

Force Field Based n -Scan Alignment

Rolf Lakaemper* Nagesh Adluru* Longin Jan Latecki*

*Department of CIS, Temple University, Philadelphia, USA

Abstract— We present a force field based approach for simultaneous alignment of multiple laser scans in robot mapping. It avoids sensitive behavior to wrong data associations and sparse sensing, which are the main challenges e.g. in multi robot mapping under the constraints given in autonomous search and rescue robotics. The presented algorithm solves the alignment problem utilizing a gradient descent approach motivated by physics, but exchanges laws of physics with constraints given by human perception. Experiments on different real world data sets show the successful application of the algorithm.

Index Terms— Scan alignment, force fields, SLAM

I. INTRODUCTION

The problem of aligning n scans has been treated as estimating sets of poses [26]. Since estimating the sets of poses involves estimating the consistency of the associated observations (maps), this joint estimation is called Simultaneous Localization and Mapping. The conditional independence between these two estimations is e.g. the key for Rao-Blackwellization (factoring the posterior of maps) of particle filters for SLAM [27].

There have been several algorithms for the estimation [30, 17, 18, 34]. The underlying framework for all such techniques is to optimize a constraint-graph, in which nodes are features, poses and edges are constraints built using various observations and measurements like odometry, scan-matching of range scans. These techniques differ in

- how they represent graphs - e.g. [17] uses a sophisticated data structure called Tree-map. [34] represents using inverse covariance matrix called information matrix instead of covariance matrix. The information matrix *naturally* happens to be sparse for the SLAM posterior, which gives them computational advantages.
- how they build constraints - e.g. [26] uses linearized constraints obtained from scan-matching and odometry, [30] works with non-linear constraints.
- how they optimize the graphs - e.g. [30] uses stochastic gradient descent for approximate optima, borrowing the ideas from learning theory. While [26] solves for exact optima using brute-force, as noted earlier their graph uses linear approximation of the non-linear constraints that occur in SLAM. [18] use Gauss-Seidel relaxation again for approximate optima.

All these approaches have performed well in many practical cases but they have one drawback that is they are sensitive to behavior of error models of sensors because of several assumptions and approximations which might not hold with sparse sensing. For e.g. [26] linearizes constraints by linearizing pose-relations, solving a linear equation of the form $AX = B$ to estimate X , the set of poses. This requires that A is invertible,

so they conjecture that A is invertible if the constraint-graph is fully connected and the errors of the observations behave in a gaussian/normal way. [10] extends the same technique for 3D scans.

[30] presents an approximate optimization of non-linear constraints and demonstrate that their approach of approximating the optimization process in non-linear state space yields superior results compared to finding exact optima by approximating a non-linear state space (SLAM) to a linear state space.

Another strategy of attacking the problem is to treat the problem of SLAM from a perspective of aligning n scans *simultaneously*. The algorithms exploiting this perspective build from image registration techniques, the most famous being Iterative Closest Point (ICP) [5, 12] and it's numerous variants to improve speed and converge basins [33] and [25, 8]. Such algorithms are not very sensitive to error models of the sensors as they do not explicitly depend on the error models. Basically all these techniques do search in transformation space trying to find the set of pair-wise transformations of scans by optimizing some function defined on transformation space. The techniques vary in defining the optimization functions that range from being error metrics like "sum of least square distances" to quality metrics like "image distance" as in [9] or probabilistic metrics [23]. Their optimization process itself can be gradient descent or hill climbing or using genetic programming strategy as in [31] or using maximum a posteriori (MAP) [28]. All of these techniques have one major limitation, which is they search in *pair-wise* transformation space. Though in some variants of ICP the error from all pair-wise transformations is spread across all transformations to simultaneously align all scans, the procedure can be highly sensitive to outliers [32]. Fig. 1 for example shows the difference between the results of aligning a hypothetical set of 3 scans using classical ICP and our approach.

II. MOTIVATION

Though we also adapted the perspective of aligning n scans, we differ from any other typical registration algorithm in that we do not search explicitly in transformation space for optimization. We embed the transformations (poses of the scans) into a higher dimensional space of global configurations of the scans and search for poses. As a stable configuration evolves the optimal set of transformations of scans is reached. This search in high dimensional space at first sight seems very complicated, demanding computation of a high dimensional gradient; but fortunately using potential field simulation for various computer vision tasks like contour detection, segmentation, registration has been empirically successful [38, 22, 37]




Fig. 1. The top row shows 3 steps of the alignment of 3 scans (each scan consists of a single corner only) by classical ICP, the bottom-row shows the results of the proposed approach. The alignment progress can be seen from left to right in both cases. The square boxes show the robot poses of the scans.

and [2]. Since mapping is closely related to registration, the approaches whose motivations are closely related to our approach are [7, 2, 14]. In [14] they align range scans by moving them simultaneously. The movements are not just based on the minimizing error of transformation computed using correspondences but on the simulated fields generated by imaginary springs attached to the corresponding points. Our technique differs from [14] in that the force field is generated not just by closest point correspondences but using perceptual principles and gaussian fields similar to [11]. Since our force field (Eq. (3)) is radial, it is conservative and hence the simulated *free* movements of scans with weights defined using our perceptual principles will be towards minimizing potential energy (Eq. (1)). Thus our main contribution is the novel design of force fields for simultaneous multi-scan alignment. [7] also performs search in $3n$ dimensional space. For each configuration they compute energy as the sum of the Normal Distribution Transforms (NDT) [6] of all the scans in the configuration and update the configuration using Newton's optimization algorithm that involves the first and second derivatives of the energy. Their approach is very closely related to ours but does not use perceptual features and rigid body dynamics and hence in principle can be more sensitive to outliers.

III. FORCE FIELD SIMULATION FFS

A. Rigid bodies & conservation of energy

Rigid bodies: Each scan is a set of points (x, y) (obtained from range data) at a pose (x, y, θ) . The pose can be either obtained by odometry or odometry independent by relative pose estimation using pre-alignment techniques described in [1]. The set of sample points and the pose are termed 'rigid body', since the relative positions of the points in the single scan do not change. These rigid bodies constitute the physical system on our approach. A global configuration or global map of n rigid bodies ($rb_1, rb_2, \dots, rb_n$) at time t is represented as RB_t . FFS tries to find the global map which is optimal with respect to an underlying fitness measure. It uses an iterative gradient approach with a decreasing step width control. The

following chapters will motivate the fitness measure as well as the gradient technique.

Conservation of energy: The law of conservation of energy is one of the most important results in Physics [16, 29]. Free moving bodies in a conservative force field, which forms a closed system, do not result in energy gain or loss. Bodies are said to move freely if they move under the influence of forces inside the system. The kinetic energy gained by the bodies under these forces is equivalent to the loss of potential energy of the system. This indicates that if we let bodies move freely in a conservative force field they always try to attain a configuration with minimal potential. This characteristic of dynamics of bodies is the key for the movements of scans in our approach. We define the fitness measure of a global map by the potential energy of the system defining an underlying gradient field. FFS is an iterative simulation, and since we reset the kinetic energy in the system to zero after each iteration, we minimize the total energy, i.e. the potential energy.

The following section proceeds to define the potential and force field. It will explain the potential update and influence of perceptual properties of the scans as they move.

The simulation of dynamics of bodies is explained in III-D.

B. Potential & force field

For simplicity and without loss of generality, we define the potential and gradient field of our system first assuming a unit mass for points on the rigid bodies. The potential $\mathcal{P}$ at a point $p_1 = (X, Y)$ on a rigid body in a global configuration of n scans at time t is defined as

$$\mathcal{P}_t(p_1) = \int_{-\infty}^{\infty} \frac{1}{\sigma_t \sqrt{2\pi}} e^{-\frac{z^2}{2\sigma_t^2}} dz \quad (1)$$

where r is the Euclidean distance of p_1 to all points in the configuration, $r = \{\sqrt{(x-X)^2 + (y-Y)^2} | (x, y) \in P\}$, P =set of all points. σ_t is a parameter that decreases over time, its significance is explained in section III-D. σ_t does not violate any assumptions needed for free moving bodies to converge towards minima. Motivation for choosing this potential, and not the hyperbolic $1/r^2$ model from physics, is to avoid the over representation of close points with influence towards infinity when the distance goes towards zero. The total potential $\mathcal{E}$ of a given configuration RB_t is the sum of potentials at all points. It is given as

$$\mathcal{E}_t(RB_t) = \sum_{p_i \in RB_t} \mathcal{P}_t(p_i) \quad (2)$$

The negative gradient of this potential, gives us the force field ($\mathcal{F}_t$) induced by $p_1 = (X, Y)$ at a point $p_2 = (x, y) \in P$ as

$$\begin{aligned} \mathcal{F}_t(p_1, p_2) &= -\nabla \mathcal{P} \\ &= \begin{bmatrix} -\frac{\partial \mathcal{P}}{\partial x} \\ -\frac{\partial \mathcal{P}}{\partial y} \end{bmatrix} = \begin{bmatrix} \frac{1}{\sigma_t \sqrt{2\pi}} e^{-\frac{r^2}{2\sigma_t^2}} \cdot \frac{X-x}{r} \\ \frac{1}{\sigma_t \sqrt{2\pi}} e^{-\frac{r^2}{2\sigma_t^2}} \cdot \frac{Y-y}{r} \end{bmatrix} = \frac{1}{\sigma_t \sqrt{2\pi}} e^{-\frac{r^2}{2\sigma_t^2}} \cdot \vec{u} \end{aligned} \quad (3)$$

where $\vec{u} = \frac{p_1 - p_2}{\|p_1 - p_2\|}$. The force field depends on the distance, it is radial and hence a gradient field. The force acting on these bodies is the negative gradient of the potential, we therefore

can exploit the free movement of rigid bodies (as explained in section III-A) to minimize the potential. The key step is to define the masses of data points, so that the rigid bodies move in a perceptually consistent manner, see section III-C).

C. Perceptual Masses, the Correspondence between data points

As described in the previous section, the basic idea of our registration method is to use a Gaussian field to define a strength of correspondence between data points. In this chapter we extend the measure to incorporate both spatial proximity and visual similarity of two points.

To define the correspondence, $\mathcal{F}_t$ is extended by two factors, the masses of two points m_1, m_2 , as well as their parallelity $\cos(\angle(p_1, p_2))$.

$$\mathcal{F}_t^c(p_1, p_2) = \mathcal{F}_t(p_1, p_2) m_1 m_2 \cos(\angle(p_1, p_2)) \quad (4)$$

with m_i being the mass assigned to p_i , and the angle $\angle(p_1, p_2)$ being the angle between the *directions of points* p_1, p_2 . The direction of a point is the direction of the underlying line segment. The strength of correspondence is weighted by the mass of each data point and depends on the angle between point directions, i.e. it is 0 for orthogonal directions, 1 for parallel directions.

A major difference to the pure physics simulation is that the mass values assigned to the data points are not assumed to be constant. The mass m_d for a point p is used to compute the force, yet it can be reassigned a different value for the computation of movement of the scan (we are not modeling physics but perception, hence freedom from Newton's laws is given). Steering the mass enables the algorithm to react better to perceptual properties: there is not perceptual reason for an 'important point', e.g. a corner point, assigned a high mass for force computation, to be less mobile than other points during movement computation (caused by its high mass). This observation suggests using different masses during the computation of forces than during the computation of the movement. The current FFS system has built in different techniques to compute the point masses to model different mid level perceptual processes, e.g. regions of interest, corner detection. In the experiments the masses were steered by point density.

The total force acting on each single point of a scan is

$$F_t^c(p_i) = \sum_{p_j \in P \setminus p_i} (F_t^c(p_i, p_j))$$

The effect of forces on rigid bodies can be seen in Fig. 2.

D. Dynamics

The force field simulation algorithm finds a (locally) optimal rigid transformation due to the gradient $\mathcal{F}_t^c$. However, in our setting, a rigid transformation does not allow the single data points to move independently, but each scan is transformed allowing translation and rotation. Hence $\mathcal{F}_t^c$, which acts on the single data points, has to be translated into rigid body




Fig. 2. FFS on four rigid bodies. The significant forces acting among the bodies are shown in double-headed red arrows. Left the initial configuration. Middle after two iterations. Right stable configuration with minimum potential.

movement, as defined by the constraints of rigid body dynamics.

The details are out of this paper's scope, we will give the basic idea. For each rigid body rb_i , the translational and rotational acceleration has to be determined. The translational acceleration $a_T(rb_i)$ of a rigid rb_i is defined by:

$$a_T(rb_i) = \frac{\sum_{p \in rb_i} F_t^c(p)}{\sum_{p \in rb_i} m_p}$$

with m_p =mass of point p .

The rotational acceleration a_R is computed by torque and moment of inertia. Torque and inertia play the role of force and mass respectively, but take into account the distance to the rotational center c_R .

$$\begin{aligned} inertia &= \sum_{p \in rb_i} m_i \|p_i - c_R\|^2 \\ torque &= \sum_{p \in rb_i} \|p_i - c_R\| \times F(p) \end{aligned}$$

a_R is defined as

$$a_R = \frac{torque}{inertia}$$

The rotational center c_R is either defined as the robot's position, or by the center of mass. Experiments show, that in the first iteration steps it is useful to set the rotational center to the center of mass, while in later steps the robot's position is preferable. The first choice enables easier rotation, the second is modeling the actual scan setting more precisely. Hence, the closer the global map is to the solution, the more preferable is the robot's position as rotational center.

With a_T and a_R the transformation $t_k = (x_k, y_k, \theta_k)$ for rigid body rb_k is defined by:

$$(x_k, y_k) = \frac{1}{2} a_T \Delta_t^2 \quad (3)$$

$$\theta_k = \frac{1}{2} a_R \Delta_t^2 \quad (4)$$

Δ_t is the step width of the gradient descent (see below). t_k is a vector in gradient direction, length depending on Δ_t .

With these constraints, the transformation for each rigid body in each iteration is computed by the following steps:

- 1) for each point $p_i \in \mathcal{P}$ compute $F_t^c(p_i)$

- 2) for each rigid body rb_k compute the transformation $t_k = (x_k, y_k, \theta_k)$ using the points $p_i^k \in rb_k$.

We transform each rb_k , using its previous transformation g_t^k and add the current transformation t_k .

$$g_{t+1}^k = g_t^k + t_k$$

As in all gradient descent methods, the determination of the step width Δ_t is crucial. Also, gradient methods imply the danger of being trapped in local minima. We tackle both problems with the determination of step with Δ_t and σ as described in the following section.

E. Cooling Down the Motion: Time Stepping Δ_t and Distance Influence Parameter σ_t

The determination of step width parameter Δ_t in any gradient descent approach is a well known problem. Δ_t chosen too small results in inapplicably slow convergence behavior and is not robust to noise, Δ_t chosen too big might miss the optimum. In FFS, the step width Δ_t is used as a steering parameter of the algorithm in connection with the parameter σ_t , which determines the influence of distance in the correspondence function. We designed Δ_t as exponentially decreasing, σ_t linearly decreasing.

A large Δ_t allows the scans to be massively relocated (shuffled), they overshoot their correct position in the direction of the correspondence gradient. Naturally, a small Δ_t moves the scans less (the amount of replacement is directly proportional to Δ_t^2 , as defined by the laws of movement). We chose the strategy of decreasing Δ_t and σ_t experimentally, having analogies of the cooling behavior of algorithms like simulated annealing in mind. The imprecise, non optimal large Δ_t at the beginning allows the system to possibly escape from local minima. Observe that in contrast to a technique like simulated annealing we cool down a gradient guided process, not a random state change or a random walk technique that would not be applicable in our high dimensional search space. We therefore avoid the problems with a high computational load (high number of iteration steps) that tend to appear in simulated annealing due to unguided selection of the next state.

The parameter σ_t in equation steers the influence of distance in the computation of point correspondences. A large σ_t enhances the relative influence of data correspondences with greater distances, and, since it equalizes this spatial proximity property, favors the influence of spatial similarity. A small σ_t emphasizes local proximity, which is useful if the global map is already close to an optimum.

The effect of cooling is demonstrated in the experiments using the 'Apartment' data set, where FFS escapes a local minimum due to the behavior of Δ_t and σ_t .

IV. COMPUTATIONAL COMPLEXITY

A. Time complexity

Precise computation of force fields requires a quadratic time complexity ($O(m^2)$) in the total number of points (m) of all the rigid bodies in a configuration. Note that m is the number of resampled points, which can be a drastically

reduced number compared to the original data set. Below we show two possible ways of improving the complexity:

Neighborhood approximation: For each point only its local neighborhood must be examined, since the forces between points rapidly decrease with distance. So any neighborhood search algorithm like using k -d trees [4] would reduce the complexity to $O(m \log m)$. We further can improve by using the line segments of the bodies as objects and use bounding box overlaps to find neighboring points. [21, 36] gives an $O(d \log^2 d + s)$ where d is the number of segments and s is the number of line segments in pair-wise overlaps of the bounding boxes of the line segments. Observe that $d \ll m$. [13] reduces the complexity further down to $O(d + s)$ by exploiting temporal and geometric coherence. A system of bodies is said to be coherent if the relative changes in the configurations are small. They can achieve this speed up because they can cache the sorted lists of neighboring objects and hence sort the neighbors in an expected time of $O(d)$ as in [3, 35].

Fast Gauss Transform (FGT): Greengard and Strain introduced FGT [20] which is in turn based on Fast Multipole Methods introduced for high speed simulation of particle dynamics in potential fields [19]. The main advantage with FGT is that the force field can be computed in $O(m)$ time thus making it linear in n . The constant term depends only on the precision required in computation of the field. Details on choosing the constant can be found in [20]. The main trick is to compute the force field using divide and conquer strategy and exploiting Hermite and Taylor expansions of Eq. (3). FGT has been successfully applied, with several improvements as well, in several applications since it was first introduced [15, 2].

Our current implementation uses the neighborhood approximation using the bounding box technique as in [21, 36].

V. EXPERIMENTAL RESULTS

In this section we present the results of applying our approach to three different data sets, collected using laser range scanners. For each data set the initial configuration, the configuration after a few iterations of FFS and the final configuration is shown. To indicate the convergence speed and convergence behavior, the plot of the potential (fitness measure) during the iteration is given speed. The scan poses are shown as ' x' ' marks. For the experiments we preprocessed the data with a line fitting algorithm to model the data set by line segments, using a robust line-fitting algorithm as described in [24]. This step is followed by equidistant resampling. Line fitting is used to represent the data with equal density, as well as to add geometric information to each data point: each data point is assigned a point-direction, which is the direction of the underlying line segment. Since the algorithm in [24] adapts even to small linear structures in the environment, it results in a robust description of directions even in non-structured environments like outdoor or search and rescue scenarios. Comparison experiments without using line segments showed the drastic improvement in run time as well as in alignment performance when line segments are added. However, the system was able to converge (slower and to less plausible local minima) without line segments. In the experiments, the

masses were steered by point density, motivated by the simple assumption that regions with lesser density are more likely to represent noise. This is true for the NIST Data Set I, bottom left. In this case steering the masses by point density improved the results and can be seen as proof of concept. More sophisticated ways to influence the masses are topic of future research.

A. NIST Data Set I

Figures 3, 4 and 7 refer to this experiment. The NIST disaster data set used in this experiment simulates a typical data set of multi robot mapping in rescue scenarios. It is especially complicated, as it matches the complicated constraints imposed by these settings, which contains poorly estimated pre alignment, no landmarks and very little overlap. The data contains 4 scans taken at 16 positions (total 64 scans), the 4 scans (N,E,S,W) have an overlap of 5 degrees only, the positions differ by 2m. The resulting little overlap made it impossible for sequential approaches to work properly. To compare our approach with a state of the art implementation of the simultaneous mapping approach in Lu/Milios style, we conducted experiments with the 3D Slam system [10]. Both systems behave similarly in terms of convergence speed and accuracy.




Fig. 3. Left, initial configuration of NIST's disaster data (64 scans). Right, after 15 iterations of FFS.




Fig. 4. Left, final map with our approach. Right, the final map obtained by the Lu & Milios technique as reported in [10]. The systems lead to results of comparable quality.

B. Apartment Data Set

See figures 5, 6 and 7 for this experiment. This is the IROS 2006 test data set, consisting of 2000 scans from which we selected every 10th scan. We introduced this experiment to show the ability to escape local minima. We used a pre alignment that shows a clear loop closing error. Initializing the step width Δ_t and σ_t to relatively large values helped to realign the data set and to fix the error. Figure 7 shows the iteration vs. potential curve of experiment 1 (NIST) and

2 (Apartment). The non monotonically decreasing behavior of the second potential curve indicates the escape from the local minimum. In both cases, FFS needed about 30 iterations before a stable state was reached.




Fig. 5. Left, initial configuration of the Apartment data (91 scans). Encircled: a huge alignment error due to a loop closing effect. Right, after 5 iterations of FFS: a large step parameter Δ_t leads to a 're-shuffling' of segments. This is needed to escape the local optimum as shown in the left configuration. Compare also the potential curve in 7




Fig. 6. Left, After 15 iterations. Right, configuration with minimum potential. The encircled regions in this and Fig. 5 show how a typical loop closing problem can be solved using our approach.




Fig. 7. Left, potential vs. iterations of FFS for disaster data. Right, potential for Apartment data set. The potential (encircled) of the apartment data is not monotonically decreasing, indicating a possible escape from a local minimum

C. NIST Maze Data

This data set consists of 16 scans with similar structures, a typical indoor environment, yet again scanned with minimal overlap. See figure 8 and 7 for this experiment.




Fig. 8. Left, initial configuration of NIST's maze data (16 scans). Right, after 5 iterations of FFS.




Fig. 9. Left, final map obtained with FFS. Right, the potential vs. iterations plot.

VI. CONCLUSIONS AND FUTURE WORK

We presented a gradient descent based simultaneous multi-scan alignment, based on physical simulation of force fields. The framework is designed to be extended to incorporate higher-level cognitive features of scans for alignment. The future work involves extension of the approach for general purpose registration problems (e.g. registration of bio-medical images), object recognition and 3D mapping.

ACKNOWLEDGEMENT

We thank Andreas Nüchter from University of Osnabrück, Germany, for his helpful comments and the code for the experiments on 3D-Lu/Milios-SLAM. We thank Raj Madhavan of National Institute of Standards and Technology (NIST) for providing us with the data sets.

REFERENCES

- [1] Nagesh Adluru, Longin Jan Latecki, Rolf Lakaemper, and Raj Madhavan. Robot mapping for rescue robots. In *Proc. of the IEEE Int. Workshop on Safety, Security and Rescue Robotics (SSRR)*, Gaithersburg, Maryland, USA, August 2006.
- [2] Venkat R. Ayyagari, Faysal Boughorbel, Andreas Koschan, and Mongi A. Abidi. A new method for automatic 3d face registration. In *CVPR '05: Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Workshops*, page 119, Washington, DC, USA, 2005. IEEE Computer Society.
- [3] David Baraff. Analytical methods for dynamic simulation of non-penetrating rigid bodies. *Computer Graphics*, 23(3):223–232, 1989.
- [4] Jon Louis Bentley. K-d trees for semidynamic point sets. In *SCG '90: Proceedings of the sixth annual symposium on Computational geometry*, pages 187–197, New York, NY, USA, 1990. ACM Press.
- [5] P.J. Besl and N.D. McKay. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.
- [6] Peter Biber and Wolfgang Strasser. The normal distributions transform: A new approach to laser scan matching. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2003.
- [7] Peter Biber and Wolfgang Strasser. nscan-matching: simultaneous matching of multiple scans and application to slam. In *Robotics and Automation. ICRA Proceedings IEEE International Conference on*, May 2006.
- [8] A. Birk and S. Carpin. Merging occupancy grid maps from multiple robots. In *Proceedings of the IEEE*, volume 94, July 2006.
- [9] Andreas Birk. Learning geometric concepts with an evolutionary algorithm. In *Proc. of The Fifth Annual Conference on Evolutionary Programming*. The MIT Press, Cambridge, 1996.
- [10] D. Borrmann, J. Elseberg, K. Lingemann, A. Nüchter, and J. Hertzberg. The Extension of Lu and Milios Style SLAM to 6 Degree of Freedom. In *IROS 2007, (submitted)*, 2007.
- [11] F. Boughorbel, A. Koschan, B. Abidi, and M. Abidi. Gaussian fields: a new criterion for 3d rigid registration. *Pattern Recognition*, 37, July 2004.
- [12] Yang Chen and Gérard Medioni. Object modelling by registration of multiple range images. *Image Vision Comput.*, 10(3), 1992.
- [13] Jonathan D. Cohen, Ming C. Lin, Dinesh Manocha, and Madhav K. Ponamgi. I-COLLIDE: An interactive and exact collision detection system for large-scale environments. In *Symposium on Interactive 3D Graphics*, pages 189–196, 218, 1995.
- [14] David W. Eggert, Andrew W. Fitzgibbon, and Robert B. Fisher. Simultaneous registration of multiple range views for use in reverse engineering of CAD models. *Computer Vision and Image Understanding: CVIU*, 69(3):253–272, 1998.
- [15] Ahmed Elgammal, Ramani Duraiswami, and Larry S. Davis. Efficient kernel density estimation using the fast gauss transform with applications to color modeling and tracking. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25(11):1499–1504, 2003.
- [16] Richard Feynman. *The Feynman Lectures on Physics: Volume 3*, volume 3 of *The Feynman Lectures on Physics*. Addison-Wesley, Boston, 1963.
- [17] Udo Frese. Treemap: An $o(\log n)$ algorithm for indoor simultaneous localization and mapping. *Auton. Robots*, 21(2):103–122, 2006.
- [18] Udo Frese, Per Larsson, and Tom Duckett. A multilevel relaxation algorithm for simultaneous localization and mapping. *Robotics, IEEE Transactions on Robotics and Automation*, 21, April 2005.
- [19] L. Greengard and V. Rokhlin. A Fast Algorithm for Particle Simulations. *Journal of Computational Physics*, 73:325–348, 1987.
- [20] Leslie Greengard and John Strain. The fast gauss transform. *SIAM J. Sci. Stat. Comput.*, 12(1):79–94, 1991.
- [21] J.E. Hopcroft, J.T. Schwartz, and M. Sharir. Efficient detection of intersections among spheres. *International Journal of Robotics Research*, 2(4), 1983.
- [22] Andrei C. Jalba, Michael H.F. Wilkinson, and Jos B.T.M. Roerdink. Cpm: A deformable model for shape recovery and segmentation based on charged particles. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(10):1320–1335, 2004.
- [23] B. Jensen and R. Siegwart. Scan alignment with probabilistic distance metric. In *Intelligent Robots and Systems, 2004. (IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on*, volume 3, pages 2191–2196, Oct 2004.
- [24] L. J. Latecki and R. Lakaemper. Polygonal approximation of laser range data based on perceptual grouping and em. In *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2006.
- [25] F. Lu and E. Milios. Robot pose estimation in unknown environments by matching 2D range scans. In *CVPR94*, pages 935–938, 1994.
- [26] F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping, 1997.
- [27] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM: A factored solution to the simultaneous localization and mapping problem. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, Edmonton, Canada, 2002. AAAI.
- [28] L. Montesano, J. Miguez, and L. Montano. Probabilistic scan matching for motion estimation in unstructured environments. In *Intelligent Robots and Systems, 2005. (IROS 2005). 2005 IEEE/RSJ International Conference on*, pages 3499–3504, Aug 2005.
- [29] Emmy Noether. Invariant variation problems, m. a. t. avels english translation of invariante variationsprobleme, nachr. d. könig. gesellsch. d. wiss. zu göttingen, math-phys. klasse, 235257 (1918), which originally appeared in transport theory and statistical physics, 1 (3), 183207 (1971)., March 2005.
- [30] Edwin Olson, John Leonard, and Seth Teller. Fast iterative optimization of pose graphs with poor initial estimates. In *Proceedings of ICRA*, pages 2262–2269, 2006.
- [31] C. Robertson and R. Fisher. Parallel evolutionary registration of range data, 2002.
- [32] Szymon Rusinkiewicz, Benedict Brown, and Michael Kazhdan. 3d scan matching and registration, *ICCV Short Course*, 2005.
- [33] Szymon Rusinkiewicz and Marc Levoy. Efficient variants of the ICP algorithm. In *Proceedings of the Third Intl. Conf. on 3D Digital Imaging and Modeling*, 2001.
- [34] Thrun S., Koller D., Ghahramani Z., Durrant-Whyte H., and Ng A.Y. Simultaneous mapping and localization with sparse extended information filters. In J.-D. Boissonnat, J. Burdick, K. Goldberg, and S. Hutchinson, editors, *Proceedings of the Fifth International Workshop on Algorithmic Foundations of Robotics*, Nice, France, 2002. Forthcoming.
- [35] M. I. Shamos. Geometric intersection problems. In *Proc. 17th IEEE Annu. Symp. Found. Comput. Sci.*, pages 208–215, Oct. 1976.
- [36] Hans-Werner Six and Derick Wood. Counting and reporting intersections of -ranges. *IEEE Trans. Computers*, 31(3):181–187, 1982.
- [37] C. Xu and J. Prince. Snakes, shapes, and gradient vector flow. *IEEE Transactions on Image Processing*, pages 359–369, March 1998.
- [38] Ronghua Yang, Majid Mirmehdi, and Xianghua Xie. A charged active contour based on electrostatics. In *Advanced Concepts for Intelligent Vision Systems, ACIVS*, pages 173–184. Springer-Verlag LNCS 4179, September 2006.