

INTERACTIVE MAPPING USING RANGE SENSOR DATA UNDER LOCALIZATION UNCERTAINTY*

Pedro Vieira, Rodrigo Ventura

Abstract:

Several methods have been proposed in the literature to address the problem of automatic mapping by a robot using range scan data, under localization uncertainty. Most scan matching methods rely on the minimization of the matching error among individual range scans. However, uncertainty in sensor data often leads to erroneous matching, hard to cope with in a purely automatic approach. This paper proposes a semi-automatic approach, denoted interactive mapping, involving a human operator in the process of detecting and correcting erroneous matches. Instead of allowing the operator complete freedom in correcting the matching in a frame by frame basis, the proposed method facilitates the adjustment along the directions with more ambiguity, while constraining the others. Experimental results using LIDAR data are presented to validate empirically the approach, together with a preliminary user study to evaluate the benefits of the approach.

Keywords: 2D Mapping, Interactive Interface, ICP, Geometric Constraint Analysis, Iterative Methods

1. Introduction

The problem of automatic environment mapping by a robot has been an active field of research for a long time (see [1] for a review). Methods vary both in terms of sensors used (e.g., sonars [2], LIDAR [3], vision [4], and more recently the Kinect [5], [9]), and in methodologies (e.g., probabilistic [6], scan matching [3]). However, most of these methods are prone to local minima, originated, for instance, by ambiguity or by locally periodic patterns in the environment. Moreover, the loop closure problem remains a not completely solved problem [7]. Approaches aiming at global consistency [3] have been proposed, however, they are often computationally complex, and yet prone to error when faced with missing data.

In this paper we propose an alternative approach where we consider the human-in-the-loop of the scan matching process. In particular, the user is asked to interact with the matching process, by adjusting the match of individual pairs of scans. This adjustment is however constrained by favouring adjustments along the directions of greater ambiguity. Take for instance a case of pairs of identical scans taken from a homogeneous corridor (Fig. 1): the system favours the adjustment of the scans along the corridor, while disfavouring movements orthogonal to the corridor.

The proposed method is based on a graphical user interface (GUI), where the user interacts with the system using a common computer interface (mouse and keyboard).

*This work was supported by the FCT projects [PEst-OE/EEI/LA0009/2011] and [PTDC/EIA-CCO/113257/2009].

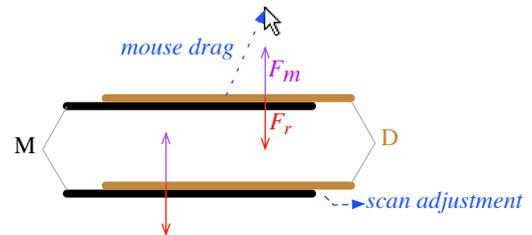


Fig. 1: Two corridors being align. The scan D is adjusted such that the reaction force F_r , computed from the cost function gradient, balances the force F_m imposed by the mouse drag.

Consider a pair of range scans denoted M (for model) and D (for data). Fig. 1 illustrates the situation for the corridor example. When the user drags one of the range scans, say scan D, using the mouse, a corresponding virtual force F_m is produced (proportional to the drag vector). Opposing it, a reaction force F_r is computed based in the match between scans M and D. Considering the scan matching cost function as a potential function, the symmetric of its gradient with respect to the shift of scan D corresponds to the reaction force F_r . This reaction force “attracts” scan D towards M, along the direction of less ambiguity. Scan D is then moved such that both forces become balanced, $F_m + F_r = 0$. In the corridor example of Fig. 1, the reaction force is (mostly) orthogonal to the corridor axis.

Although the method has been implemented for 2D mapping, it can be easily extended to 3D. On the one hand, the matching cost function is applicable to 3D, and on the other, 2D mouse movements can be mapped into 3D forces, depending on the map viewpoint in the GUI, for instance. The work presented here is a preliminary study towards a full 3D implementation.

Such method can be used to make accurate maps, that can then be used, for instance, in search and rescue missions. This missions can vary from searching for victims within risky areas, to analysing certain areas for better mission planning. The availability of a map of these areas can improve dramatically the efficiency and effectiveness of these operations.

Little related work can be found in the literature concerning interactive mapping. There is however some work on the interactive integration of semantic information in maps. Systems which have explored this topic include SemanticSLAM [10] and Interactive SLAM [11], both prototypes not developed beyond their early stages.

This paper is organized as follows. In Section 2 the

method for interactive mapping is proposed, followed by Section 3 showing the results, together with a small user study. Finally, Section 4 draws some conclusions and discuss possible future work directions.

2. Interactive Alignment

Consider two scans in the 2D space, $M = \{\mathbf{m}_k\}$ and $D = \{\mathbf{d}_k\}$, for $\mathbf{m}_k, \mathbf{d}_k \in \mathbb{R}^2$, and an initial rigid transformation (R, \mathbf{t}) which align D with M , where R is a rotation matrix parametrized by an angle θ :

$$R(\theta) = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}, \quad (1)$$

and \mathbf{t} is a translation vector. Although this transformation can be initialized with a default value (e.g. $R|_{\theta=0} = I_{2 \times 2}$ and $\mathbf{t} = (0, 0)^\top$), we start by using the Iterative Closest Points (ICP) algorithm (see [8]) to get the best initial transformation automatically. By applying this transformation to D , a transformed scan, D' , is obtained. The two scans (M and D') are then presented in a viewer (GUI) for alignment analysis. If the alignment isn't well done, the user interacts with the viewer using common computer mouse and keyboard, and apply either translations or rotations to D' , in order to correct it. These actions are carried out separately using a designated key to choose which mode to use.

Scans are aligned by balancing a virtual force originated by a mouse drag, henceforth called mouse force. This drag is defined by two points: \mathbf{p}_o , the mouse pointer position where the user clicked on the mouse, and \mathbf{p}_f , the current pointer position. We consider the mouse force to be proportional to the difference between \mathbf{p}_f , and the initial point \mathbf{p}_o transformed by the alignment under way. In the general case, we assume D' rotated by θ with respect to a center of mass \mathbf{c} , and then translated by \mathbf{t} . Then, point \mathbf{p}_o is transformed into $\mathbf{p}_o' = R(\theta)(\mathbf{p}_o - \mathbf{c}) + \mathbf{c} + \mathbf{t}$. Now, we define a potential function J_m that grows with the distance between points \mathbf{p}_o' and \mathbf{p}_f :

$$J_m = \frac{1}{2} \|\mathbf{p}_f - \mathbf{p}_o'\|^2. \quad (2)$$

By taking the gradient of this potential, with respect to either translation \mathbf{t} and rotation R , one obtains the virtual forces and torques induced by the mouse drag.

The potential function that minimizes the distances between corresponding points of the two point clouds, and is responsible for creating the reaction force to the mouse drag, is given by:

$$J_r = \frac{1}{2} \sum_{k=1}^N \|\mathbf{m}_k - [R(\theta)(\mathbf{d}_k - \mathbf{c}) + \mathbf{c} + \mathbf{t}]\|^2, \quad (3)$$

where $\{\mathbf{m}_k\}$ and $\{\mathbf{d}_k\}$ are pairs of closest points from M and D' respectively, and N is the number of these pairs. This function is similar to the cost function used in ICP. The closest points are computed in the same fashion as in ICP (e.g. using Kd-trees), and only the pairs sufficiently close are considered.

2.1. Translations

In this mode, the alignment consists in a translation by \mathbf{t} . Thus, $R(\theta)$ is the identity ($\theta = 0$), and the potential function (2) is simplified. The mouse force \mathbf{F}_m is computed from the gradient of the cost function (2) with respect to the translation \mathbf{t} and evaluated at $\theta = 0$:

$$\begin{aligned} \mathbf{F}_m &= -k_m \nabla_{\mathbf{t}} J_m|_{\theta=0} \\ &= k_m (\mathbf{p}_f - \mathbf{p}_o - \mathbf{t}), \end{aligned} \quad (4)$$

where k_m is the proportionality constant. Opposing this force, a reaction force \mathbf{F}_r is computed from the gradient of the cost function (3) with respect to the translation \mathbf{t} , and evaluated at $\theta = 0$:

$$\begin{aligned} \mathbf{F}_r &= -k_r \nabla_{\mathbf{t}} J_r|_{\theta=0} \\ &= k_r \sum_{k=1}^N [\mathbf{m}_k - (\mathbf{d}_k + \mathbf{t})], \end{aligned} \quad (5)$$

where k_r is the proportionality constant. Fig. 2 illustrates the forces involved.

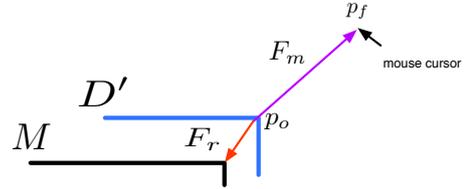


Fig. 2: Geometry of the force balance for translations before alignment ($\mathbf{t} = 0$).

To find a scan adjustment that balances the mouse and the reaction forces, translations are iteratively performed, since each time scan D' moves, the correspondences among points may change. So, for each iteration, a translation \mathbf{t} is computed by solving the equation

$$\mathbf{F}_m + \mathbf{F}_r = 0, \quad (6)$$

with respect to \mathbf{t} . This equation has the following closed form solution:

$$\mathbf{t} = \frac{k_m (\mathbf{p}_f - \mathbf{p}_o) + k_r \sum_{k=1}^N (\mathbf{m}_k - \mathbf{d}_k)}{k_m + N k_r}. \quad (7)$$

The scan adjustment results from the following algorithm:

- 1) Compute translation \mathbf{t} according to (7);
- 2) Compute the new correspondences $\{\mathbf{m}_k\}$ and $\{\mathbf{d}_k\}$ from scans M and D' ;
- 3) Unless the correspondences are the same, go to step 1).

The obtained $\mathbf{t} = [t_x \ t_y]^T$ corresponds to the homogeneous transformation:

$$T_t = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix}. \quad (8)$$

2.2. Rotations

Rotation mode comprises a rotation of scan D' , with respect to the center of mass \mathbf{c} , by an angle θ (Fig. 3). The center of mass of a scan is set to its centroid.

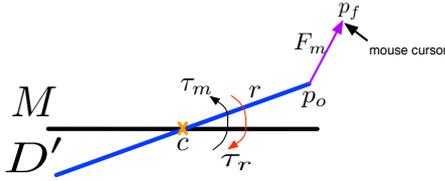


Fig. 3: Torque produced by force originated from a mouse drag prior to alignment ($\theta = 0$).

When the user clicks on the scan D' and attempts to drag it, a force \mathbf{F}_m is created using (2), for $\mathbf{t} = 0$. However, unlike translations, the balance is formulated here in terms of virtual torques. The mouse torque τ_m is obtained by computing the gradient of the cost function (2) with respect to θ , and evaluated at $\mathbf{t} = 0$:

$$\tau_m = -k_m \nabla_{\theta} J_m|_{\mathbf{t}=0}. \quad (9)$$

The opposing torque τ_r is obtained by computing the gradient of the cost function (3) with respect to θ , and evaluated at $\mathbf{t} = 0$:

$$\tau_r = -k_r \nabla_{\theta} J_r|_{\mathbf{t}=0}. \quad (10)$$

Both gradients can be computed using the chain rule:

$$\nabla_{\theta} J_m|_{\mathbf{t}=0} = \text{tr} \left[(d_m)^T d_R \right], \quad (11)$$

$$\nabla_{\theta} J_r|_{\mathbf{t}=0} = \text{tr} \left[(d_r)^T d_R \right]. \quad (12)$$

where

$$d_m = \frac{\partial J_m}{\partial R(\theta)}, \quad (13)$$

$$d_r = \frac{\partial J_r}{\partial R(\theta)}, \quad (14)$$

$$d_R = \frac{\partial R(\theta)}{\partial \theta}. \quad (15)$$

The partial derivative (15) can be computed from the derivative of (1) with respect to θ :

$$d_R = \frac{\partial R(\theta)}{\partial \theta} = A R(\theta), \quad (16)$$

where $A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$.

The partial derivatives (13) and (14) are trivially obtained by computing the derivative of (2) and (3) respectively, with respect to R :

$$d_m = -\mathbf{p}_f' \mathbf{r}_i^T, \quad (17)$$

$$d_r = -\sum_{k=1}^N \mathbf{m}_k' (\mathbf{d}_k')^T, \quad (18)$$

where $\mathbf{r}_i = (\mathbf{p}_o - \mathbf{c})$, $\mathbf{p}_f' = (\mathbf{p}_f - \mathbf{c})$, $\mathbf{d}_k' = (\mathbf{d}_k - \mathbf{c})$ and $\mathbf{m}_k' = (\mathbf{m}_k - \mathbf{c})$.

Both torques are obtained by computing the trace of (11) and (12) and using the partial derivatives (16), (17) and (18):

$$\tau_m = k_m \mathbf{r}_i^T R^T A^T \mathbf{p}_f', \quad (19)$$

$$\tau_r = k_r \sum_{k=1}^N (\mathbf{d}_k')^T R^T A^T \mathbf{m}_k', \quad (20)$$

Note that τ_m can also be interpreted as the norm of the cross product between the arm $\mathbf{r} = R(\theta)(\mathbf{p}_o - \mathbf{c})$ and the mouse force \mathbf{F}_m .

As in the case of translations, the scan D' is iteratively rotated until convergence of the correspondences is reached. In each iteration, the balance of torques

$$\tau_m + \tau_r = 0, \quad (21)$$

has a closed form solution. Using (19) and (20) in (21), we can obtain

$$\tan(\theta) = \frac{k_m \mathbf{r}_i^T A^T \mathbf{p}_f' + k_r \sum_{k=1}^N (\mathbf{d}_k')^T A^T \mathbf{m}_k'}{k_m \mathbf{r}_i^T \mathbf{p}_f' + k_r \sum_{k=1}^N (\mathbf{d}_k')^T \mathbf{m}_k'}. \quad (22)$$

However, this only allows us to compute θ up to a π congruence. To resolve this ambiguity, which is caused by the existence of two solutions π radians apart, one has to determine which one corresponds to a stable solution. This can be easily determined from the sign of the derivative of the total torque $\tau = \tau_m + \tau_r$,

$$\frac{\partial \tau}{\partial \theta} = -k_m \mathbf{r}_i^T R^T(\theta) \mathbf{p}_f' - k_r \sum_{k=1}^N (\mathbf{d}_k')^T R^T(\theta) \mathbf{m}_k'. \quad (23)$$

A solution is stable if and only if the sign of this derivative is negative. A positive derivative implies that a small perturbation in θ will swing the scan π radians towards the other solution. Note that (23) has opposite signs for angles θ and $\theta + \pi$, and therefore there is always a single negative solution.

The scan adjustment follows a similar algorithm as in the case of translations:

- 1) Compute rotation θ according to (22), choosing the solution θ or $\theta + \pi$ with negative derivative (23);
- 2) Compute the new correspondences $\{\mathbf{m}_k\}$ and $\{\mathbf{d}_k\}$ from scans M and D' ;
- 3) Unless the correspondences are the same, go to step 1).

A homogeneous transformation reflecting the rotation performed, is then created using the value θ :

$$T_{\theta} = \begin{bmatrix} \cos(\theta) & -\sin(\theta) & b_x \\ \sin(\theta) & \cos(\theta) & b_y \\ 0 & 0 & 1 \end{bmatrix}, \quad (24)$$

$$\mathbf{b} = \begin{bmatrix} b_x \\ b_y \end{bmatrix} = [I - R(\theta)] \mathbf{c}. \quad (25)$$

Note that translation \mathbf{b} accounts for the fact that the rotation is performed with respect to center \mathbf{c} , rather than to the origin.

3. Results

A dataset was, initially, obtained with a Nomadic Scout robot with a Hokuyo LIDAR onboard. Each scan was associated with a robot location estimated by the robot. This estimation was performed using a simple implementation of EKF together with scan matching using the LIDAR. A partial vector map of the environment was used. As a consequence, the localization algorithm got lost upon entering an office outside of its map (upper part in Fig. 4).

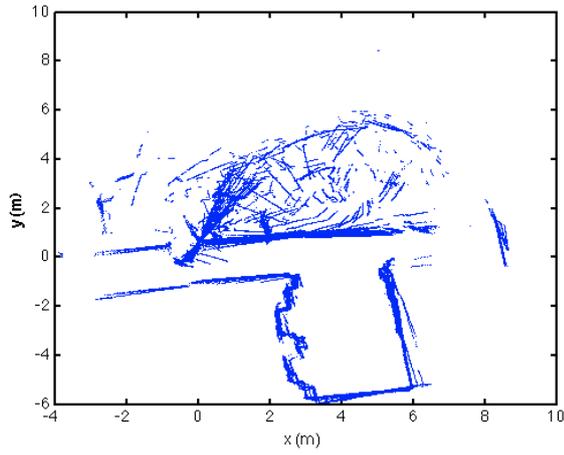


Fig. 4: Initial dataset of a map with 118 scans, obtained using a simple implementation of EKF together with scan matching using the LIDAR

In order to see how the ICP algorithm performs, we run it over the dataset of Fig. 4, obtaining the result shown in Fig. 5. The ICP was able to align most scans, however some of them converged to a local minimum, resulting into a wrong alignment .

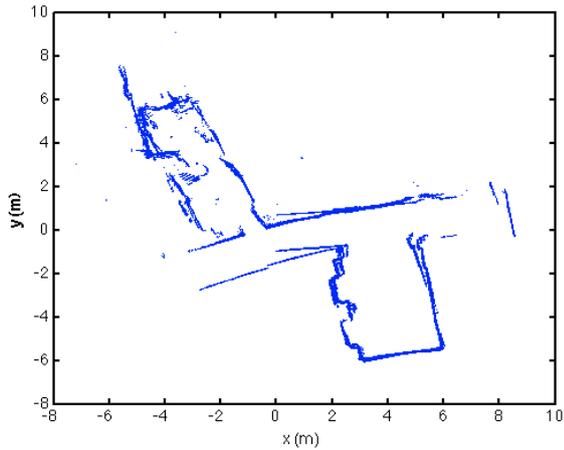


Fig. 5: Map obtained aligning the 118 frames of the initial dataset with ICP.

In Fig. 9 and 10 is shown the alignment of two scans with and without the use of forces for the two interactive modes, translations and rotations respectively. On the left sub-figures of Fig. 9 and 10, we can see that, without the use of forces, the scan movement tends to follow the mouse cursor. On the contrary on the right side, the use of forces favours the adjustment of scan along the corridor, and disfavours the movement orthogonal to it.

To validate the proposed method a user study was conducted by taking 8 users selected among the MSc students from Instituto Superior Técnico (IST). They were asked to correct the alignment of 118 frames, from a map, in two trials: with and without the use of forces. Users had no prior knowledge of the map, and everyone performed the correction of the alignment on the same initial dataset (Fig. 4).

The user performance was measured by (i) the time

they took to finish the task, and by (ii) the value of the cost function, that measures how good the corrections were made. This cost function is the sum of the cost function for all the scan pairs:

$$f = \sum_{i=1}^{n-1} J_i, \quad (26)$$

where n is the number of scans and J_i is the cost function between two scans S_i and S_{i+1} , computed from expression (3).

The forces proportionality constants (k_m and k_r) were empirically adjusted, and in the following experiments, the values of $k_m = 0.1$ and $k_r = 0.001$ were used, together with a closest point threshold of 0.2 meters.

After doing each one of the trials, users obtain knowledge about the initial alignment of each frame which will influence the second experiment, so we mitigate this bias by alternating the order of the experiences between users. Half of them were asked to perform the trial without using forces first, while the other half performed the other one first. Fig. 6 shows an example of a map after the interactive alignment, for a random user.

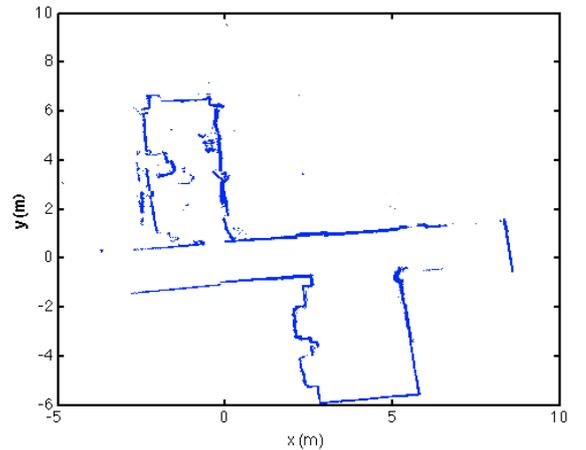


Fig. 6: One of the maps obtained by a user after using the GUI to manually align the scans (compare with Fig. 4).

The results obtained for each user are presented in Fig. 7.

Tab. 1: Average and standard deviation values for time and cost function

	Without Forces		With Forces	
	Time (min)	Cost function Value	Time (min)	Cost function Value
μ	16.27	36.86	10.74	30.60
σ	5.45	8.09	2.77	3.33

As it is visible in Fig. 7 and table 1, users were, in average, able to correct the alignment faster and better using forces than not using them. Paired t-tests over the results have shown that both the average time and average cost functions are lower when forces are used ($p < 0.005$ for time and $p < 0.03$ for cost).

At the end of the trials, users were asked to answer a

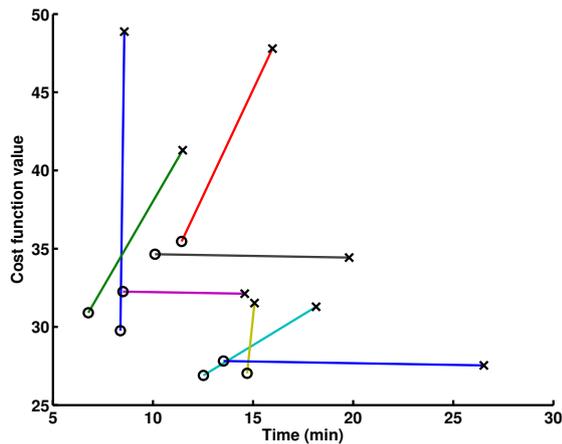


Fig. 7: Results in terms of task completion time and cost function obtained for each user. Each line corresponds to a user, and the “O” and the “X” markers are the results of using or not using forces.

questionnaire, using the well-known Likert scale as a rating criterion (from 1 to 5). The Likert statements were:

- 1) The interface used for the interactive mapping was easy to use
- 2) The use of forces makes the alignment faster
- 3) The use of forces makes the alignment easy

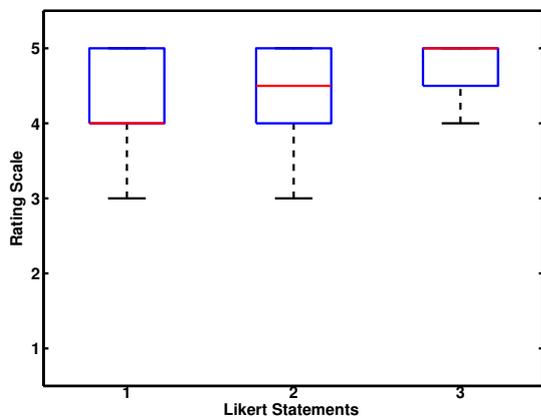


Fig. 8: Box plot of the responses to the questionnaire.

In Fig. 8 is shows the box plot responses given by the users to each statement the questionnaire. Overall, users responded positively to the usage of the GUI, in line with the goal of providing a system to facilitate manual alignment of range scans.

4. Conclusions

In this paper we have proposed interactive method: a method using virtual forces with the purpose of helping users to adjust the alignment of range scan data. The use of forces proved to be a valuable aid in the correction of the alignment, making it faster and more easy to use. A small user evaluation has corroborated this method. As future work we propose extending the concepts proposed in this paper to 3D.

AUTHORS

Pedro Vieira* – Institute for Systems and Robotics, Instituto Superior Técnico, e-mail: pedro.s.vieira@ist.utl.pt

Rodrigo Ventura – Institute for Systems and Robotics, Instituto Superior Técnico, e-mail: rodrigo.ventura@isr.ist.utl.pt

* Corresponding author

References

- [1] S. Thrun, “Exploring artificial intelligence in the new millennium”, Morgan Kaufmann, ch. Robotic mapping: A survey, pp. 1–35, 2002.
- [2] A. Elfes, “Sonar-based real-world mapping and navigation”, *IEEE J. Robot. Automat.*, vol. 3, no. 3, pp. 249–265, June 1987.
- [3] F. Lu and E. Milios, “Globally consistent range scan alignment for environment mapping”, *Autonomous robots*, vol. 4, no. 4, pp. 333–349, 1997.
- [4] D. Murray and C. Jennings, “Stereo vision based mapping and navigation for mobile robots”, in *Proceedings of IEEE International Conference on Robotics and Automation (ICRA’97)*, vol. 2, pp. 1694–1699, April 1997.
- [5] P. Henry, M. Krainin, E. Herbst, X. Ren, and D. Fox, “RGB-D mapping: Using depth cameras for dense 3d modeling of indoor environments”, in *The 12th International Symposium on Experimental Robotics (ISER)*, 2010.
- [6] A. Elfes, “Occupancy grids: A probabilistic framework for robot perception and navigation”. Ph.D. dissertation, Carnegie Mellon University, 1989.
- [7] P. Newman, G. Sibley, M. Smith, M. Cummins, A. Harrison, C. Mei, I. Posner, R. Shade, D. Schroeter, D. Cole, and I. Reid, “Navigating, recognising and describing urban spaces with vision and laser”, *The International Journal of Robotics Research*, vol. 28, no. 11–12, pp. 1406–1433, 2009.
- [8] J. Besl and N. D. McKay, “A method for registration of 3-d shapes”. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(2):239–256, 1992.
- [9] K. Lai, L. Bo, X. Ren, and D. Fox, “A large-scale hierarchical multi-view RGB-D object dataset”, in *Proceedings of IEEE International Conference on Robotics and Automation (ICRA’11)*, 2011, pp. 1817–1824.
- [10] F. Dellaert and D. Brummer, “Semantic slam for collaborative cognitive workspaces”, in *AAAI Fall Symposium Series 2004: Workshop on The Interaction of Cognitive Science and Robotics: From Interfaces to Intelligence*, 2004.
- [11] A. Diosi, G. Taylor, and L. Kleeman, “Interactive SLAM using laser and advanced sonar,” in *Proceedings of IEEE International Conference on Robotics and Automation (ICRA’05)*, 2005, pp. 1103–1108.

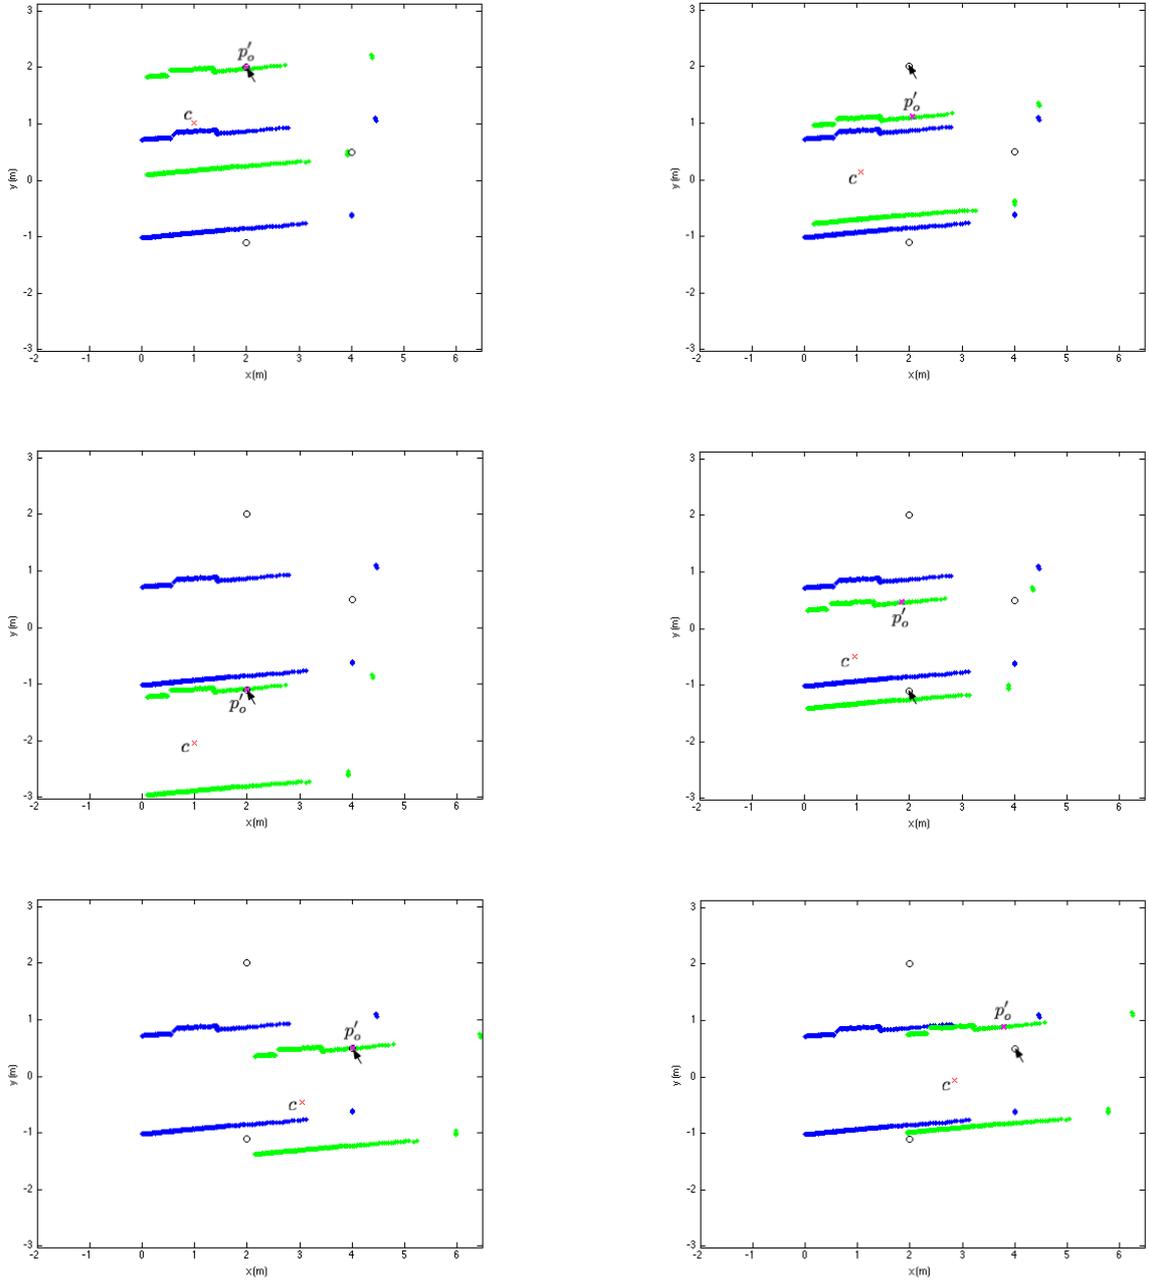


Fig. 9: Two corridors being aligned by applying translations. Left three figures: alignment without forces. Right three figures: alignment with forces. Black arrow is the mouse current position, p_o' is the initial point clicked, c is the green scan center of mass and the back "O" are references points.

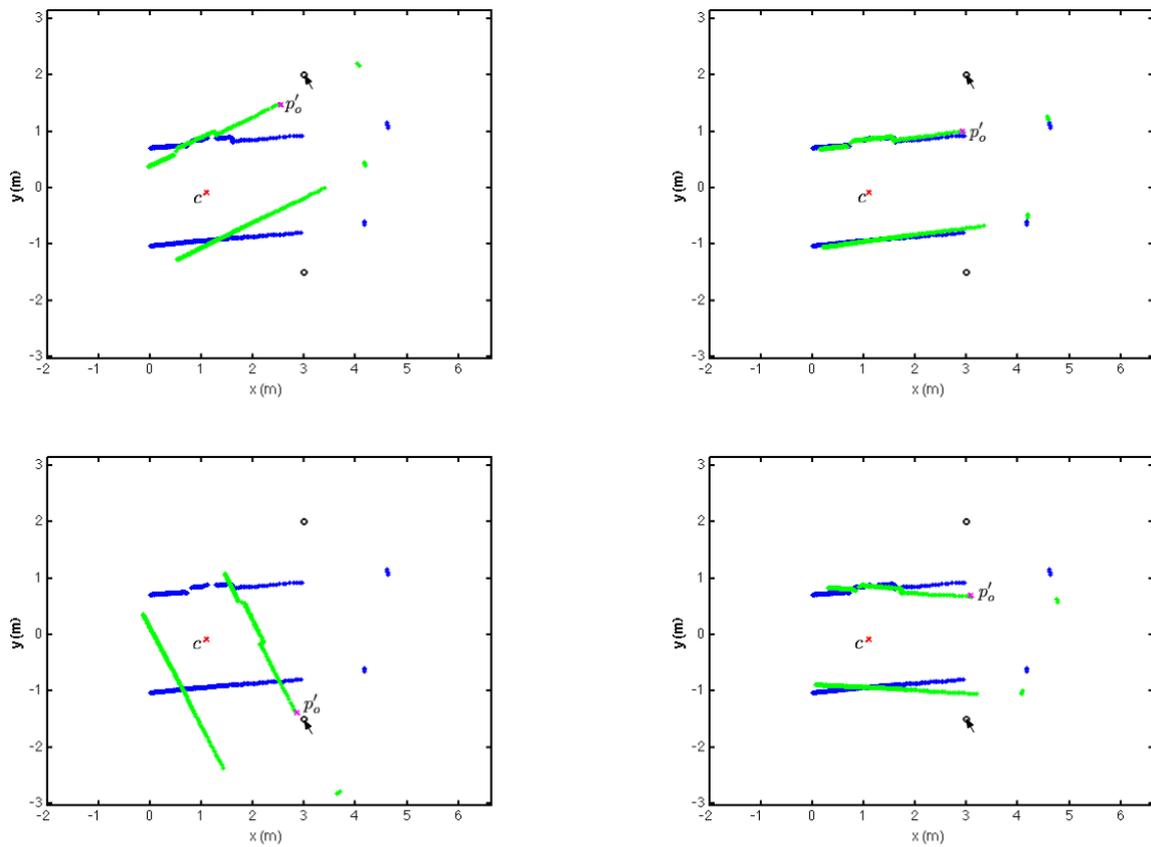


Fig. 10: Two corridors being aligned by applying rotations. Left two figures: alignment without forces. Right two figures: alignment with forces. Black arrow is the mouse current position, p_o' is the initial point clicked, c is the green scan center of mass and the back "O" are references points.