

Accelerated Gravitational Point Set Alignment with Altered Physical Laws*

Vladislav Golyanik¹
¹MPI for Informatics

Christian Theobalt¹
²University of Kaiserslautern

Didier Stricker^{2,3}
³DFKI

Abstract

This work describes Barnes-Hut Rigid Gravitational Approach (BH-RGA) — a new rigid point set registration method relying on principles of particle dynamics. Interpreting the inputs as two interacting particle swarms, we directly minimise the gravitational potential energy of the system using non-linear least squares. Compared to solutions obtained by solving systems of second-order ordinary differential equations, our approach is more robust and less dependent on the parameter choice. We accelerate otherwise exhaustive particle interactions with a Barnes-Hut tree and efficiently handle massive point sets in quasilinear time while preserving the globally multiply-linked character of interactions. Among the advantages of BH-RGA is the possibility to define boundary conditions or additional alignment cues through varying point masses. Systematic experiments demonstrate that BH-RGA surpasses performances of baseline methods in terms of the convergence basin and accuracy when handling incomplete, noisy and perturbed data. The proposed approach also positively compares to the competing method for the alignment with prior matches.

1. Introduction

Alignment of point sets provided in different reference frames is an essential algorithmic component in multiple fields including but not limited to visual odometry [33], 3D reconstruction and augmented reality [40], [39], robot navigation and localisation [46], computer graphics [48], CAD modelling [52], cultural heritage [8] and medical technology [11, 25]. Even though rigid point set registration is a well-studied area [6, 11, 43, 20, 12, 18, 25, 50, 41, 38, 14, 21, 15], multiple challenges are remaining — on the top of the list are the broader basin of convergence, the robustness to disturbing effects in the samples (noise, missing data, clustered outliers) and handling of massive point sets. We are considering rigid alignment (RA) relative to a fixed reference frame \mathcal{R} (perhaps an observer’s reference frame), and coordinates of one of the point sets \mathcal{X} are known in \mathcal{R} .

* supported by the ERC Consolidator Grant 4DReply (770784) and the BMBF projects DYNAMICS (01IW15003) and VIDETE (01IW18002).




Figure 1: An overview of our BH-RGA approach. In every external iteration, we compute a Barnes-Hut tree for the reference and template jointly and specify an optimisation problem. Next, an NLLS solver runs for a few internal iterations while requiring $\mathcal{O}(M \log N)$ time to approximate gravitational potential energy of globally multiply-linked particle interactions.

We refer to \mathcal{X} as a *reference*. The objective of RA is the recovery of a non-injective, non-surjective correspondence function $f: \mathcal{Y} \rightarrow \mathcal{X}$ together with the rigid transformation parameters, *i.e.*, rotation \mathbf{R} and translation \mathbf{t} (6 DoF) registering the *template* point set \mathcal{Y} to \mathcal{X} .

All RA techniques can be divided into local and multiply-linked policies. An RA method is called *globally multiply-linked* if and only if every point of \mathcal{X} interacts with every point of \mathcal{Y} , with an example of the recently proposed particle dynamics based gravitational approach (GA) [21]. In GA, every point of \mathcal{Y} interacts with every point of \mathcal{X} through a force field. The method was shown to be more robust to large amounts of noise. This robustness comes at the cost of quadratic complexity and parameter sensitivity which impedes its applicability to real-world large and dense point sets. Moreover, the solution based on physical simulation is slow and unstable. We believe that considerable improvements are possible in the field of physics-based RA. In practical applications, we are interested in fast convergence and well-posedness with respect to the parameter settings, and less in the physically accurate modelling. On the other hand, feasible acceleration techniques for particle dynamics based RA remained unexplored so far and are worth to investigate.

Contributions. In the proposed exposition, we present a stable, robust and accelerated globally multiply-linked particle dynamics based RA approach and overcome the limitations of previous works, see Fig. 1 for an overview. Recall that behind the scenes, GA finds transformations corre-

sponding to the locally minimal gravitational potential energy (GPE). Along with that, the physically accurate GPE is not the optimal form to minimise by general-purpose iterative optimisation techniques (*e.g.*, Gauss-Newton method).

Thus, we **first** apply a negative elementwise reciprocal transform (Sec. 4) and enable efficient solution in the framework of non-linear least squares (NLLS). The proposed transform changes the form of the energy functional to a weighted sum of squared residuals and inverts physics while preserving the advantages of multiply-linked particle dynamics (*e.g.*, notion of the particle mass). Thanks to the new form of GPE, optimisation requires fewer iterations and can be performed with a general-purpose Gauss-Newton solver, thus eliminating the need for the explicit updates of forces, acceleration and displacements. Moreover, the number of parameters is reduced by the factor of three while showing a wider convergence basin and outperforming the baseline GA [21] in the accuracy.

Second, we adopt an acceleration technique based on the idea of Barnes and Hut [5] initially developed for N -body simulations [1] (Sec. 4.2). Applied to RA, a Barnes-Hut (BH) tree allows to efficiently handle large point sets with hundreds of thousands of points while preserving the globally multiply-linked nature of interactions. In other words, this happens not at the expense of constraining interactions to local vicinities but rather by accumulating contributions of reference points into clusters. The BH tree encompasses multiple extractable representations of \mathcal{X} with spatially varying cluster configurations.

We call the new method *Barnes-Hut Rigid Gravitational Approach* (BH-RGA). **Finally**, we design a set of systematic tests for RA with heightened complexity and variety. We compare BH-RGA to multiple RA methods [6, 18, 38, 32, 21, 23] on our extensive evaluation benchmark and apply it to real data originating from cultural heritage and a lidar sensor (Sec. 5).

2. Related Work

Early point set registration algorithms were motivated by emerging 3D scanners producing partial point clouds that need to be aligned. The seminal *iterative closest point* (ICP) algorithm for aligning two point clouds [6, 11] alternates between transformation estimation [28, 29] and local correspondence inference [17]. Even though ICP is easy to implement, its fundamentally heuristic local correspondence search makes it prone to erroneous local convergence if badly initialised, and sensitive to outliers. Different improvements were subsequently proposed for ICP, ranging from accelerating policies for nearest neighbour search [26, 41] and relaxation of one-to-one correspondences [20] to more efficient optimisation schemes [45, 18].

Another class of methods models source and target point clouds as probability density functions [12, 38, 32], such

as Gaussian Mixture Models (GMMs). Coherent Point Drift (CPD) [38] explicitly incorporates a model of uniform noise. A wider convergence basin for alignment was obtained by Gaussian mixture decoupling [16] or combining GMM representation and continuous domain mapping with a support vector machine [10]. The multiply-linked Kernel Correlation (KC) approach minimises the Renyi’s quadratic entropy of the joint system of the reference and the transformed template [50]. In contrast to BH-RGA, only local neighbourhoods are involved in one-to-many interactions.

More recently, approaches relying on analogies to physical processes gained attention [14, 21, 3]. Deng *et al.* cast point sets into the Schrödinger distance transform representation, and align them by minimising a geodesic distance on a unit Hilbert sphere [14]. Golyanik *et al.* [21] interpret point sets as rigid swarms of particles with masses moving in a gravitational force field. An optimal alignment corresponds to the state of minimal GPE of the system. Quadratic complexity and parameter sensitivity impede applicability of GA in practice. We develop the idea of global multiply-linking further and accelerate all-to-all interactions with a BH tree resulting in quasilinear computational complexity. We are altering the law of particle dynamics and simulate an inverse world. In the inverse world, a state with locally optimal GPE can be found iteratively by NLLS more easily. Thus, we remove viscosity (an energy dissipator in GA), and the gravitational constant becomes redundant as a multiplicative factor in the energy.

Large point sets in RA are often treated by subsampling [25], and rarely raw, let alone in the globally multiply-linked fashion. Some methods naturally allow embedding of prior matches [6, 23]. Similarly to GA, BH-RGA uses varying masses to define different boundary conditions. Particle masses can be set proportional to the reliabilities of superimposed prior correspondences. Some methods additionally use colours as an alignment cue [13]. In BH-RGA, point intensities can be likewise mapped to masses.

3. Newtonian Gravitational Approach [21]

Consider globally multiply-linked gravitational alignment of two point sets, with the reference $[\mathbf{x}_j] = \mathbf{X} \in \mathbb{R}^{D \times N}$, $j \in \{1, \dots, N\}$ and the template $[\mathbf{y}_i] = \mathbf{Y} \in \mathbb{R}^{D \times M}$, $i \in \{1, \dots, M\}$. N and M denote the number of points in the reference and template, respectively, and D is point set dimensionality. The objective is to recover the parameter set $\vartheta = \{\mathbf{R}, \mathbf{t}\}$, *i.e.*, rotation \mathbf{R} ($\mathbf{R}^{-1} = \mathbf{R}^\top$, $\det(\mathbf{R}) = 1$) and translation \mathbf{t} aligning \mathbf{Y} to \mathbf{X} . In [21], point sets are aligned by minimising the mutual *gravitational potential energy* (GPE) \mathbf{U} of the corresponding system of particles in the force field induced by \mathbf{X} :

$$\mathbf{U}(\mathbf{R}, \mathbf{t}) = -G \sum_{i,j} \frac{m_{\mathbf{y}_i} m_{\mathbf{x}_j}}{\|\mathbf{R} \mathbf{y}_i + \mathbf{t} - \mathbf{x}_j\|_2 + \epsilon}, \quad (1)$$

where $m_{\mathbf{y}_i}$ and $m_{\mathbf{x}_j}$ denote masses, G is the gravitational constant and ϵ is a softening parameter for preventing local near-field singularities. For the reasons stressed in [24], we do not include scaling in (1) and our energy functionals (5) and (9). Golyanik *et al.* [21] minimise (1) implicitly by updating the forces $\vec{\mathbf{f}}_i$ acting on particles \mathbf{y}_i , accelerations (using Newton’s second law of motion $\vec{\mathbf{f}}_i = m_{\mathbf{y}_i} \vec{\mathbf{a}}_i$), velocities v_i^{t+1} and individual point displacements d_i^{t+1} [1]:

$$\vec{\mathbf{f}}_i = -Gm_{\mathbf{y}_i} \sum_j m_{\mathbf{x}_j} (\|\mathbf{y}_i - \mathbf{x}_j\|^2 + \epsilon^2)^{-3/2} \hat{\mathbf{n}}_{i,j} - \eta v_i^t, \quad (2)$$

$$v_i^{t+1} = v_i + \Delta t \frac{\vec{\mathbf{f}}_i}{m_{\mathbf{y}_i}} \quad \text{and} \quad d_i^{t+1} = \Delta t v_i^{t+1}. \quad (3)$$

In (2), η denotes a dissipation constant which determines the portion of the kinetic energy which is dragged from the system per template particle. In (3), Δt is the the forward integration step (time interval of the simulation). In every iteration, once updated, the unconstrained displacements are added to the current positions. A consensus rigid transformation is found using Procrustes analysis [29], which relates the previous and current poses. The algorithm converges when the difference in the GPE of two last successive system states is below some small threshold ρ_E or terminates when the maximum number of iterations is reached.

4. The Proposed Approach

4.1. Our Gravitational Potential Energy Functional

GPE (1) includes a reciprocal relation and is minimised implicitly by ordinary differential equations (ODE) of second order by updating forces acting on the template particles, point accelerations and displacements. The forces are inversely proportional to distances, and at the optimal alignment, the GPE is limited by $-\infty$. For the transformation parameters \mathbf{R} and \mathbf{t} to converge, additional regularisation constants are introduced in (2), *i.e.*, softening distance ϵ to avoid near-field singularities, and an energy dissipation constant η to, *de facto*, prevent infinite oscillations. The combination of all these properties (reciprocal relation, regularising energy dissipation term, ϵ) makes GA ill-posed with respect to the parameter values. Thus, every scenario demands different parameters — otherwise, the method will not converge. Another disadvantage of the implicit solution with second-order ODEs is that a large number of iterations is required until convergence (300 – 1000 iterations are not rare). Moreover, \mathbf{R} is updated from unconstrained point displacements regularised by Procrustes fitting. Ideally, we would like to optimise directly in the $SO(D)$ group.

We propose to directly optimise a GPE and avoid implicit solution by the second-order ODEs. Moreover, we eliminate all disadvantages of (1) and reduce the number of parameters while preserving the advantages of gravitational

alignment. Since we are solving a computer vision problem and are not interested in accurate physical simulations, we are free to alter laws of simulated physics, as long as it brings us closer to the desired properties.

Consider a *negative elementwise reciprocal transform operator* $\eta^-(\cdot)$ which acts on every element c of an arbitrary sum or a vector as $\eta^-(\frac{1}{c}) = c$. $\eta^-(\cdot)$ preserves function monotonicity and is reversibleⁱ. We apply $\eta^-(\cdot)$ to (1) and obtain:

$$\xi^-(\mathbf{U}(\mathbf{R}, \mathbf{t})) = \sum_{i,j} \frac{1}{Gm_{\mathbf{y}_i} m_{\mathbf{x}_j}} \|\mathbf{R}\mathbf{y}_i + \mathbf{t} - \mathbf{x}_j\|_2 + \epsilon. \quad (4)$$

Eq. (4) changes the model of the simulated world. In the inverse world, the potential between two particles is inversely proportional to the product of their masses and directly proportional to the distance between the points. The meaning of a mass has changed — now, the mass is comprehended as a property of the matter so that the less its value, the more significant is the interaction. Nevertheless, we can preserve the notion of mass and replace $\frac{1}{m_{\mathbf{y}_i}^{\text{old}}} = m_{\mathbf{y}_i}^{\text{new}}$ and $\frac{1}{m_{\mathbf{x}_j}^{\text{old}}} = m_{\mathbf{x}_j}^{\text{new}}$. Note that the softening parameter ϵ can be omitted, due to the vanished near-field singularities. Moreover, we get rid of G because it is a global multiplicative constant in the energy. Thus, our final form of GPE \mathbf{E} isⁱⁱ

$$\mathbf{E}(\mathbf{R}, \mathbf{t}) = \sum_i \sum_j m_{\mathbf{y}_i} m_{\mathbf{x}_j} \|\mathbf{R}\mathbf{y}_i + \mathbf{t} - \mathbf{x}_j\|_2. \quad (5)$$

With the new GPE (5), the optimal alignment is still achieved when GPE is locally minimal, while there is no reciprocity. It is noteworthy that in both cases, *i.e.* $\mathbf{U}(\mathbf{R}, \mathbf{t})$ and $\xi^-(\mathbf{U}(\mathbf{R}, \mathbf{t}))$, *the further are two particles apart from each other, the higher is the GPE between them:*

$$\lim_{\|\mathbf{R}\mathbf{Y} + \mathbf{T}\|_{\mathcal{F}} \rightarrow \|\mathbf{X}\|_{\mathcal{F}}} \mathbf{U}(\mathbf{R}, \mathbf{t}) = -\infty, \quad \text{and} \quad (6)$$

$$\lim_{\|\mathbf{R}\mathbf{Y} + \mathbf{T}\|_{\mathcal{F}} \rightarrow \|\mathbf{X}\|_{\mathcal{F}}} \xi^-(\mathbf{U}(\mathbf{R}, \mathbf{t})) = 0, \quad (7)$$

where $\|\cdot\|_{\mathcal{F}}$ denotes Frobenius norm, and we use a shorthand $\mathbf{T}_{D \times M} = [\mathbf{t} \quad \mathbf{t} \quad \dots \quad \mathbf{t}]$. Note that (5) still represents a globally multiply-linked GPE, with no explicit encoding of correspondences. These two factors reveal the deceptive visual similarity of the transformed GPE with the classic ICP formulation [6]. The further difference is that ICP requires alternating between correspondence and transformation updates, whereas (5) does not change throughout the optimisation (when not accelerated).

ⁱReciprocal transform is one of the widely-used transforms in signal processing and applied mathematics [7]. It has several modifications such as reciprocal-wedge transform in image processing [49].

ⁱⁱNext, we simplify the notation for masses and omit the superscript ^{new}.

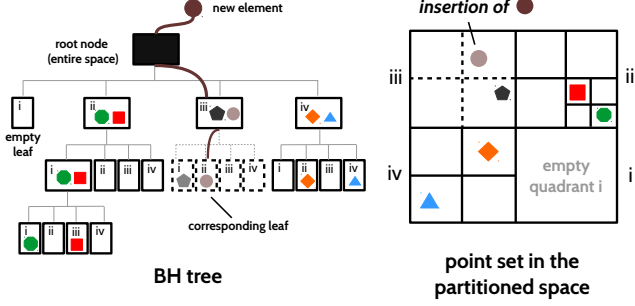


Figure 2: Building a BH tree exemplified. Each point is of different shape and colour for visualisation purposes. A new element either resides in a non-occupied leaf or initiates cell partitioning (while updating cluster masses and centres of mass) until it is placed in its own external leaf.

4.2. Acceleration with a Barnes-Hut 2^D -Tree

To surpass the quadratic computational complexity while computing potentials between points, we introduce a new variant of a BH tree or a 2^D -tree which is applied in the acceleration of N -body problems in astrodynamics [5]. We approximate exhaustive interactions of every template point \mathbf{y}_i with N reference points by interaction with $\mathcal{O}(\log N)$ many clusters which are fetched from the BH tree which we build. Whereas previous approaches *de facto* restrict interactions to local neighbourhoods [50, 38], we preserve global multiply-linking in our approximation.

Suppose \mathbf{y}_i interacts with K sufficiently distant particles \mathbf{x}_k , $k \in \{1, \dots, K\}$ and \mathbf{x}_k are sufficiently close to each other. Let us denote $\hat{\mathbf{y}}_i = \mathbf{R}\mathbf{y}_i + \mathbf{t}$ for brevity. Consider that variance of the distances from \mathbf{y}_i to every \mathbf{x}_k is below some small ζ . If located sufficiently far away, the total impact of all elements \mathbf{x}_k in a volume of space V to \mathbf{y}_i can be well approximated by the impact of the combined single particle $\tilde{\mathbf{x}}$. The mass of $\tilde{\mathbf{x}}$ equals to the mass integral of \mathbf{x}_k over V , and the position of $\tilde{\mathbf{x}}$ equals to the centre of mass of the elements \mathbf{x}_k in V [5]:

$$\sum_{k=1}^K m_{\mathbf{y}_i} m_{\mathbf{x}_k} \|\hat{\mathbf{y}}_i - \mathbf{x}_k\|_2 \approx m_{\mathbf{y}_i} \left[\sum_{k=1}^K m_{\mathbf{x}_k} \right] \|\hat{\mathbf{y}}_i - \tilde{\mathbf{x}}\|_2. \quad (8)$$

Building BH Tree. We build a joint BH tree on the point set union $\mathcal{X} \cup \mathcal{Y}$. Since template points are not influenced by each other, we set all $m_{\mathbf{y}_i}$ to zeroes. This allows us to include all points of the template into the tree but exclude their effect (mass) to other template points while calculating potentials. If \mathcal{X} and \mathcal{Y} contain duplicated points, the subdivision in a BH tree will continue infinitely long. In practice, the depth of a BH tree is restricted to avoid infinite splits.

A BH tree is initialised as a root node with 2^D empty external nodes. An insertion of every new \mathbf{y}_i always results in a new leaf (an occupied external node) and can initiate clusters (internal nodes). The size of daughter cells is always the half size of the parent cell along all dimensions. Every \mathbf{y}_i affects the mass and the centre of mass of every cluster

which includes \mathbf{y}_i . Fig. 2 schematically visualises BH tree building. We strictly follow the insertion rules proposed in [5]. Additionally, see our supplement for more details.

Using BH Tree. The potential at position \mathbf{y}_i is calculated by traversing the BH tree starting from the root towards the leaves and requires a single parameter — the distance threshold γ . Suppose l is the size of the current examined cell, and μ is the distance from \mathbf{y}_i to the centre of the cell. If $l / \mu < \gamma^{-1}$, then the cell’s influence on \mathbf{y}_i is accumulative. Otherwise, subcells on the next level are visited, and the examination is performed recursively until the condition $l / \mu < \gamma^{-1}$ is satisfied or the tree is entirely traversed [5].

To obtain a non-zero potential for the current \mathbf{y}_i , we restore its original mass. After the potential is computed, we set $m_{\mathbf{y}_i}$ again to zero in order to avoid interaction between template points. The earlier the node is used, the higher is the amount of individual particle-particle interactions which are bypassed. The time and memory requirements of BH tree building for a given $\mathcal{X} \cup \mathcal{Y}$ depend on how homogeneous point distribution is and have an upper bound, and the speed of GPE calculation depends on γ (see Sec. 4.4).

Fig. 3 visualises cluster configurations fetched during registration for a single template point and one instance of the *clean-500* experiment (see Sec. 5.1) with $\gamma = 5.0$. The input point clouds contain 817 points each, and the cardinality of the configurations fetched from the BH octree lies in the range [49; 62] points. Some of the clusters are massless and can be removed from the respective configuration as they comprise only massless template points. Distinct is the trajectory of the exemplary point over all iterations. Some distant particles are joined into clusters which are persistent over all iterations, whereas compositions of other clusters change dynamically from iteration to iteration while the template undergoes a rigid transformation.

GPE with a BH tree. The final GPE of BH-RGA with 2^D -tree acceleration for particle interactions reads

$$\mathbf{E}(\mathbf{R}, \mathbf{t}) = \sum_{\forall \mathbf{y}_i} m_{\mathbf{y}_i} \sum_{\mathbf{k}_j \in \mathcal{K}(\mathbf{y}_i)} m_{\mathbf{k}_j} \|\mathbf{R}\mathbf{y}_i + \mathbf{t} - \mathbf{k}_j\|_2, \quad (9)$$

where $\mathbf{k}_j \in \mathcal{K}(\mathbf{y}_i)$ denotes fetchable representations of the reference for every \mathbf{y}_i with clusters \mathbf{k}_j and masses $m_{\mathbf{k}_j}$.

4.3. Gravitational Potential Energy Optimisation

The proposed approach is summarised in Alg. 1. BH-RGA minimises a sum of squared residuals $f_r(\vartheta) = m_{\mathbf{y}_i} m_{\mathbf{k}_j} \|\mathbf{R}\mathbf{y}_i + \mathbf{t} - \mathbf{k}_j\|_2^2$ weighted by the product of masses. The number and composition of residuals depend on the BH tree configuration and γ . Since minimising GPE (9) is an overconstrained non-linear optimisation problem, we solve it in the least-squares sense with the Levenberg-Marquardt (LM) algorithm [34, 36] by iteratively linearising the GPE in the vicinity of the current solution. For the




Figure 3: Clusters fetched during alignment (*clean-500* experiment, see Sec. 5.1) from a BH octree: **a/** cluster configuration at the beginning; **b/** overlaid cluster configurations for all seven iterations. The colours in **a/** and **b/** encode the distance to the exemplary template particle (the darker, the further away), and massless clusters are shown in black; **c/** cluster configuration in the first iteration, for a single template point (shown as a triangle) and three different γ values. The colour scheme is given on the right. Each colour encodes a mass range of the cluster. Initially, all points were assigned unit masses.

Algorithm 1 Barnes-Hut Rigid Gravitational Approach

Input: reference $\mathbf{X} \in \mathbb{R}^{D \times N}$ and template $\mathbf{Y} \in \mathbb{R}^{D \times M}$

Output: rigid transformation $\vartheta = \{\mathbf{R}, \mathbf{t}\}$ aligning \mathbf{Y} to \mathbf{X}

- 1: **Initialisation:** $\mathbf{R} = \mathbf{I}$, $\mathbf{t} = \mathbf{0}$; masses $m_{\mathbf{y}_i}$ and $m_{\mathbf{x}_j}$, the distance threshold γ for cluster fetching from a 2^D -tree
 - 2: **while** not converged **do**
 - 3: build a 2^D -tree on the union of point sets $\{\mathbf{X}, \mathbf{R}\mathbf{Y} + \mathbf{T}\}$, $\mathbf{T}_{D \times M} = [\mathbf{t} \ \mathbf{t} \ \dots \ \mathbf{t}]$
 - 4: **for** every template point \mathbf{y}_i **do**
 - 5: fetch a set of clusters $\mathcal{K}(\mathbf{y}_i)$ from the 2^D -tree
 - 6: **end for**
 - 7: minimise (9), i.e., the GPE $\mathbf{E}(\mathbf{R}, \mathbf{t}) = \sum_{\mathbf{y}_i} m_{\mathbf{y}_i} \sum_{\mathbf{k}_j \in \mathcal{K}(\mathbf{y}_i)} m_{\mathbf{k}_j} \|\mathbf{R}\mathbf{y}_i + \mathbf{t} - \mathbf{k}_j\|_\epsilon$ with non-linear least squares
 - 8: **end while**
-

higher resistance to outliers, we apply Huber loss in the evaluation of every $f_r(\vartheta)$ defined as:

$$\|a\|_\epsilon = \begin{cases} \frac{1}{2}a^2, & \text{if } |a| \leq \epsilon \\ \epsilon(|a| - \frac{1}{2}\epsilon), & \text{otherwise.} \end{cases} \quad (10)$$

The BH tree is renewed after every successful $\Delta\vartheta$ update, and the objective is formulated for the new set of $f_r(\vartheta)$ in every external iteration. The number of internal LM iterations is kept in the range [3; 10], see Fig. 1. BH-RGA converges when the difference in two consecutive GPE values is below some threshold ρ_E .

4.4. Computational Complexity

Building a BH tree jointly for the union $\{\mathcal{X}, \mathcal{Y}\}$ in every iteration is performed in $\mathcal{O}((M + N) \log(M + N))$ time. The complexity of fetching a set of clusters from the reference for one template point is $\mathcal{O}(\log N)$. In total, this operation is performed for all M template points, and we arrive at $\mathcal{O}(M \log N)$. Thus, the computational complexity of BH-RGA is $\mathcal{O}((M + N) \log(M + N) + M \log N) \approx \mathcal{O}(M \log N)$ (assuming $M \sim N$). This is a significant improvement compared to the quadratic complexity of GA [21], especially for large point sets.

4.5. Defining Boundary Conditions With Masses

Setting varying masses of template and reference points in BH-RGA allows for different effects such as embedding of prior correspondences, additional cues and counterbalancing inhomogeneous point densities.

Sparse Prior Matches. Different masses in BH-RGA can encode prior correspondences as the optimal alignment is more likely when points with larger masses are close to each other. Suppose $(i, j) \in N_c \subset \mathbb{N}^2$ is the set of points for which correspondences are known in advance. The GPE of BH-RGA which considers prior matches is then specified as

$$\mathbf{E}^p(\mathbf{R}, \mathbf{t}) = \begin{cases} \mathbf{E}(\mathbf{R}, \mathbf{t}) \text{ as in (9)}, & \forall \mathbf{y}_i : i \notin N_c, \\ m_{\mathbf{y}_i}^p m_{\mathbf{x}_j}^p \|\mathbf{R}\mathbf{y}_i + \mathbf{t} - \mathbf{x}_j\|_2, & \text{else,} \end{cases} \quad (11)$$

with $m^p = m_{\mathbf{y}_i}^p = m_{\mathbf{x}_j}^p$ being, as a rule, several orders of magnitude larger than default masses (depending on the total mass distribution and the reliability of prior matches). Note that for points in N_c , multiply-linking and cluster fetching are avoided. We refer to *anchor points* as two subsets of \mathcal{X} and \mathcal{Y} for which it is known that they correspond among each other. Anchor points can be embedded into BH-RGA by setting their masses to m_p , with no changes in (9). Thus, the effect of anchor points in BH-RGA is closer to *alignment in alignment* or a hierarchical registration.

Point Colours. Additional alignment cues such as point colour can be converted to masses in BH-RGA. For the GPE to be locally minimal in this setting, points with similar masses have to coincide or be close to each other.

Handling Varying Point Densities. Handling of point sets with inhomogeneous densities has been indicated as challenging for many approaches, with specialised techniques and pre-processing steps (such as re-sampling) required to handle the setting efficiently [30]. For BH-RGA, we propose to normalise total point mass per unit volume in linear time. In *volumetric mass normalisation* (VMN), the bounding box of a point set is split into multiple non-overlapping voxels, and in every voxel, the pre-defined mass is uniformly distributed among all points belonging to it.

5. Experiments

In this section, we describe the evaluation methodology and results. BH-RGA is implemented in C++ with *eigen3* [27] and *Galois* [42] libraries. For NLLS, we use *ceres* solver [2]. The experiments are performed on a system with 32 Gb RAM and quad-core Intel *i7-6700K* processor running at 4GHz. We compare our BH-RGA to Iterative Closest Point (ICP) [6], LM-ICP [18], Coherent Point Drift (CPD) [38], GMM Registration (GMR) [32] and GA [21]. CPD and GMR are evaluated both as 6 DoF (\mathbf{R} , \mathbf{t}) and 7 DoF (\mathbf{R} , \mathbf{t} and scale s) variants. We also compare BH-RGA against rigid Extended CPD (E-CPD) guided by prior matches [23]. ICP is taken from the *Matlab* repository [51], and implementations of several methods are publicly available [37, 31, 35]. In all experiments, we report the mean *root-mean-square error* (RMSE) and the standard deviation of the RMSE denoted by σ .

5.1. Quantitative Evaluation

Main Tests. In the tests with synthetic data, we use a sub-sampled *Stanford bunny* [47]. In the first *clean-500* test, we sample the 3D rotation space by the angular displacement of $\frac{2\pi}{10}$ radians and synthesise 500 different initial configurations (the duplicated states are avoided). In this experiment, we test the methods for their ability to resolve rotation in scenarios with noiseless data. In contrast, the second test is designed to evaluate the ability of RA to converge under severe amounts of noise. We add 100% of uniform noise to the *clean-500* and obtain *N500-U100* dataset. The range of the noise generating function is a bounding sphere of the template. Several methods demonstrate poor performance on *N500-U100*, and we also generate *N500-U50* dataset with 50% of uniform noise according to the same rule.

Next, we choose a single initial configuration which can be successfully resolved by all methods and add three different noise patterns to the template — 100% of uniform noise, 100% of Gaussian noise (the range volume is the same as in *N500-U100*) and 100% of the Gaussian noise with the mean values coinciding with the template point locations. The resulting datasets are abbreviated as *100U*, *100G* and *100GS*, respectively. Table 1 summarises the outcomes of the experiments with the new datasets. If RMSE is below 0.1, the registration is assumed to succeed. This value is not set arbitrarily — in the case of successful alignments, RMSE was, in most cases, much lower than 0.1. Otherwise, it was in most cases considerably larger than 0.1.

Upon our expectation, BH-RGA is not the most accurate method on the *clean-500* dataset (it is the second-best regarding the success rate). Perfectly matching data is, however, not common in practical applications. With an increasing noise, the relative accuracy and performance of BH-RGA compared to all other methods increases. While CPD (7 DoF) achieves a lower RMSE than BH-RGA on

clean-500, the situation inverts on *N500-U50* and *N500-U100*, with the success rates of BH-RGA close to those of CPD (and the lowest RMSE compared to CPD). At the same time, ICP, LM-ICP, CPD (6 DoF) and GMR considerably lower their success rates on *N500-U50* and *N500-U100*. BH-RGA was the only successful method on *U100* — it resolves all cases with the lowest RMSE. The second-best method is CPD (7 DoF, 48 cases out of 50), while demonstrating a 50% to 66% higher RMSE on *N500-U100* and *U100*, respectively. GMR is able to hold on par with ICP/LM-ICP only in the 6 DoF mode. On *G100* and *GS100*, most methods successfully recover transformations for all cases, except GMR (6 DoF) and GA which recover 52%, 25%, 2% and 20% of transformations, respectively.

The performance of baseline GA [21] in our tests is, in most cases, inferior to the results of other methods. Our experimental setting is different from the one proposed in [21] where the angle of initial misalignment is randomly selected in the range $[0; \frac{\pi}{2}]$. Golyanik *et al.* [21] report that GA was starting to fail in the angular misalignment range $[\frac{\pi}{4}, \frac{\pi}{2}]$. In our *clean-500* and *N500-* experiments, the angles are sampled with the 36° step implying that starting from the second value (out of ten), the angle is always $\geq 72^\circ$. Regarding *U100*, *G100* and *GS100*, the excessive amount of noise hinders GA to converge, though the angle of initial misalignment is in the convergence basin of GA.

Perturbed Data. Next, we propose a test for evaluating how robust are the methods to the corrupted data — the *U256* and *G256* experiments. We perturb template points with uniform and Gaussian noise. In total, there are 256 magnitude indexes generating states with the increasing degree of perturbation. The maximum displacement in every state is equal to the scaled magnitude index times the length of the bounding box of the object in the x -direction. In both experiments, no point of the template physically coincides with any reference point at the optimum. RMSE between the reference and the template transformed by the recovered \mathbf{R} , \mathbf{t} is reported in Table 2. CPD and GMR have higher RMSE than ICP and BH-RGA, *i.e.*, both probabilistic methods are sensitive to point perturbations, which agrees with our expectations. ICP shows the lowest RMSE. Especially for small perturbations, conditions for ICP are optimal. BH-RGA slightly concedes to ICP as its RMSE is a little higher for small perturbation magnitudes. On the other side, BH-RGA has the smallest σ .

Prior Matches and Anchor Points. We repeat the *N500-U50* test with up to three prior matches and anchor points, and compare BH-RGA to Extended CPD (E-CPD) [23], see Table 3 for the summary. The latter approach is a correspondence-conditioned CPD variant. Especially with only one or two reliable prior matches, E-CPD has advantages [23]. Thus, CPD with no prior matches resolves 17% of the cases and its success rate increases to 18% and 75%

Table 1: Summary of the quantitative evaluation of the compared methods on *clean-500*, *N500-U50*, *N500-U100*, *U100*, *G100* and *GS100* datasets.

methods and metrics		<i>clean-500</i>	<i>N500-U50</i>	<i>N500-U100</i>	<i>U100</i>	<i>G100</i>	<i>GS100</i>
ICP [6]	success (in %)	62 (12.4%)	36 (7.2%)	19 (3.8%)	33 (66%)	50 (100%)	50 (100%)
	RMSE (σ)	0.005 (0.016)	0.022 (0.3)	0.042 (0.031)	0.091 (0.081)	0.007 (0.002)	0.002 (5E-4)
LM-ICP [18]	success (in %)	123 (24.6%)	82 (16.4%)	72 (14.4%)	49 (98%)	50 (100%)	50 (100%)
	RMSE (σ)	0.002 (6.3E-4)	0.015 (0.009)	0.023 (0.017)	0.025 (0.021)	0.006 (0.003)	0.003 (0.001)
CPD (7 DoF) [38]	success (in %)	130 (26%)	128 (25.6%)	109 (21.8%)	48 (96%)	50 (100%)	50 (100%)
	RMSE (σ)	0.04 (6E-5)	0.064 (0.003)	0.088 (0.003)	0.098 (0.066)	0.027 (1.7E-3)	0.046 (1.7E-3)
CPD (6 DoF) [38]	success (in %)	143 (28.6%)	98 (19.6%)	62 (12.4%)	48 (96%)	50 (100%)	50 (100%)
	RMSE (σ)	0.006 (0.009)	0.025 (0.014)	0.034 (0.017)	0.061 (0.148)	0.01 (0.003)	7.5E-3 (2.3E-3)
GMR (7 DoF) [32]	success (in %)	126 (25.2%)	113 (22.6%)	0 (0%)	0 (0%)	50 (100%)	50 (100%)
	RMSE (σ)	7E-5 (8E-5)	0.084 (0.005)	n / a	n / a	0.04 (7.5E-3)	1.5E-3 (4.7E-4)
GMR (6 DoF) [32]	success (in %)	131 (26.2%)	79 (15.8%)	87 (17.4%)	36 (72%)	26 (52%)	25 (25%)
	RMSE (σ)	8E-5 (2E-4)	5E-4 (2E-4)	6E-4 (2E-4)	0.16 (0.077)	0.37 (0.4)	0.37 (0.38)
GA [21]	success (in %)	21 (4.2%)	6 (1.2%)	3 (0.6%)	19 (38%)	1 (2%)	10 (20%)
	RMSE (σ)	0.029 (0.021)	0.049 (0.025)	0.03 (0.012)	0.164 (0.082)	0.289 (0.087)	0.163 (0.072)
BH-RGA (ours)	success (in %)	132 (26.4%)	132 (26.4%)	100 (20%)	50 (100%)	50 (100%)	50 (100%)
	RMSE (σ)	0.009 (0.004)	0.032 (0.013)	0.059 (0.021)	0.056 (0.017)	0.04 (0.01)	0.022 (0.006)

Table 2: RMSE and σ (in parentheses) in *U256* and *G256* experiments.

method	<i>U256</i>	<i>G256</i>
ICP [6]	9E-3 (7E-3)	0.015 (0.012)
LM-ICP [18]	0.077 (0.041)	0.113 (0.063)
CPD (7 DoF) [38]	0.051 (9E-3)	0.064 (0.021)
CPD (6 DoF) [38]	0.016 (0.014)	0.045 (0.058)
GMR (7 DoF) [32]	0.019 (0.028)	0.065 (0.051)
GMR (6 DoF) [32]	0.853 (1.16)	1.027 (1.226)
GA [21]	0.149 (0.143)	0.207 (0.158)
BH-RGA (ours)	0.015 (4E-3)	0.019 (9.5E-3)

with one and two prior matches, respectively. BH-RGA resolves 25% with no priors and $\sim 39\%$ with one or two available anchor points. Recall that anchor points have higher masses and interact with all other points, which is the reason that the success with three anchor points drops to 33%. In contrast, BH-RGA resolves all cases with two or three prior correspondences and outperforms E-CPD. Prior matches are not involved in the interaction with other points, which is the strongest boundary condition.

Runtime and Convergence. We evaluate the runtime of BH-RGA for the inputs of different sizes with different thresholds γ . For each approximation level of the BH tree, we report the achieved accuracy and runtime, among others.

We take a frame from the SINTEL *sleeping2* sequence [9] as a reference and its translated and rotated clone as a template. The translation roughly amounts to one-third of the point cloud extent in the x -direction, and the rotation is either 5° or 24° around all axes simultaneously. In total, there are seven versions of the point cloud obtained by subsampling, with the cardinalities of *ca.* 5k, 12k, 25k, 50k, 105k, 205k and 446k (the original resolution). For each combination of the resolution and template transformation, we test $\gamma \in [0.0625; 64.0]$. The runtime evaluation metrics

Table 3: Comparison of E-CPD [23] and BH-RGA.

evaluated configuration	RMSE	σ	cases (success rate)	
E-CPD [23], prior matches only	[no priors]	0.08	0.013	85 (17%)
	1 prior	0.084	0.012	89 (18%)
	2 priors	0.076	0.014	376 (75%)
	3 priors	0.056	0.018	488 (97%)
BH-RGA (ours), anchor points	[no priors]	0.05	0.018	124 (25%)
	1 prior	0.043	0.013	199 (40%)
	2 priors	0.011	0.017	191 (38%)
	3 priors	0.0086	0.005	165 (33%)
BH-RGA (ours), prior matches	1 prior	0.0191	1.1E-4	435 (87%)
	2 priors	0.0055	1.72E-5	500 (100%)
	3 priors	0.0001	4.5E-9	500 (100%)

for 5° rotations are summarised in Fig. 4, *i.e.*, runtimes (full and per iteration), runtime ratio of the BH tree generation, RMSE and the number of residuals.

Several patterns can be noticed in the runtime statistics. When γ is fixed, RMSE is similar for a different number of points. Conversely, with the increasing γ , the alignment accuracy increases for all resolutions, and smaller and smaller runtime fraction is spent on building BH tree while not being influenced much by the number of points. For $\gamma > 2.0$, BH tree runtime ratio is $\leq 0.5\%$. In comparison, the solver runtime ratio is always $\sim 80\%$ suggesting that an up to five-fold acceleration with a Gauss-Newton solver on a GPU is possible. Next, as γ increases, the number of iterations decreases and stabilises. For $\gamma \geq 6$, the number of iterations is ≤ 4 . The average number of residuals (number of point-to-cluster interactions) reaches $\approx 1.8E7$ for $\gamma = 6.0$ and 446k points. The graph metrics for the case of 24° rotations coarsely follow the shapes of the graph metrics for the visualised case of 5° rotations. Larger misalignments require larger γ for successful registration and the overall runtime increases by the factor of two (note that in this experiment, we do not perform a translational pre-alignment). The per-iteration runtime increases by $\sim 40\%$, and the number of iterations for $\gamma \geq 6$ increases by one-two.

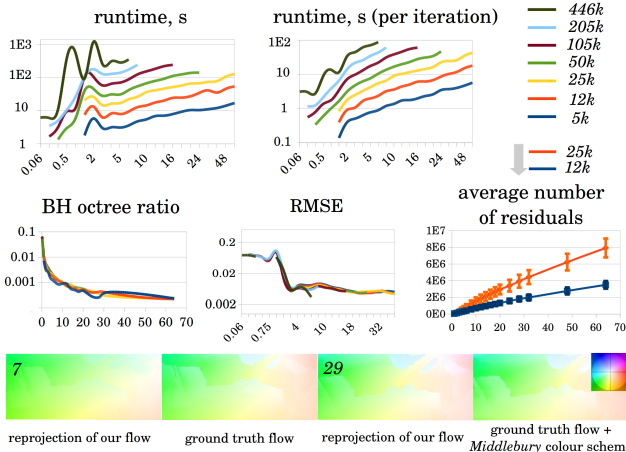


Figure 4: Graphs of different runtime evaluation metrics as the functions of the BH tree threshold γ (SINTEL *sleeping2* sequence [9], different subsampling rates and a different number of points in the point clouds). In the bottom row, we show recovered projected displacements for frames 7 and 29 following [22] and the used *Middlebury* optical flow encoding [4].

The study shows that BH-RGA can register two point sets of 446k points each in 11 seconds with RMSE of 0.068 ($\gamma = 0.25$). RMSE of 0.0065 requires 264 seconds ($\gamma = 3.0$). A comparable RMSE of 0.0067 for the point sets with 25k points takes 41 seconds ($\gamma = 12.0$). BH-RGA registers large point sets in a globally multiply-linked manner in tolerable time, whereas tested implementations of competing methods cannot cope with the data. The CPD paper [38] reports 10 minutes for two point sets with $\approx 35k$ points each with acceleration techniques. BH-RGA would require ~ 1 second for this task with $\gamma = 0.25$.

5.2. Evaluation with Real-World Data

BH-RGA is well suitable for registration of real-world data, including large partially overlapping point sets and lidar measurements. Fig. 5 comprises several such scenarios. The *Vestalin* and *pedestal* datasets represent partial 3D reconstructions of statues acquired with a structured light technique [44]. Due to the scanning process, they are represented in different reference frames. *Vestalin* and *pedestal* are processed with three and two anchor points, respectively. For *Vestalin*, we additionally vary masses based on point intensities. In *pedestal*, the anchor points are selected in the lower left and right corners of the plate. *Universe* dataset represents two 3D reconstructions of sculpture (the shapes differ mostly in the lower parts).

We also test BH-RGA on lidar data from the KITTI dataset [19]. Fig. 5-(right) shows an excerpt from the *2011_09_26_drive_0001* sequence. We take the first and third frames of the sequence so that displacements are moderate and the outlier ratio amounts to *ca.* 20%. Due to highly varying point densities in the scans, we apply VMN (see




Figure 5: BH-RGA is applicable to partially overlapping real-world data. *Vestalin* is aligned with the colour cue and three anchor points, and for *pedestal*, we use two anchor points. For the lidar data, we apply VMN helping to counterbalance varying point densities. For the *Vestalin* and the *universe*, the initialisation is shown on the left, and the alignments are shown on the right. For the *pedestal* and *driving scenario*, the initialisations are on the top, and the alignments are on the bottom.

Sec. 4.5) which initialises masses uniformly per unit volume. It can be well noticed on buildings and cars that the scans are accurately registered. In this experiment, no anchor points or further prior knowledge are used. Both scans contain *ca.* 120k points, and the registration is performed in ≈ 90 seconds on a CPU without subsampling or parallelisation, with $\gamma = 1.5 \cdot 10^4$ (at the original point set scale).

6. Concluding Remarks

The main conclusion of this paper is that particle dynamics based RA can be accelerated using a BH tree and stabilised by altering laws of simulated physics. Our formulation significantly reduces the number of parameters in the gravitational alignment — by the factor of three — and eliminates the need for parameter tuning. Even for large point sets of dimensionality going beyond the extents feasible for other general-purpose RA methods, the proposed BH-RGA performs global multiply-linked updates. This contrasts with other known techniques which *de facto* only follow locally multiply-linked policies for the approximation of point interactions. As a result, BH-RGA is notably robust to noise and achieves state-of-the-art performance whenever existing techniques do. Furthermore, BH-RGA accepts such additional cues as prior matches, anchor points and even point intensities, and efficiently handles different point densities with VMN.

There are multiple avenues for further investigation of particle dynamics based RA. The most promising ones are multibody registration, parallelisation and adaption of BH-RGA for CAD and automotive applications.

Our supplementary material contains more details on the 2^D -tree generation and projective scene flow visualisation used in Fig. 4. Besides, we visualise statistics of resolved misalignments for *clean-500* and *N500-U100* experiments.

References

- [1] Sverre J. Aarseth. *Gravitational N-body Simulations: Tools and Algorithms*. Cambridge University Press, 2003. 2, 3
- [2] Sameer Agarwal, Keir Mierle, and Others. Ceres solver. <http://ceres-solver.org>. 6
- [3] Sk A. Ali, Vladislav Golyanik, and Didier Stricker. Nrga: Gravitational approach for non-rigid point set registration. In *International Conference on 3D Vision (3DV)*, 2018. 2
- [4] Simon Baker, Daniel Scharstein, J. P. Lewis, Stefan Roth, Michael J. Black, and Richard Szeliski. A database and evaluation methodology for optical flow. *International Journal of Computer Vision (IJCV)*, 92(1), 2011. 8
- [5] Josh Barnes and Piet Hut. A hierarchical $O(n \log n)$ force-calculation algorithm. *Nature*, 324:446–449, 1986. 2, 4
- [6] Paul J. Besl and Neil D. McKay. A method for registration of 3-d shapes. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 14(2):239–256, 1992. 1, 2, 3, 6, 7
- [7] J. Martin Bland and Douglas G. Altman. Statistics notes: The use of transformation when comparing two means. *British Medical Journal (BMJ)*, 312(7039):1153, 1996. 3
- [8] Benedict Brown and Szymon Rusinkiewicz. Global non-rigid alignment of 3-D scans. *ACM Transactions on Graphics (TOG)*, 26(3), 2007. 1
- [9] Daniel Jonas Butler, Jonas Wulff, Garrett B. Stanley, and Michael J. Black. A naturalistic open source movie for optical flow evaluation. In *European Conference on Computer Vision (ECCV)*, pages 611–625, 2012. 7, 8
- [10] Dylan Campbell and Lars Petersson. An adaptive data representation for robust point-set registration and merging. In *International Conference on Computer Vision (ICCV)*, 2015. 2
- [11] Yang Chen and Gérard Medioni. Object modelling by registration of multiple range images. *Image and Vision Computing (IVC)*, 10(3):145 – 155, 1992. 1, 2
- [12] Haili Chui and Anand Rangarajan. A feature registration framework using mixture models. In *Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA)*, pages 190–197, 2000. 1, 2
- [13] Martin Danelljan, Giulia Meneghetti, Fahad Shahbaz Khan, and Michael Felsberg. A probabilistic framework for color-based point set registration. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [14] Yan Deng, Anand Rangarajan, Stephan Eisenschenk, and Baba C. Vemuri. A riemannian framework for matching point clouds represented by the schrödinger distance transform. In *Computer Vision and Pattern Recognition (CVPR)*, pages 3756–3761, 2014. 1, 2
- [15] Ben Eckart, Kihwan Kim, and Jan Kautz. Hgmr: Hierarchical gaussian mixtures for adaptive 3d registration. In *The European Conference on Computer Vision (ECCV)*, 2018. 1
- [16] Ben Eckart, Kihwan Kim, Alejandro Troccoli, Alonzo Kelly, and Jan Kautz. Mlmd: Maximum likelihood mixture decoupling for fast and accurate point cloud registration. In *International Conference on 3D Vision (3DV)*, 2015. 2
- [17] Jan Elseberg, Stephane Magnenat Rol, and Siegwart Andreas Nüchter. Comparison of nearest-neighbor-search strategies and implementations for efficient shape registration. *Journal of Software Engineering for Robotics*, pages 2–12, 2012. 2
- [18] Andrew W. Fitzgibbon. Robust registration of 2D and 3D point sets. In *British Machine Vision Conference (BMVC)*, pages 662–670, 2001. 1, 2, 6, 7
- [19] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013. 8
- [20] Steven Gold, Anand Rangarajan, Chien-Ping Lu, and Eric Mjolsness. New algorithms for 2d and 3d point matching: Pose estimation and correspondence. *Pattern Recognition*, 31:957–964, 1997. 1, 2
- [21] Vladislav Golyanik, Sk A. Ali, and Didier Stricker. Gravitational approach for point set registration. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 1, 2, 3, 5, 6, 7
- [22] Vladislav Golyanik, Kihwan Kim, Robert Maier, Matthias Nießner, Didier Stricker, and Jan Kautz. Multiframe scene flow with piecewise rigid motion. In *International Conference on 3D Vision (3DV)*, 2017. 8
- [23] Vladislav Golyanik, Bertram Taetz, and Didier Stricker. Joint pre-alignment and robust rigid point set registration. In *International Conference on Image Processing (ICIP)*, pages 4503–4507, 2016. 2, 6, 7
- [24] Vladislav Golyanik and Christian Theobalt. Optimising for scale in globally multiply-linked gravitational point set registration leads to singularities. In *International Conference on 3D Vision (3DV)*, 2019. 3
- [25] Sébastien Granger and Xavier Pennec. Multi-scale em-icp: A fast and robust approach for surface registration. In *Euro-pean Conference on Computer Vision (ECCV)*, pages 418–432, 2002. 1, 2
- [26] Michael A. Greenspan and Guy Godin. A nearest neighbor method for efficient icp. In *International Conference on 3-D Digital Imaging and Modeling*, pages 161–168, 2001. 2
- [27] Gaël Guennebaud, Benoît Jacob, et al. Eigen v3. <http://eigen.tuxfamily.org>, 2010. 6
- [28] Berthold K. P. Horn. Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America A*, 4(4):629–642, 1987. 2
- [29] Berthold K. P. Horn, Hugh M. Hilden, and Shahriar Negahdaripour. Closed-form solution of absolute orientation using orthonormal matrices. *Journal of the Optical Society of America*, 5(7):1127–1135, 1988. 2, 3
- [30] Felix J. Lawin, Martin Danelljan, Fahad S. Khan, Per-Erik Forssén, and Michael Felsberg. Density adaptive point set registration. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. 5
- [31] Bing Jian. Matlab implementation of gmr. <https://github.com/bing-jian/gmmreg/tree/master/MATLAB>, 2008. 6
- [32] Bing Jian and Baba C. Vemuri. Robust point set registration using gaussian mixture models. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 33(8):1633–1645, 2011. 2, 6, 7
- [33] Christian Kerl, Jürgen Sturm, and Daniel Cremers. Dense visual slam for rgb-d cameras. In *International Conference on Intelligent Robot Systems (IROS)*, 2013. 1

- [34] Kenneth Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly Journal of Applied Mathematics*, II(2):164–168, 1944. 4
- [35] Matlab computer vision routines. https://github.com/markeroon/matlab-computer-vision-routines/tree/master/third_party/lmicp, 2012. 6
- [36] Donald W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *SIAM Journal on Applied Mathematics*, 11(2):431–441, 1963. 4
- [37] Andriy Myronenko. Coherent point drift (cpd) project page. <https://sites.google.com/site/myronenko/research/cpd>. 6
- [38] Andriy Myronenko and Xubo Song. Point-set registration: Coherent point drift. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2010. 1, 2, 4, 6, 7, 8
- [39] Richard A. Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J. Davison, Pushmeet Kohli, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 127–136, 2011. 1
- [40] Matthias Nießner, Michael Zollhöfer, Shahram Izadi, and Marc Stamminger. Real-time 3d reconstruction at scale using voxel hashing. *ACM Transactions on Graphics (TOG)*, 32(6):169:1–169:11, 2013. 1
- [41] Andreas Nüchter, Kai Lingemann, and Joachim Hertzberg. Cached k-d tree search for icp algorithms. In *International Conference on 3-D Digital Imaging and Modeling (3DIM)*, pages 419–426, 2007. 1, 2
- [42] Keshav Pingali, Donald Nguyen, Milind Kulkarni, Martin Burtscher, M. Amber Hassaan, Rashid Kaleem, Tsung-Hsien Lee, Andrew Lenharth, Roman Manevich, Mario Méndez-Lojo, Dimitrios Proutzos, and Xin Sui. The tao of parallelism in algorithms. In *Programming Language Design and Implementation (PLDI)*, pages 12–25, 2011. 6
- [43] Anand Rangarajan, Haili Chui, and Fred L. Bookstein. The softassign procrustes matching algorithm. In *Information Processing in Medical Imaging (IPMI)*, pages 29–42, 1997. 1
- [44] Claudio Rocchini, Paolo Cignoni, Claudio Montani, Paolo Pingi, and Roberto Scopigno. A low cost 3 d scanner based on structured light. In *Eurographics*, 2001. 8
- [45] Szymon Rusinkiewicz and Marc Levoy. Efficient variants of the icp algorithm. In *International Conference on 3-D Digital Imaging and Modeling*, pages 145–152, 2001. 2
- [46] Renato F. Salas-Moreno, Richard A. Newcombe, Hauke Strasdat, Paul H. J. Kelly, and Andrew J. Davison. Slam++: Simultaneous localisation and mapping at the level of objects. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1352–1359, 2013. 1
- [47] The stanford 3d scanning repository. <http://graphics.stanford.edu/data/3Dscanrep/>. 6
- [48] Gary K. Tam, Zhi-Quan Cheng, Yu-Kun Lai, Frank Langbein, Yonghuai Liu, A. David Marshall, Ralph Martin, Xianfang Sun, and Paul Rosin. Registration of 3d point clouds and meshes: A survey from rigid to nonrigid. *Transactions on Visualization and Computer Graphics (TVCG)*, 19(7):1199–1217, 2013. 1
- [49] Frank Tong and Li Ze-Nian. Reciprocal-wedge transform for space-variant sensing. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 17(5):500–511, 1995. 3
- [50] Yanghai Tsin and Takeo Kanade. A correlation-based approach to robust point set registration. In *European Conference on Computer Vision (ECCV)*, pages 558–569, 2004. 1, 2, 4
- [51] Jakob Wilm. Iterative closest point. <https://de.mathworks.com/matlabcentral/fileexchange/27804-iterative-closest-point>. Version 1.14. 6
- [52] Ruigang Yang and Peter K. Allen. Registering, integrating, and building cad models from range data. In *International Conference on Robotics and Automation (ICRA)*, pages 3115–3120, 1998. 1