupervised map estimation from multiple point clouds. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (2019)
5. Chen, Y., Medioni, G.: Object modelling by registration of multiple range images. *Image Vision Comput.* **10** (1992) 145–155
6. Rusinkiewicz, S., Levoy, M.: Efficient variants of the icp algorithm. In: *Proceedings Third International Conference on 3-D Digital Imaging and Modeling*. (2001) 145–152
7. Pomerleau, F., Colas, F., Siegwart, R.: *A Review of Point Cloud Registration Algorithms for Mobile Robotics*. (2015)
8. Segal, A., Hähnel, D., Thrun, S.: Generalized-icp. In *Trinkle, J., Matsuoka, Y., Castellanos, J.A., eds.: Robotics: Science and Systems, The MIT Press* (2009)
9. Jian, B., Vemuri, B.C.: Robust point set registration using gaussian mixture models. *IEEE Trans. Pattern Anal. Mach. Intell.* **33** (2011) 1633–1645
10. Eckart, B., Kim, K., Kautz, J.: Hgmr: Hierarchical gaussian mixtures for adaptive 3d registration. In: *The European Conference on Computer Vision (ECCV)*. (2018)
11. Gao, W., Tedrake, R.: Filterreg: Robust and efficient probabilistic point-set registration using gaussian filter and twist parameterization. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR, Computer Vision Foundation / IEEE* (2019) 11095–11104
12. Chetverikov, D., Stepanov, D., Krsek, P.: Robust euclidean alignment of 3d point sets: the trimmed iterative closest point algorithm. *Image and Vision Computing* **23** (2005) 299 – 309
13. Bouaziz, S., Tagliasacchi, A., Pauly, M.: Sparse iterative closest point. In: *Proceedings of the Eleventh Eurographics/ACMSIGGRAPH Symposium on Geometry Processing. SGP '13* (2013) 113–123
14. Rusinkiewicz, S.: A symmetric objective function for ICP. *ACM Transactions on Graphics (Proc. SIGGRAPH)* **38** (2019)
15. Fitzgibbon, A.W.: Robust registration of 2D and 3D point sets. In: *British Machine Vision Conference*. (2001) 662–670
16. Rangarajan, A., Chui, H., Bookstein, F.L.: The softassign procrustes matching algorithm. In: *Biennial International Conference on Information Processing in Medical Imaging*. (1997) 29–42
17. Granger, S., Pennec, X.: Multi-scale em-icp: A fast and robust approach for surface registration. In: *European Conference on Computer Vision, Springer* (2002) 418–432
18. Tsin, Y., Kanade, T.: A correlation-based approach to robust point set registration. In: *European conference on computer vision, Springer* (2004) 558–569
19. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017) 652–660

20. Aoki, Y., Goforth, H., Srivatsan, R.A., Lucey, S.: Pointnetlk: Robust & efficient point cloud registration using pointnet. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2019)
21. Lucas, B.D., Kanade, T.: An iterative image registration technique with an application to stereo vision. In: Proceedings of the 7th International Joint Conference on Artificial Intelligence - Volume 2. IJCAI'81 (1981) 674–679
22. Wang, Y., Solomon, J.M.: Deep closest point: Learning representations for point cloud registration. In: The IEEE International Conference on Computer Vision (ICCV). (2019)
23. Vinyals, O., Fortunato, M., Jaitly, N.: Pointer networks. In Cortes, C., Lawrence, N.D., Lee, D.D., Sugiyama, M., Garnett, R., eds.: Advances in Neural Information Processing Systems 28. (2015) 2692–2700
24. Hertz, A., Hanocka, R., Giryes, R., Cohen-Or, D.: Pointgmm: A neural gmm network for point clouds. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). (2020) 12051–12060
25. Lu, W., Wan, G., Zhou, Y., Fu, X., Yuan, P., Song, S.: Deepvcp: An end-to-end deep neural network for point cloud registration. In: The IEEE International Conference on Computer Vision (ICCV). (2019)
26. Sarode, V., Li, X., Goforth, H., Aoki, Y., Srivatsan, R.A., Lucey, S., Choset, H.: Pernet: Point cloud registration network using pointnet encoding. *ArXiv abs/1908.07906* (2019)
27. Plötz, T., Roth, S.: Neural Nearest Neighbors Networks. Proceedings of Advances in Neural Information Processing Systems (NeurIPS) (2018)
28. Bentley, J.L.: Multidimensional binary search trees used for associative searching. *Commun. ACM* **18** (1975) 509–517
29. Turk, G., Levoy, M.: Zipped polygon meshes from range images. In: Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques. SIGGRAPH '94, ACM (1994) 311–318
30. Sturm, J., Engelhard, N., Endres, F., Burgard, W., Cremers, D.: A benchmark for the evaluation of RGB-D SLAM systems. In: IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, Vilamoura, Algarve, Portugal, October 7-12, 2012. (2012) 573–580
31. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: Conference on Computer Vision and Pattern Recognition (CVPR). (2012)
32. Lu, W., Zhou, Y., Wan, G., Hou, S., Song, S.: L3-Net: Towards learning based LiDAR localization for autonomous driving. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE (2019)
33. Myronenko, A., Song, X.: Point set registration: Coherent point drift. *IEEE Trans. Pattern Anal. Mach. Intell.* **32** (2010) 2262–2275
34. Rusu, R.B., Cousins, S.: 3d is here: Point cloud library (pcl). In: IEEE International Conference on Robotics and Automation (ICRA), Shanghai, China (2011)
35. Pavlov, A.L., Ovchinnikov, G.W., Derbyshev, D.Y., Tsetserukou, D., Oseledets, I.V.: Aa-icp: Iterative closest point with anderson acceleration. In: 2018 IEEE International Conference on Robotics and Automation (ICRA). (2018) 3407–3412
36. Stoyanov, T., Magnusson, M., Andreasson, H., Lilienthal, A.J.: Fast and accurate scan registration through minimization of the distance between compact 3d ndt representations. *I. J. Robotics Res.* **31** (2012) 1377–1393
37. Yew, Z.J., Lee, G.H.: 3DFeat-Net: Weakly supervised local 3d features for point cloud registration. In: European Conference on Computer Vision, Springer (2018) 630–646