A Non-Rigid Approach to Scan Alignment and Change Detection Using Range Sensor Data

Ralf Kaestner[1], Sebastian Thrun[1], Michael Montemerlo[1], Matt Whalley[2]

[1] Robotics Laboratory
Computer Science Department
Stanford University
Stanford, CA
ralfk@cs.stanford.edu, thrun@stanford.edu, mmde@stanford.edu

[2] Army/NASA Rotorcraft Division
Aeroflightdynamics Directorate (AMRDEC)
US Army Research, Development and Engineering Command,
Ames Research Center, CA
mwhalley@mail.arc.nasa.gov

We present a probabilistic technique for alignment and subsequent change detection using range sensor data. The alignment method is derived from a novel, non-rigid approach to register point clouds induced by pose-related range observations that are particularly erroneous. It allows for high scan estimation errors to be compensated distinctly, whilst considering temporally successive measurements to be correlated. Based on the alignment, changes between data sets are detected using a probabilistic approach that is capable of differentiating between likely and unlikely changes. When applied to observations containing even small differences, it reliably identifies intentionally introduced modifications.

1 Introduction

We provide a unified probabilistic technique for alignment and subsequent change detection using range sensor data. Our work has been motivated by the goal to identify even small changes of the size of a little box by (airborne) vehicle observations. To allow for an exhaustive application of our approach, a terrain shall be sensed as infrequently as possible.

An autonomous Yamaha RMAX helicopter (see Fig. 1) that has been developed within the scope of the NASA Autonomous Rotorcraft Project (ARP), serves as model and experimental platform for our alignment and change detection approach. Tests were performed at the Ames Disaster Assistance and Rescue Team (DART) Collapsed Structure Rescue Training Site (see Fig. 2) located at NASA Ames Research Center in Mountain View, California.

The vehicle's pose is estimated by a Kalman filter integrating GPS position and velocity. The RMAX has been modified to include an avionics payload which, for our purpose, contains a SICK PLS laser range scanning unit that is mounted under the helicopter's nose, pointing straight down at an angle of 90 degrees. The device provides centimeter-accuracy range measurement every one degree over a field-of-view of 180 degrees and at a frequency of 75 Hz. Details on ARP as well as the helicopter's operational issues are to be found in [14].

Fig. 1. The NASA/ARP helicopter equipped with a range measurement device from SICK.

Fig. 2. The DART Collapsed Structure Rescue Training Site located in Mountain View, CA.

2 Related Work

In the past, automatic change detection based on probabilistic techniques has widely been studied in the field of computer vision and image processing. In [1] the history of change detection from remotely sensed digital images is summarized comprehensively. A variety of approaches addressing interesting purposes, like environmental monitoring [3], urban [4], [5], [6] and forest surveillance [7], etc. have further been developed. These studies exclusively base on optical and range imagery acquired by camera or aerial mapping radar and lidar systems.

Recent major breakthroughs in the field of high-precision range sensor technologies led to an increasing availability of inexpensive scanning devices. As a consequence thereof, applying change detection methods to spatially interpreted range sensor data as presented in [8] and [9] has become more and more attractive. However, compared to the major efforts and breakthroughs in the imagery-related theories, these approaches have only been studied to some unsatisfactory extent. Nevertheless, they promise to open up a wide area of potentially superior applications.

In this paper, we are therefore deriving a novel unsupervised technique to evaluate changes in spatial point clouds. Our method deeply relates to those based on the so-called "difference image" as proposed in [1]. In principle, these methods analyze spatial or spatio-temporal distributions of a distance metric between previously aligned reference and sample imagery. But in spite of their relative simplicity and widespread use, they usually exhibit a major weakness: As shown in [2], change detection accuracy strongly correlates with precise registration results that again suffer from the prevailing neglection of building sufficient models.

In contrast to related work that merely makes insufficient effort to precisely register data sets by only few rigid transformations, the theory presented in this paper focuses on enhancing alignment results significantly. This goal is achieved by applying a novel algorithm based on the non-rigid registration approach as presented in [13]. We therefore explicitly derive and consider models of the measuring processes involved in data acquisition.

3 Models

3.1 Helicopter Model

Whenever a range scan is acquired, the helicopter is at a specific pose relative to a global GPS coordinate system. Let us denote the pose by two sets of variables, pertaining to the

x-y-z coordinates of the vehicle and its three Euler angles. Specifically, we assume that at time t, for $t = 1, \ldots, N$, the helicopter's location is given by the variables x_t, y_t, and z_t. We denote the angles at time t by ϕ_t, θ_t and ψ_t. Thus, the state vector of the helicopter at time t is given by $\boldsymbol{\xi_t} = \left(x_t \; y_t \; z_t \; \phi_t \; \theta_t \; \psi_t \right)^T$.

An estimate of the helicopter's pose is provided by the onboard Kalman filter. Specifically, at time t we receive from the EKF a Gaussian pose estimate. The mean of an estimate will be denoted $\boldsymbol{\mu_t}$ and the covariance Σ_t. The sequence of all poses acquired during a single flight of the helicopter is represented by a high-dimensional Gaussian trajectory estimate $\boldsymbol{\xi}$ with mean $\boldsymbol{\mu}$ and covariance Σ. Stated differently, we have $\boldsymbol{\xi} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$. This implies that the corresponding negative logarithm of the pose trajectory distribution is given by

$$-\log pr(\boldsymbol{\xi}) = \text{const.} + \frac{1}{2} \; (\boldsymbol{\xi} - \boldsymbol{\mu})^T \; \Sigma^{-1} \; (\boldsymbol{\xi} - \boldsymbol{\mu}) \tag{1}$$

The constant term in Eq. 1 is the logarithm of the normalizer $|2\pi\Sigma|^{-\frac{1}{2}}$. However, this normalizer does not depend on the target variable $\boldsymbol{\xi}$. Henceforth, its actual value will play no important role, and it can safely be omitted.

3.2 Measurement Model

The helicopter acquires ground data using a range scanner. Each scan consists of $M = 180$ range measurements. The i-th measurement is oriented at angle α_i. Let the actual range measurement be denoted $r_{t,i}$, where t is once again the time index, and i is the index of the measurement beam acquiring the range scan at time t.

Basic geometry suggests that the projection of this measurement into spatial coordinates is now obtained by

$$\mathbf{p_{t,i}}(\boldsymbol{\xi_t}, r_{t,i}) = R_{\phi_t\theta_t\psi_t} \; R_i \; \left(0 \; 0 \; r_{t,i} \right)^T + \left(x_t \; y_t \; z_t \right)^T \tag{2}$$

Here $R_{\phi_t\theta_t\psi_t}$ is the joint rotation matrix that maps the orientation of the helicopter's local coordinate system back into global world coordinates. The vector (x_t, y_t, z_t) represents the helicopter's location in 3-D world coordinates. Finally, the matrix R_i is a rotation matrix that captures the angle α_i of the i-th measurement beam.

In practice, even in static environments each measurement will be corrupted by noise. To model the noise, we assume the existence of a "true" surface point, denoted $\hat{\mathbf{p}}_{\mathbf{t,i}}$. Of course, $\hat{\mathbf{p}}_{\mathbf{t,i}}$ is unknown. However, this surface point induces a true range, denoted $\hat{r}_{t,i}$, which is unknown as well, but it will play an important role in determining whether two scans refer to the same static surface patch in the differencing process. Again, we assume the noise probability distribution elongating along the measurement beam to be Gaussian with mean $\hat{r}_{t,i}$ and variance s. Thus, we define $r_{t,i} \sim \mathcal{N}(\hat{r}_{t,i}, s)$. The corresponding negative logarithm shall then be given by

$$-\log pr(r_{t,i}) = \text{const.} + \frac{1}{2} \; (r_{t,i} - \hat{r}_{t,i}) \; s^{-1} \; (r_{t,i} - \hat{r}_{t,i}) \tag{3}$$

Once again, the constant term in Eq. 3 is the logarithm of the normalizer $|2\pi s|^{-\frac{1}{2}}$ and can be omitted. We remark that the assumed measurement error s plays a critical role in change detection. It characterizes the normal variation we are expecting when measuring the ground surface.

From Eq. 2 we learn that the distribution over probable surface points $\mathbf{p_{t,i}}$ depends on the time-related pose estimation $\boldsymbol{\xi_t}$ and the measurement model for $r_{t,i}$. Both models are represented by Gaussians and we therefore infer that the $\mathbf{p_{t,i}}$ are also distributed

normally. Hence, we define the joint probability $\mathbf{p_{t,i}} \sim \mathcal{N}(\hat{\mathbf{p}}_{\mathbf{t,i}}, Q_{t,i})$ and the negative logarithm

$$-\log pr(\mathbf{p_{t,i}}) = \text{const.} + \frac{1}{2} (\mathbf{p_{t,i}} - \hat{\mathbf{p}}_{\mathbf{t,i}})^T Q_{t,i}^{-1} (\mathbf{p_{t,i}} - \hat{\mathbf{p}}_{\mathbf{t,i}}) \tag{4}$$

Here, the mean $\hat{\mathbf{p}}_{\mathbf{t,i}}$ is once again the "true" surface point we have sensed, and the covariance $Q_{t,i}$ incorporates the helicopter model as stated in Sec. 3.1 along with the projection of our measurement model from Sec. 3.2 into spatial coordinates. The constant part is omitted as usual.

3.3 Practical Considerations

In change detection, the same ground is purposely sensed more than once. To distinguish variable values arising from two independent scanning runs, we introduce a special notation remarking that estimations occurred at different times. Therefore, poses, measurements and induced surface points originating from an earlier run will further be referred to as $\boldsymbol{\xi}_{\mathbf{t'}}$, $\mathbf{r}_{\mathbf{t',i'}}$ and $\mathbf{p}_{\mathbf{t',i'}}$ respectively. Observations acquired during a subsequent pass will remain denoted $\boldsymbol{\xi}_{\mathbf{t}}$, $\mathbf{r}_{\mathbf{t,i}}$ and $\mathbf{p}_{\mathbf{t,i}}$. The noise models will be modified analogously.

4 Scan Alignment

In this section, we are deriving a probabilistic model for the alignment process. Based on that, we will show why classical approaches addressing the problem of registering range scan data that originates from measurements carried out by a helicopter must perform poorly. Hence, we will introduce a novel and superior method that is much more capable of aligning maps of the discussed type.

4.1 Alignment Model

Based on the assumptions made in Sec. 3, we can now formally consider alignment as the problem of maximizing the posterior over locations of points $\mathbf{p} = \{\mathbf{p_{t,i}}\}$ in the world, although they might have been induced by different sensor readings acquired in different runs. Put differently, we seek to find

$$\underset{\boldsymbol{\xi},r}{\operatorname{argmax}}\, pr(\mathbf{p}) \tag{5}$$

where $\boldsymbol{\xi}$ denotes the helicopter's trajectory estimate, whilst $r = \{r_{t,i}\}$ refers to all the measurements acquired during the entire duration of the observation flight.

As common, our approach minimizes the negative likelihood. From Eq. 5, we can therefore derive under independent sensor noise:

$$\underset{\boldsymbol{\xi},r}{\operatorname{argmax}}\, pr(\mathbf{p}) \;=\; \underset{\boldsymbol{\xi},r}{\operatorname{argmin}} -\log pr(\mathbf{p}) \;=\; \underset{\boldsymbol{\xi},r}{\operatorname{argmin}} -\log \prod_{t,i} pr(\mathbf{p_{t,i}}) \tag{6}$$

This expression may then be converted into a simple non-linear least squares problem by applying the negative logarithm stated in Eq. 4. The optimization then resolves to

$$\underset{\boldsymbol{\xi},r}{\operatorname{argmin}} \left\{ \sum_{t,i} \frac{1}{2} (\mathbf{p_{t,i}} - \hat{\mathbf{p}}_{\mathbf{t,i}})^T Q_{t,i}^{-1} (\mathbf{p_{t,i}} - \hat{\mathbf{p}}_{\mathbf{t,i}}) + \text{const.} \right\} \tag{7}$$

A basic assumption our approach does not make is the existence of an explicit model of the environment. This is an important factor, because we are aiming at being able to detect changes by sensing parts of the world not more than twice. Hence, it seems nearly impossible to incorporate enough knowledge to actually learn an adequate probabilistic representation for it. We therefore have to refer to measurements originating from a previous observation run as the reference. Put differently, we define

$$\hat{\mathbf{p}}_{\mathbf{t},\mathbf{i}} = m\left(\mathbf{p}_{\mathbf{t},\mathbf{i}}\right) \tag{8}$$

Here, the function $m\left(\mathbf{p}_{\mathbf{t},\mathbf{i}}\right) \rightarrow \mathbf{p}_{\mathbf{t}',\mathbf{i}'}$ denotes a matching of points acquired at time t and angle α_i with corresponding surface points of the reference, that is points induced by sensor readings of an earlier pass. Hence, the posterior's optimization from Eq. 7 can be restated

$$\underset{\boldsymbol{\xi},r}{\operatorname{argmin}} \sum_{t,i} [\mathbf{p}_{\mathbf{t},\mathbf{i}} - m\left(\mathbf{p}_{\mathbf{t},\mathbf{i}}\right)]^T \; Q_{t,i}^{-1} \; [\mathbf{p}_{\mathbf{t},\mathbf{i}} - m\left(\mathbf{p}_{\mathbf{t},\mathbf{i}}\right)] \tag{9}$$

where constant parts and factors that apply to the whole term have been omitted. This is safe, because they do not depend on the target variables and thus, they will play no role in the overall optimization.

One might have noticed that we have not yet derived the joint covariance matrix $Q_{t,i}$ that appears in our distribution over probable surface points. From Eq. 2 we learn that inferring the elements of $Q_{t,i}$ involves several rotational and translational transformations to be applied to the distributions introduced by our helicopter and our measurement model. However, for the purpose of simplicity, the correct form of $Q_{t,i}$ will play no role in the overall alignment. Its determination is therefore dispensable.

4.2 Aligning Rigid Models

Iterative Closest Point Algorithm

Besl and McKay suggest in [10] to solve the alignment problem by using a special iteration scheme. They prove that, with respect to a given metric, an optimal matching between a reference and a sample point cloud can be achieved by repeatedly optimizing the registration parameters $R_{\Delta\phi\Delta\theta\Delta\psi}$ and $\Delta\mathbf{t}$ such that they minimize the distance metric between the sample and the best matching parts of the reference at a time. Here, $R_{\Delta\phi\Delta\theta\Delta\psi}$, or written briefly ΔR, denotes a joint transformation matrix that rotates all sample points about the three spatial axes, whilst $\Delta\mathbf{t} = \left(\Delta x \; \Delta y \; \Delta z\right)^T$ performs a translation. Put formally, ICP therefore seeks to find

$$\underset{\Delta R,\Delta\mathbf{t}}{\operatorname{argmin}} \sum_{t,i} [\Delta R \; \mathbf{p}_{\mathbf{t},\mathbf{i}} + \Delta\mathbf{t} - m\left(\mathbf{p}_{\mathbf{t},\mathbf{i}}\right)]^T \; Q_{t,i}^{-1} \; [\Delta R \; \mathbf{p}_{\mathbf{t},\mathbf{i}} + \Delta\mathbf{t} - m\left(\mathbf{p}_{\mathbf{t},\mathbf{i}}\right)] \tag{10}$$

Here, the desired metric that matches correspondent points of the clouds is the Mahalanobis distance. However, the classical ICP approach only allows for metrics that assume the constant isotropic and time-invariant case of our posterior. Put differently $Q_{t,i} = q \; I_3$, where I_3 is the 3×3 identity matrix. With respect to the overall minimization, $Q_{t,i}$ is now constant and can be omitted. The optimization therefore reduces to

$$\underset{\Delta R,\Delta\mathbf{t}}{\operatorname{argmin}} \sum_{t,i} [\Delta R \; \mathbf{p}_{\mathbf{t},\mathbf{i}} + \Delta\mathbf{t} - m\left(\mathbf{p}_{\mathbf{t},\mathbf{i}}\right)]^2 \tag{11}$$

Still, the problem is a non-linear least squares problem, because the rotation is a non-linear function of the angles involved.

The ICP algorithm itself is stated in [10] and its convergence is proven. We therefore waive a complete notation of the iteration scheme. The reader should note that ΔR and $\Delta\mathbf{t}$ both incorporate global transformation parameters that align reference and sample as a whole. Thus, both point clouds are considered to be rigid.

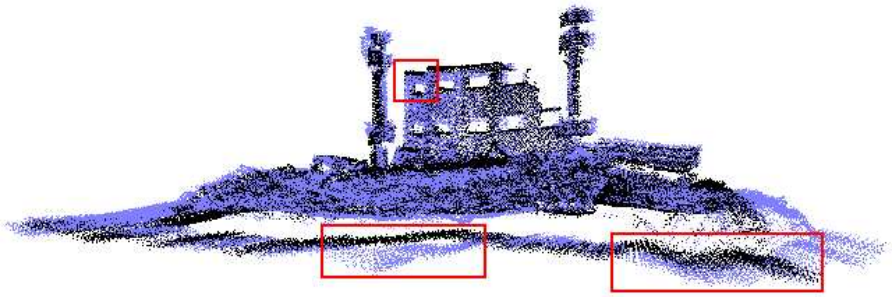

Fig. 3. Registration results obtained by considering reference (light blue) and sample (black) to be rigid. The side view reveals, that there remain huge divergences (red boxes) between numerous parts of the two point clouds.

Linear Optimization Approach

Commonly, it is convenient to linearize the optimization criterion. For that purpose, we simply expect aligning rotations at each step of ICP to be small and approximate $\cos\alpha$ by 1 and $\sin\alpha$ by α. In our case, this is a reasonable assumption, because we seek to register point clouds that are pre-aligned to some extent. Accordingly, $R_{\Delta\phi\Delta\theta\Delta\psi}$ denotes

$$R_{\Delta\phi\Delta\theta\Delta\psi} \;=\; \Delta R \;\approx\; \begin{pmatrix} 1 & -\Delta\psi & \Delta\theta \\ \Delta\psi & 1 & -\Delta\phi \\ -\Delta\theta & \Delta\phi & 1 \end{pmatrix} \tag{12}$$

Now that the optimization problem is linear, it can easily be solved exactly (relatively to the approximations). For this purpose, the partial derivatives of Eq. 11 with respect to the dependents Δx, Δy, Δz, $\Delta\phi$, $\Delta\theta$ and $\Delta\psi$ are set to zero, yielding 6 equations containing the 6 sought variables. This gives us a matrix equation of the form $A\,\mathbf{x} = -\mathbf{b}$, where A is a 6×6 matrix collecting the gradients of the optimization criterion, $\mathbf{b}$ is a 6-dimensional vector incorporating the constant components of the derivatives and $\mathbf{x} = \begin{pmatrix} \Delta x & \Delta y & \Delta z & \Delta\phi & \Delta\theta & \Delta\psi \end{pmatrix}^T$.

Here, the exact determination of A and $\mathbf{b}$ be omitted, because this solely involves basic maths. However, we shall remark the important fact that A is symmetric and positive definite. This is an important property a variety of particularly efficient numerical algorithms take advantage of. For example, $A\,\mathbf{x} = -\mathbf{b}$ may be solved using the Cholesky Decomposition.

Preliminary Results

Fig. 3 shows two point clouds, one being the reference $\mathbf{p_{t',i'}}$, the other being the sample $\mathbf{p_{t,i}}$. Both have been registered using rigid body transformations as suggested by Besl and McKay. The scans inducing the point clouds were acquired independently during two flights of our laser-equipped helicopter. The classical approach obviously performs poorly in that there remain huge divergences between numerous parts of the two data sets.

To gain more insight into the apparent lacks that occur when applying standard ICP to helicopter range observations, we need to investigate the reasons for its failing. By looking at Fig. 3, we find that the errors in our observations are nearly uncorrelated. This implies that the variances and covariances of the pose estimation and the measurement distribution are significant. The attempt to apply rigid body transformations to our sample is obviously not capable of compensating for this kind of errors.

5 Aligning Non-Rigid Models

In Sec. 4.2 we have shown that assuming an alignment model consisting of several individual scans to be rigid leads to results that implicitly induce a vast amount of false differences. Therefore, we now present a novel approach that is capable of dealing with high, randomly distributed estimation errors ICP does not consider.

The key idea is to extent the classical alignment approach by treating the sample we wish to align to a reference model as a non-rigid point cloud. More specifically, each scan is considered as rigid body, whilst it remains subject to local transformations as a whole. Based on that, we are then able to involve dependency assumptions that account for correlations between observation errors.

Extended Optimization Criterion

In a first step we simply allow for scans to be transformed independently instead of assuming a global ΔR and a global $\mathbf{\Delta t}$. For that purpose, we introduce a rotation matrix $\Delta R_t = R_{\Delta\phi_t \Delta\theta_t \Delta\psi_t}$ and a translation vector $\mathbf{\Delta t_t} = \begin{pmatrix} \Delta x_t & \Delta y_t & \Delta z_t \end{pmatrix}^T$. Both are applying to all $\mathbf{p_{t,i}}$ that were induced by readings acquired at time t. We now modify the simplified posterior from Eq. 11 to incorporate the scan-related transformation parameters and henceforth seek to find

$$\underset{\Delta R_t, \mathbf{\Delta t_t}}{\operatorname{argmin}} \sum_{t,i} [\Delta R_t \ \mathbf{p_{t,i}} + \mathbf{\Delta t_t} - m(\mathbf{p_{t,i}})]^2 \tag{13}$$

At first glance, this only seems to be a slight modification, but Eq. 13 essentially differs from Eq. 11 in that we now use a vast sequence of scan-related, local rotation matrices ΔR_t and translation vectors $\mathbf{\Delta t_t}$ instead of a single scan-common, global ΔR and a global $\mathbf{\Delta t}$ to align reference and sample.

Nevertheless, this gives us a linear optimization problem we are able to solve exactly. Setting the partial derivatives with respect to the dependents Δx_t, Δy_t, Δz_t, $\Delta\phi_t$, $\Delta\theta_t$ and $\Delta\psi_t$ for all t to zero, we yield $6N$ equations with $6N$ unknown variables. That is, the dimensionality of our optimization problem grew by the factor N. However, we collect the elements to a linear matrix equation of the form $A^* \ \mathbf{x^*} = -\mathbf{b^*}$, where A^* is once again the matrix containing all the gradients of our optimization criterion and now incorporates $6N \times 6N$ elements. As usual, the $6N$-dimensional vector $\mathbf{b^*}$ accumulates the constant components of the derivatives and $\mathbf{x^*}$ represents all $6N$ sought transformation parameters.

A closer look to the matrix equation reveals that A^* remains symmetric and positive definite. Although its dimensionality increased tremendously, non-zero elements can solely be found along the main diagonal. The reason for that is covered in our model assumption. In particular, we considered all pose estimations and measurements to be independent in terms of noise. The missing gradients to the left and the right of the main diagonal therefore account for missing dependencies in the posterior. Hence, A^* is said to be sparse and the stated problem is usually referred to as a *sparse energy minimization problem*, for which a rich family of efficient solvers exist. Amongst them, the conjugate gradient (CG) algorithm is the most prominent iterative method. Descriptions of CG can be found in contemporary textbooks on optimization [15]. The details of the algorithm are omitted for brevity.

Plausible Dependency Assumptions

The alignment approach stated so far is based on a non-rigid model that assembles rigid scan lines in a non-specific way. Estimation noise is assumed to be independent and consequently, scans are considered as incoherent. However, this is a very weak assumption that

may lead to unexpected registration results. The Kalman filter estimating the helicopter's pose at each step t, for example, emits sequences of state vectors $\boldsymbol{\xi_t}$ that are known to be statistically dependent over time.

To account for correlated observation errors of this type, we therefore assume non-zero covariances between temporally successive observations. Put differently, the idea may be compared to attaching optimization constraints in the form of little springs between scans arising from time t and time $t-1$. We modify the posterior from Eq. 13 accordingly to take advantage of a translational covariance parameter Δq_t and a rotational covariance Δq_r. Henceforth, we seek to find

$$\begin{aligned} &\underset{\Delta R_t, \mathbf{\Delta t_t}}{\operatorname{argmin}} \sum_{t,i} [\Delta R_t \ \mathbf{p_{t,i}} + \mathbf{\Delta t_t} - m\left(\mathbf{p_{t,i}}\right)]^2 \\ &+ \frac{1}{\Delta q_t} \sum_t [\mathbf{\Delta t_t} - \mathbf{t_{\Delta t-1}}]^2 + \frac{1}{\Delta q_r} \sum_t [\mathbf{\Delta r_t} - \mathbf{\Delta r_{t-1}}]^2 \end{aligned} \tag{14}$$

Here, we introduced another vector $\mathbf{\Delta r_t} = \begin{pmatrix} \Delta\phi_t & \Delta\theta_t & \Delta\psi_t \end{pmatrix}^T$ that simply accumulates the corresponding Euler angles from our rotation matrix $\Delta R_t = \Delta R_{\Delta\phi_t \Delta\theta_t \Delta\psi_t}$.

To exactly solve the linear optimization problem stated in Eq. 14, we are once again setting the partial derivatives with respect to the dependents Δx_t, Δy_t, Δz_t, $\Delta\phi_t$, $\Delta\theta_t$ and $\Delta\psi_t$ for all t to zero. Furthermore, we collect the gradients and constants to yield a linear matrix equation of the familiar form and structure. Again, the gradient matrix remains sparse as well as symmetric and positive definite. Hence, applying the CG algorithm to solve the system of linear equations remains a feasible approach.

Online Approach

The approach presented so far constructs a global optimization problem that incorporates knowledge on the complete data sets. However, a lot of applications exist where an online algorithm is a much more favorable solution. We are therefore briefly presenting a slightly modified, iterative approach that incorporates range scans as they occur.

The key idea of our alternative, real-time computable alignment method, which can directly be derived from Eq. 14, additionally considers a translational variance parameter s_t as well as a rotational variance s_r. Thus, the online posterior optimization shall be denoted

$$\begin{aligned} &\underset{\Delta R_t, \mathbf{\Delta t_t}}{\operatorname{argmin}} \sum_{i} [\Delta R_t \ \mathbf{p_{t,i}} + \mathbf{\Delta t_t} - m\left(\mathbf{p_{t,i}}\right)]^2 \\ &+ \frac{1}{\Delta q_t} [\mathbf{\Delta t_t} - \mathbf{t_{\Delta t-1}}]^2 + \frac{1}{\Delta q_r} [\mathbf{\Delta r_t} - \mathbf{\Delta r_{t-1}}]^2 + \frac{1}{s_t} \mathbf{\Delta t_t}^2 + \frac{1}{s_r} \mathbf{\Delta r_t}^2 \end{aligned} \tag{15}$$

and has to be determined for every arising observation at time t.

5.1 Alignment Results

Alignment results are depicted in Fig. 4. Compared to the preliminary results shown in Fig. 3 that were obtained by applying classical ICP, our approach provides a near-perfect registration of reference and sample, even though the helicopter's pose estimation is exceptionally imprecise and noisy.

In order to yield a quantitative evaluation of our alignment results, we assume a normal distribution of all difference vectors between registered sample points and corresponding, that is nearest, reference points. This approach allows for the mean and the

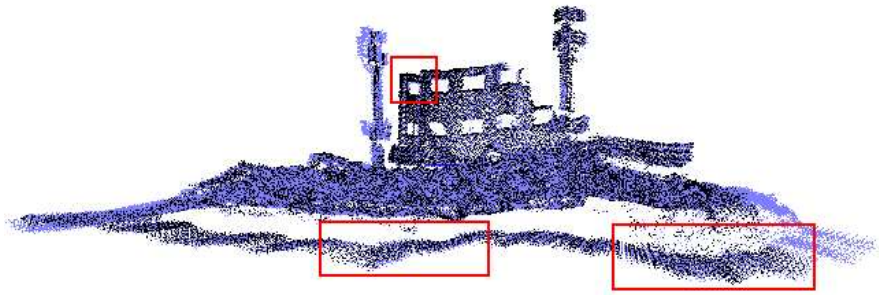

Fig. 4. Registration results obtained by applying the non-rigid approach presented in this paper. Reference (light blue) and sample (black) are aligned near-perfectly.

Alignment/Error Mean		**Std. Deviation**
None	$\begin{pmatrix} 0.017 & 0.032 & -0.063 \end{pmatrix}$	$\begin{pmatrix} 0.117 & 0.129 & 0.212 \end{pmatrix}$
Rigid	$\begin{pmatrix} 0.001 & 0.000 & 0.002 \end{pmatrix}$	$\begin{pmatrix} 0.094 & 0.090 & 0.153 \end{pmatrix}$
Non-Rigid, Offline	$\begin{pmatrix} 0.002 & 0.002 & 0.004 \end{pmatrix}$	$\begin{pmatrix} 0.058 & 0.054 & 0.061 \end{pmatrix}$
Non-Rigid, Online	$\begin{pmatrix} 0.002 & 0.003 & 0.000 \end{pmatrix}$	$\begin{pmatrix} 0.053 & 0.052 & 0.058 \end{pmatrix}$

Table 1. Alignment errors (in meters) obtained by applying the classical registration as well as the approaches presented. Significantly smaller standard deviations of the error distributions achieved by applying the non-rigid methods clearly indicate better alignment results.

standard deviation of the Gaussian alignment error for each of the spatial dimensions to be determined empirically.

Table 1 presents the alignment errors achieved by applying the classical, rigid method as well as the non-rigid offline and the non-rigid online approach proposed in this paper to the helicopter's observations. Obviously, all methods perform fairly well in that they are capable of shifting the distributions' means into the origin. However, the rigid registration does not compensate for the error's variant components, whilst both non-rigid approaches significantly narrow the Gaussians, thus clearly indicating better alignment results.

6 Change Detection

This section of the paper deals with the question of how to actually identify relevant changes in the previously aligned models. It will become apparent that in order to reliably detect changes in naturally erroneous and noisy sensor data, several assumptions are to be made. Therefore a probabilistic approach with the capability of considerably improving the detection results is presented.

6.1 Change Model

The critical question in change detection is whether a surface in the real world has changed. In our context, a change is manifested by the fact that two "true" measurements differ, that is the environment has been altered between two consecutive scans. The probability for this to happen shall be given by the expression

$$pr\left(\hat{\mathbf{p}}_{\mathbf{t},\mathbf{i}} \neq \hat{\mathbf{p}}_{\mathbf{t}',\mathbf{i}'}\right) \tag{16}$$

We introduce a joint probability distribution giving an estimate of changes in the environment. Therefore, we are considering two distinct cases.

Case 1: The world has not changed. Consider the range scan $r_{t,i}$ and its induced point $\mathbf{p_{t,i}}$ along with the acquisition pose $\boldsymbol{\xi}_\mathbf{t}$. Furthermore, let $\bar{\mathbf{p}}_{\mathbf{t'},\mathbf{i'}}$ be the nearest neighbor to $\mathbf{p_{t,i}}$, and $\bar{r}_{t',i'}$ the associated range sensed relative to the same pose. Then our approach gives us $pr\left(r_{t,i} \mid \bar{\mathbf{p}}_{\mathbf{t'},\mathbf{i'}}, \boldsymbol{\xi}_\mathbf{t}, \text{false}\right) \sim \mathcal{N}\left(\bar{r}_{t',i'}, 2s\right)$. The variance of $2s$ accounts for the two measurement noise variables involved in this process, one from each of the scans. We assume independence in measurement noise, hence the variances of both noise variables are simply added.

Case 2: The world has changed. Since we make no assumptions as to how the world changes, the best we can assume is a uniform distribution over scan lengths. Thus, we define $pr\left(r_{t,i} \mid \bar{\mathbf{p}}_{\mathbf{t'},\mathbf{i'}}, \boldsymbol{\xi}_\mathbf{t}, \text{true}\right) \sim \mathcal{U}\left(0, r_{\max}\right)$, with $\mathcal{U}$ denoting the desired uniform distribution.

These two probabilities enable us to arrive at a probabilistic expression to estimate when the world changed. In particular, Bayes rule suggests:

$$\begin{aligned} & pr\left(\text{true} \mid r_{t,i}, \bar{\mathbf{p}}_{\mathbf{t'},\mathbf{i'}}, \boldsymbol{\xi}_\mathbf{t}\right) \\ & = \frac{pr\left(r_{t,i} \mid \bar{\mathbf{p}}_{\mathbf{t'},\mathbf{i'}}, \boldsymbol{\xi}_\mathbf{t}, \text{true}\right) \cdot p\left(\text{true}\right)}{pr\left(r_{t,i} \mid \bar{\mathbf{p}}_{\mathbf{t'},\mathbf{i'}}, \boldsymbol{\xi}_\mathbf{t}, \text{true}\right) \cdot p\left(\text{true}\right) + pr\left(r_{t,i} \mid \bar{\mathbf{p}}_{\mathbf{t'},\mathbf{i'}}, \boldsymbol{\xi}_\mathbf{t}, \text{false}\right) \cdot p\left(\text{false}\right)} \end{aligned} \tag{17}$$

Here $p\left(\text{true}\right)$ is the prior probability of a change, and $p\left(\text{false}\right) = 1 - p\left(\text{true}\right)$ the probability that the world remains static at any given location.

The logarithm of this expression can be approximated using Jensen's inequality:

$$\begin{aligned} & \log pr\left(\text{true} \mid r_{t,i}, \bar{\mathbf{p}}_{\mathbf{t'},\mathbf{i'}}, \boldsymbol{\xi}_\mathbf{t}\right) \\ & = \log[pr\left(r_{t,i} \mid \bar{\mathbf{p}}_{\mathbf{t'},\mathbf{i'}}, \boldsymbol{\xi}_\mathbf{t}, \text{true}\right) \cdot p\left(\text{true}\right)] - \log[pr\left(r_{t,i} \mid \bar{\mathbf{p}}_{\mathbf{t'},\mathbf{i'}}, \boldsymbol{\xi}_\mathbf{t}, \text{true}\right) \cdot p\left(\text{true}\right)] \\ & \quad - \log[pr\left(r_{t,i} \mid \bar{\mathbf{p}}_{\mathbf{t'},\mathbf{i'}}, \boldsymbol{\xi}_\mathbf{t}, \text{false}\right) \cdot p\left(\text{false}\right)] \\ & = - \log[pr\left(r_{t,i} \mid \bar{\mathbf{p}}_{\mathbf{t'},\mathbf{i'}}, \boldsymbol{\xi}_\mathbf{t}, \text{false}\right) \cdot p\left(\text{false}\right)] \\ & = - \log p\left(\text{false}\right) + \log\sqrt{4\pi s} + \frac{1}{2}\left(r_{t,i} - \bar{r}_{t',i'}\right)^T \ \left(2s\right)^{-1} \ \left(r_{t,i} - \bar{r}_{t',i'}\right) \\ & = \text{const.} + \frac{1}{2}\left(r_{t,i} - \bar{r}_{t',i'}\right)^T \ \left(2s\right)^{-1} \ \left(r_{t,i} - \bar{r}_{t',i'}\right) \end{aligned} \tag{18}$$

Therefore, we are now able to determine how probable a likely change is. We compute a simple quadratic distance of the type

$$d_{t,i} = \left(r_{t,i} - \bar{r}_{t',i'}\right)^T \ \left(2s\right)^{-1} \ \left(r_{t,i} - \bar{r}_{t',i'}\right) \tag{19}$$

Assuming that a change with a probability $pr\left(\text{true} \mid r_{t,i}, \bar{\mathbf{p}}_{\mathbf{t'},\mathbf{i'}}, \boldsymbol{\xi}_\mathbf{t}\right) > 0.5$ may be significant, the distance $d_{t,i}$ has to be compared to the following threshold

$$\begin{aligned} d_{t,i} & \overset{?}{>} 2 \log pr\left(\text{true} \mid r_{t,i}, \bar{\mathbf{p}}_{\mathbf{t'},\mathbf{i'}}, \boldsymbol{\xi}_\mathbf{t}\right) - 2 \text{ const.} \\ & = 2 \log 0.5 + 2 \log p\left(\text{false}\right) - 2 \log\sqrt{4\pi s} \end{aligned} \tag{20}$$

If the comparison evaluates to "true", a probable change can be marked as detected.

6.2 Change Detection Results

The change detection results shown in Figs. 5, 7 and 6 were all obtained by preliminarily registering reference and sample according to our alignment approach. A small cubical box with a maximum edge length of about 20 cm has been placed within the area examined by the helicopter.

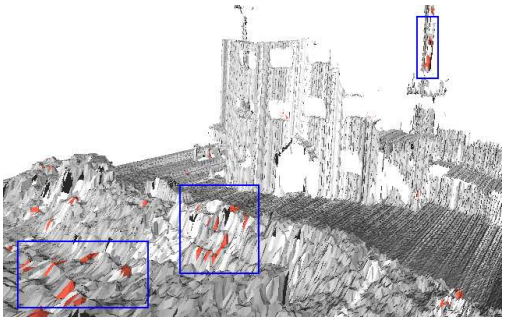

Fig. 5. Change detection results obtained by evaluating the Euclidean distance between reference and sample. Changes (red) are commonly occurring in scattered areas.

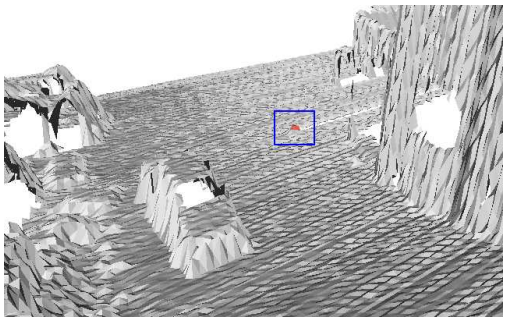

Fig. 6. The intentionally introduced change (a small box with an edge length of about 20 cm) has been detected correctly by our approach.

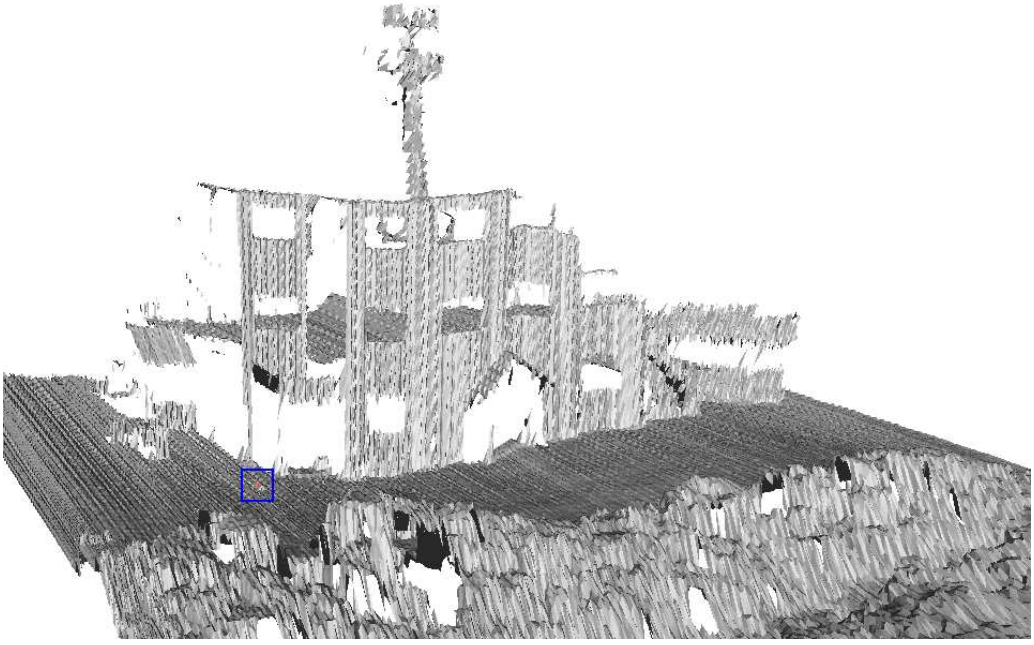

Fig. 7. By applying the probabilistic threshold analysis presented in our approach, all unlikely changes could successfully be extracted from the results. The "true" change is marked.

Fig. 5 illustrates that solely evaluating the Euclidean distances between both point clouds explicitly fails in scattered areas of the environment. The superiority of our probabilistic change detection approach is visualized in Fig. 7, where no false changes were marked within these parts. In Fig. 6 a closer look to our intentionally introduced change reveals that is has been detected correctly after all.

7 Discussion and Future Work

We presented a probabilistic approach for alignment and change detection using range sensor data. The alignment method provides a non-rigid point cloud registration that near-perfectly deals with high estimation errors. Changes between data sets are detected

using a probabilistic threshold that is capable of differentiating between likely and unlikely changes. Even small intentional modifications are detected reliably, whilst errors are excluded.

There is ample opportunity for future research. On the basis of a reliable change detection, changes could be classified. That would allow for a variety of applications in robotics to superiorly deal with dynamic environments. Another opportunity pertains to the learning component of our work: In addition to learning alignment models, it should be possible to learn the actual noise models of the sensors and the pose estimation system. We also suspect that further improvements can be achieved by better involving the exact physical noise characteristics of a sensor.

References

1. Singh, A. 1989. Digital change detection techniques using remotely-sensed data. *International Journal of Remote Sensing*, 10(6):989-1004.
2. Dai, X., Khorram, S. 1998. The effects of image misregistration on the accuracy of remotely sensed change detection. *IEEE Transactions on Geoscience and Remote Sensing*, 36(5):1566-1577.
3. Chavez, P. S., MacKinnon, D. J. 1994. Automatic detection of vegetation changes in the southwestern United States using remotely sensed images. *Photogramm Eng. Remote Sensing*, 60(5):571-582.
4. Carlotto, M. J. 1997. Detection and analysis of change in remotely sensed imagery with application to wide area surveillance. *IEEE Transactions on Image Processing*, 6(1):189-202.
5. Murakami, H., Nakagawa, K., Hasegawa, H., Shibata, T., Iwanami, E. 1999. Change detection of buildings using an airborne laser scanner. *Journal of Photogrammetry and Remote Sensing*, 54(2-3):148-152.
6. Vu, T. T., Matsuoka, M., Yamazaki, F. 2004. Lidar-based change detection of buildings in dense urban areas. *IEEE 200 International Geoscience and Remote Sensing Symposium*, CD-ROM, 4p, 2004.9.
7. Coppin, P. R., Bauer, M. E. 1996. Digital change detection in forest ecosystems with remote sensing imagery. *Remote Sensing Reviews*, 13(3-4):207-234.
8. Hsiao, K. H., Liu, J. K., Yu, M. F., Tseng, Y. H. 1999. Change detection of landslide terrains using ground-based lidar data. *Proceedings of the XXth ISPRS Congress*, Istanbul, Turkey.
9. Habib, A. F., Lee, Y.-R., Morgan, M. 2001. Surface matching and change detection using a modified Hough transformation for robust parameter estimation. *Photogrammetric Record*, 17(98):303-315.
10. Besl, P. and McKay, N. 1992. A method for registration of 3d shapes. *Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239-256.
11. Chen, Y. and Medioni, G. 1992. Object modeling by registration of multiple range images. *Image and Vision Computing*, 10:145-155.
12. Chui, H., Rangarajan, A. 2000. A new point matching algorithm for non-rigid registration. *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*.
13. Haehnel, D., Thrun, S., Burgard, W. 2003. An extension of the ICP algorithm for modeling nonrigid objects with mobile robots. *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, Acapulco, Mexico.
14. Whalley, M., Freed, M., Harris, R., Takahashi, M., Schulein, G., and Howlett, J. 2005. Design, integration, and flight test results for an autonomous surveillance helicopter. *AHS International Specialists Meeting on Unmanned Rotorcraft*, Mesa, AZ.
15. Press, W. H. 1988. *Numerical recipes in C: the art of scientific computing.* Cambridge University Press, Cambridge, New York.