

Scale-invariant registration of monocular endoscopic images to CT-scans for sinus surgery

Darius Burschka^{a,b,c,*}, Ming Li^{a,b,c}, Masaru Ishii^{a,b,c},
Russell H. Taylor^{a,b}, Gregory D. Hager^{a,b}

^a Computational Interaction and Robotics Laboratory, CIRL, The Johns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218, USA

^b Computer Integrated Surgical Systems and Technology, CISST, The Johns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218, USA

^c Department of Otolaryngology, The Johns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218, USA

Available online 11 July 2005

Abstract

In this paper, we present a novel method for intra-operative registration directly from monocular endoscopic images. This technique has the potential to provide a more accurate surface registration at the surgical site than existing methods. It can operate autonomously from as few as two images and can be particularly useful in revision cases where surgical landmarks may be absent. A by-product of video registration is an estimate of the local surface structure of the anatomy, thus providing the opportunity to dynamically update anatomical models as the surgery progresses.

Our approach is based on a previously presented method [Burschka, D., Hager, G.D., 2004. V-GPS (SLAM): – Vision-based inertial system for mobile robots. In: Proceedings of ICRA, 409–415] for reconstruction of a scaled 3D model of the environment from unknown camera motion. We use this scaled reconstruction as input to a PCA-based algorithm that registers the reconstructed data to the CT data and recovers the scale and pose parameters of the camera in the coordinate frame of the CT scan. The result is used in an ICP registration step to refine the registration estimates.

The details of our approach and the experimental results with a phantom of a human skull and a head of a pig cadaver are presented in this paper.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Image registration; Monocular reconstruction; Vision-based SLAM; Sinus surgery

1. Introduction

Sinus surgery is a procedure used to remove blockages in the sinuses (the spaces filled with air in some of the bones of the skull). These blockages cause sinusitis, a condition in which the sinuses swell and become clogged, causing pain and impaired breathing. The endonasal approach for surgical treatment of sinusitis

has become increasingly established during the last few decades. In such a procedure, information is provided primarily through the endoscope. The limited information from the endoscopic view requires from the surgeon a detailed knowledge of the anatomy (Fig. 1).

The sinuses are physically close to the brain, the eye, and major arteries. Thus, endonasal sinus surgery requires a high degree of precision, since minor misjudgments of anatomical relationships can lead to catastrophic consequences. These demands are particularly challenging when surgical landmarks used for navigation are distorted or obscured by extensive disease or previous surgeries. Surgical navigation systems, which allow for the real time tracking and localization of

* Corresponding author.

E-mail addresses: burschka@cs.jhu.edu (D. Burschka), liming@cs.jhu.edu (M. Li), mishii3@jhmi.edu (M. Ishii), rht@cs.jhu.edu (R.H. Taylor), hager@cs.jhu.edu (G.D. Hager).

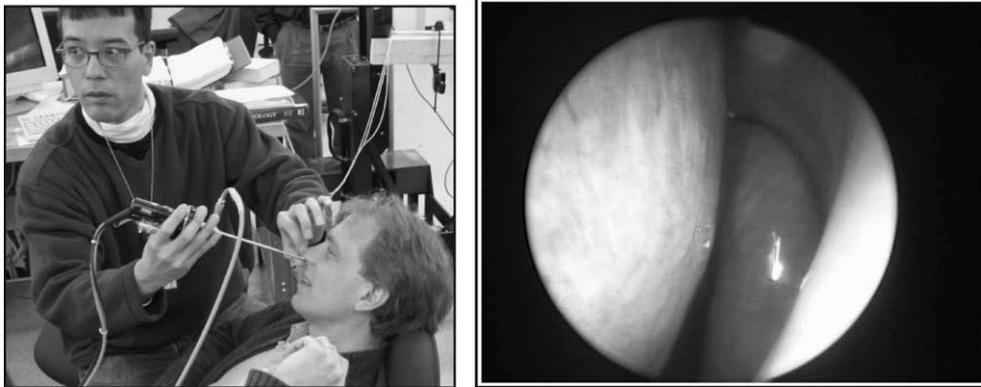


Fig. 1. Endoscopic inspection of the nasal sinus cavities depicting the limited information provided to the surgeon in the current procedures.

surgical instruments with respect to surrounding anatomical structures, are thus essential for safety and have been employed in endonasal approaches to simplify the procedure. The evolution of this field is driven in part by technical advancements such as computer-aided surgical navigation. Fig. 2 shows a typical system (here a VTI system from GE Medical Systems) for endoscopic sinus surgery.

This paper addresses two fundamental problems associated with vision-based navigation of surgical tools in the human body: (1) the reconstruction of the 3D surface geometry within the view of the endoscopic camera and (2) the registration of the camera to a pre-operative CT scan that allows a verification of the pre-operative procedure planning.

There are a variety of surgical navigation systems that rely on artificial optical and magnetic fiducials. We give a short overview over the existing approaches in the following section.

1.1. Related work

Surgical Navigation Systems were first developed in the 1980's for neurosurgical applications (e.g., Watanabe et al., 1987; Kosugi et al., 1988; Reinhardt et al., 1988). They have subsequently been applied in many

surgical fields, including neurosurgery (Smith et al., 1994; Heilbrun et al., 1993; Maciunas, 1993; Galloway et al., 1991; Barnett et al., 1993; Zinreich et al., 1993; Reinhardt, 1996; Scholtz et al., 1998; Dey et al., 2002), craniofacial surgery (Cutting et al., 1990, 1992, 1997; Burghart et al., 1999; VanderKolk et al., 1992), ENT (Adams et al., 1990, 1991, 1992; Bartz et al., 2001), spine surgery (Lavalée et al., 1994; Merloz et al., 1997; Nolte et al., 1994; Cleary, 1999), and orthopedic surgery generally (DiGioia et al., 1996; Picard et al., 2001; Simon et al., 1997; Kunz et al., 2001; van HellenMondt et al., 2001; Nolte and Visarius, 1996). Applications in other surgical disciplines, such as minimally invasive hepatic surgery (Herline et al., 1999; Stefansic et al., 2000), are also being explored. Their primary purpose is to provide intra-operative information about the geometric relationships (e.g., between a surgical tool and anatomical structures) that cannot be readily observed using the surgeon's normal visual and haptic senses.

A typical surgical navigation system consists of a computer workstation, navigational tracking device, and associated tools with marker devices whose pose (position and orientation) is continuously measured relative to the navigational tracker (Fig. 3). Usually, one or more "reference" markers are affixed to the patient's anatomy in order to eliminate the effects of patient or

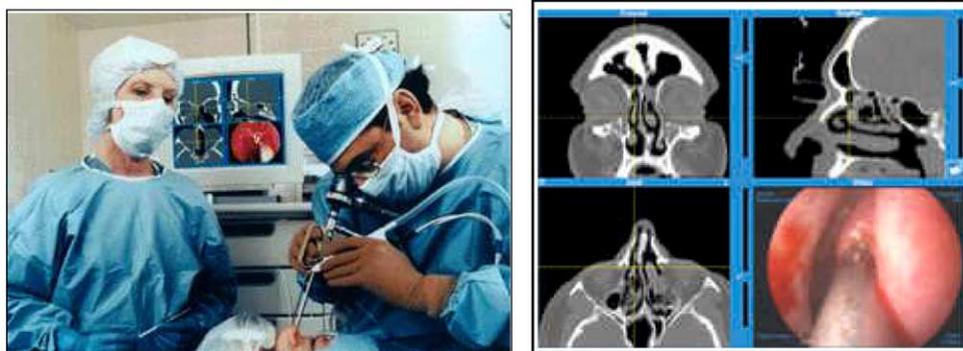


Fig. 2. Typical surgical navigation system for endoscopic sinus surgery. (Left) Operating room scene; (right) navigation screen shot (photos from VTI, Inc. – now part of GE Medical Systems).

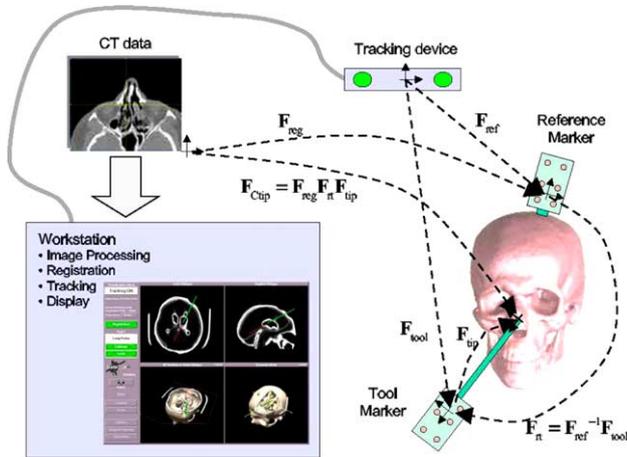


Fig. 3. In the current surgical systems, the camera position is registered indirectly relative to external fiducials (markers) measured by a tracking device.

tracker motion. The workstation typically imports pre-operative CT, MRI, or other volumetric image data associated with the patient. After suitable calibration, typically relying on either artificial or anatomical landmarks, image processing, and registration steps are performed. This allows the system to determine the transformation between volumetric image coordinates and the patient reference coordinates. Typically, the system also computes an estimate of the accuracy of the registration, and allows for re-registration if the initial solution is faulty or if re-registration is otherwise deemed necessary. Commonly used tracker technologies include specialized optical tracking systems (Maciunas, 1993; Nolte et al., 1994; Bucholz and Smith, 1993; Taylor et al., 1992; Lavallee et al., 1996; Hofstetter et al., 1999), stereo vision with conventional cameras (Heilbrun et al., 1991), electromagnetic sensors (Wu and Taylor, 2003; Poulin and Amiot, 2001), acoustic sensors (Reinhardt, 1996; Bucholz and Smith, 1993), and mechanical linkages (Kosugi et al., 1988; Galloway et al., 1991; Reinhardt, 1996). Currently, systems based on specialized optical devices such as the Polaris™ or OptoTrak™ (Northern Digital, Waterloo, Ont.) are the most accurate, but each technology has advantages and disadvantages.

While surgical navigation has been a catalyst for advancement in this field, it suffers from a number of fundamental limitations that constrain further surgical innovation. These limitations include the inability to register the surgical site to anatomical landmarks, the inability to account for changes in anatomy brought about by surgery, and the inability to repetitively and autonomously register a patient. The first limitation is important, since it decreases the practical accuracy of surgical navigation. Current registration schemes rely on surface anatomy or fiducial points (Kennedy et al., 2001), away from the surgical site, to perform the regis-

tration, often leading to relatively large registration errors in the region of greatest interest. The second limitation means that, as the surgery progresses, the navigation system becomes less and less useful, as does the preoperative data. The latter leads to degradation of the surgical system's performance with time due to patient motion, shifts in the reference frame and so forth. Currently, the operating surgeon must update the registration at several points in the operative field throughout the procedure. Typically, registration is verified by localization on known bony landmarks on the skull and in the nasal cavity.

Most navigation systems today report position errors on the order of 2 mm or less (Fried et al., 1997; Metson et al., 1998). However, the accuracy of registration can vary widely depending on the location of the surgical site relative to the landmarks. In particular, sites located far from a fiducial will generally have higher error than those near a fiducial. This is a consequence of the indirect nature of the tool-to-anatomy calculation.

1.2. Structure of the paper

In this paper, we present a novel method for intra-operative registration directly from the endoscopic images without manual inputs from the surgeon. It is especially useful in revision cases, where the surgical landmarks are usually absent. The paper is structured as follows. In Section 2, we describe the underlying image processing that allows us to recover the 3D-structure and the motion from monocular images of an endoscopic camera, and the way we perform the final alignment using a modified iterative closest point (ICP) approach. In Section 3, we present the experimental results on the phantom skull. We conclude in Section 4 with an evaluation of our approach and present our future research goals.

2. Approach

The two major problems that we address in this paper are: (1) 3D reconstruction from monocular camera images and (2) registration of the reconstructed 3D model to a pre-operative CT scan.

Our system reconstructs a scaled 3D model of the environment from a monocular camera. This reconstruction requires knowledge about the motion of the camera, which we assume to be unknown, or at least uncertain. Thus, in parallel to model reconstruction, we estimate the motion of the camera as well. We discuss the implementation of the vision-based reconstruction in Section 2.1.

The 3D structure estimated from monocular camera images is known only up to scale. The correct scale needs to be recovered from the data to align the points

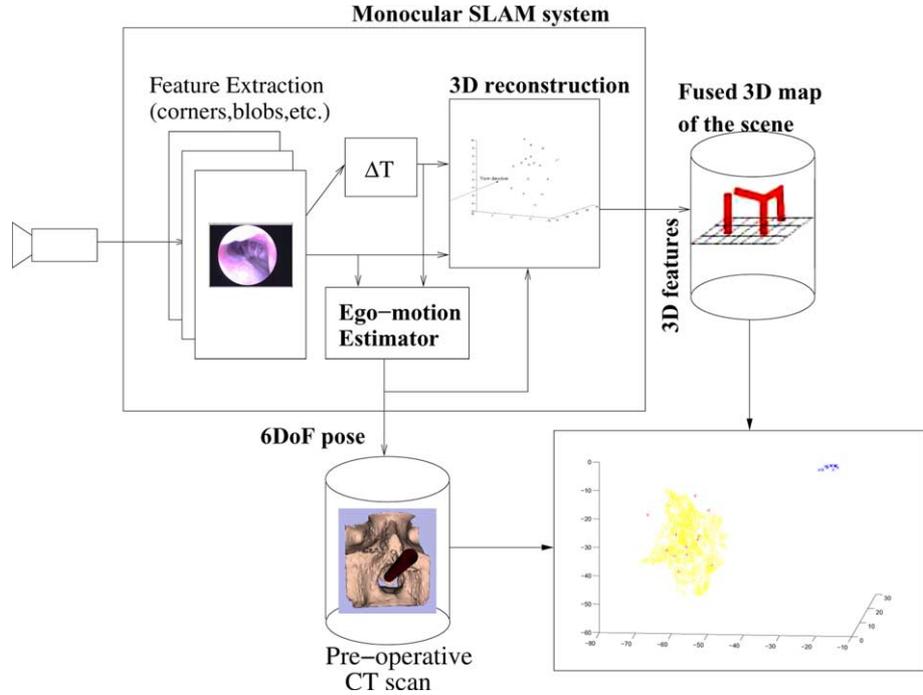


Fig. 4. The architecture of our system.

roughly with the CT scan. The remaining alignment error between the CT scan data and the reconstructed model is corrected with a modified ICP algorithm with covariance tree optimization (Section 2.2.2).

The architecture of our system, depicted in Fig. 4, thus consists of two major components solving the problems stated above: (1) a monocular VSLAM system that reconstructs a scaled version of the environment and (2) a registration module mapping the data reconstructed from the camera images to the CT scan.

In the following sections, we will describe the details of the system implementation.

2.1. Scaled 3D reconstruction

As mentioned earlier, an important component of our system is the 3D reconstruction performed by the *Monocular SLAM system* in Fig. 4. This module takes the monocular image data and extracts point features (e.g., corners or textured regions) in the *Feature Extraction* module and uses this pre-processed data to estimate the *Ego-Motion* of the camera relative to the environment in all 6 degrees of freedom. This motion data is then used to reconstruct 3D information of newly found features in the *3D reconstruction* module.

2.1.1. System initialization

Our approach requires an initial guess about the 3D structure of at least three landmarks. Two possibilities for obtaining this initialization are:

- *the eight-point-algorithm* applied to estimate the *Essential Matrix* of the system from eight-point correspondences. A relation between the projections p_i, p_i^* in two camera images with known intrinsic calibration parameters can be expressed with the Essential Matrix $\tilde{\mathbf{E}}$ as $p_i^* \tilde{\mathbf{E}} p_i = 0$.

The Essential Matrix $\tilde{\mathbf{E}}$ consists of a product of two matrices

$$\tilde{\mathbf{E}} = \tilde{\mathbf{R}} \cdot \text{sk}(\mathbf{T}) \quad \text{with} \quad \text{sk}(\mathbf{T}) = \begin{pmatrix} 0 & -T_z & T_y \\ T_z & 0 & -T_x \\ -T_y & T_x & 0 \end{pmatrix}. \quad (1)$$

Note that, given a correspondence, we can form a linear constraint on $\tilde{\mathbf{E}}$. It is only unique up to scale, therefore, we need 8 matches to calculate $\tilde{\mathbf{E}}$. From $\tilde{\mathbf{E}}$, it is possible to recover the rotation matrix $\tilde{\mathbf{R}}$ and the scaled version of the translation vector \mathbf{T}' (Brown and Ballard, 1982). This is sufficient for our approach as we describe later (Section 2.1.3).

- *manual feature selection* in the endoscope image. Here, the surgeon selects three points with known correspondences to the CT-data and the system uses this information to bootstrap the processing.

The first alternative is completely unsupervised, but it requires a significant initial movement to get a well-conditioned *Essential Matrix*. The second alternative is similar to the current image guided system (IGS) procedure, but it is necessary just for the first frame of the sequence.

2.1.2. Feature extraction

The 3D reconstruction algorithm assumes that *point features* are extracted from the images. Possible features are: intersections of contours resulting from edge filters (Brown and Ballard, 1982) or simply textured areas of the image used for template matching in a sum of square differences (SSD) matching algorithm (Hager and Belhumeur, 1998). There are several other possible matching algorithms to use in this system (Brown et al., 2003); some using local measures (Banks and Corke, 2001) and some incorporating global surface constraints (Scharstein and Szeliski, 2002; Ramey et al., 2004).

One issue in real endonasal images is the sparse density of points that can be reliably placed in correspondence and used for model reconstruction. Another problem is the moving light source, which is attached to the endoscope (Fig. 5). This changes the shadows cast in the images due to small structures on the observed surfaces. The latter necessitates a brightness independent image feature matching algorithm.

Our initial results were based on experiments with a phantom skull. This skull does not have any detectable texture, so sparsely distributed colored points were added to the surface. These are segmented in the hue space of the color representation. This way, we are able to identify and track the features in image sequences using a simple color blob tracker despite the changing lighting conditions (Fig. 6).

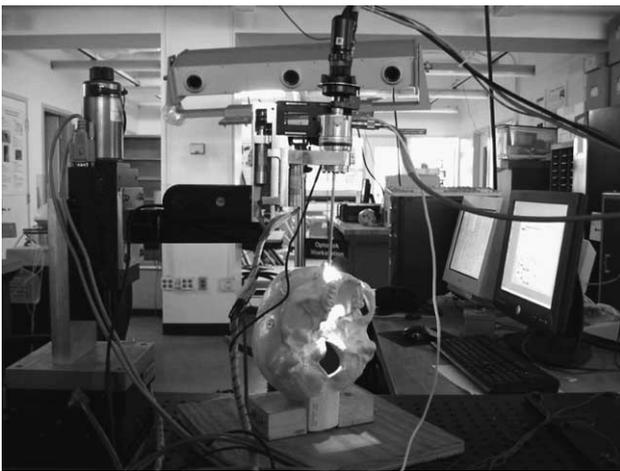


Fig. 5. Our experimental system.

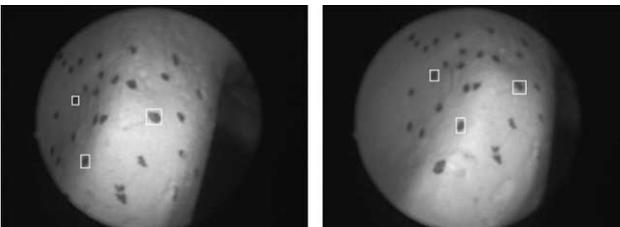


Fig. 6. Example of corresponding points on our phantom.



Fig. 7. Validation of the approach on a head of a pig cadaver with anatomical structures similar to the human sinuses. OptoTrak data is used as a ground truth to validate the motion estimation of our algorithm.

We have also obtained preliminary results with real endonasal images using an endoscopic camera (Fig. 7). We used an image tracker based on the minimization of the SSD between a reference image template and its current projection in the image. This tracker follows an image template compensating for rotation, translation, illumination changes and occlusions. The details of the implementation are presented in (Hager and Belhumeur, 1998). The tracker was able to track point features in real sinuses of the head as depicted in Fig. 8. The estimated positions of the centers of the tracked regions are used to subsequently estimate the motion of the camera. Given motion, it is possible to recover 3D information for the newly established correspondences. Additional correspondences allow us to add new points to the model to compensate for loss of visible points due to occlusions and coming out of the field of view.

2.1.3. Localization and mapping step

Since the camera motion needs to be estimated simultaneously with reconstruction, the so-called epipolar geometry from the motion between two camera frames needs to be recovered. An approach, commonly used in situations with at least eight point correspondences between images, is the *eight-point-algorithm*. The recovered *Essential Matrix* contains the information about the translation direction T and rotation R between the images. The translation information can only be recovered to a scale factor due to the homogeneous structure of the epipolar constraint (Brown and Ballard, 1982).

The number of corresponding (detectable) points between two camera frames varies significantly during the sinus surgery. There are situations when fewer than

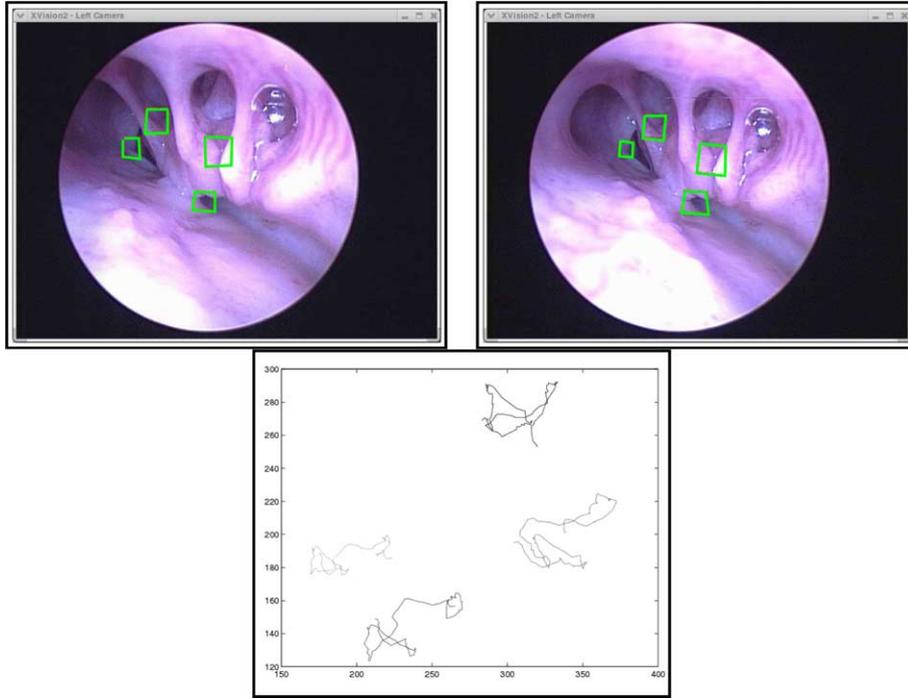


Fig. 8. Tracking result in real sinus images: (top) begin and end of a sequence with four point features being tracked; (bottom) tracking result in image coordinates.

eight points can be matched. The above approach fails in these cases, so we instead apply a method for camera localization and mapping requiring merely three point correspondences. We will sketch out the process below. The reader should consult (Burschka and Hager, 2004) for details of the algorithm.

In this approach, we assume that each 3D point P_i imaged in a normalized camera frame $p_i = (u_i, v_i, 1)^T$ can be represented as its direction vector $n_i = p_i / \|p_i\|$ and the distance to the real point D_i so that $P_i = D_i \cdot n_i$. Since, in typical applications, the scale m of the reconstruction may be unknown, the system operates with a scaled version of the distance $\lambda_i = D_i/m$. This approach calculates an estimate for the rotation $\tilde{\mathbf{R}}$ and the scaled translation \mathbf{T}'^* between the points in the current frame $\{P_i\}$ and the next frame $\{P_i^*\}$ as

$$\begin{aligned} \bar{P} &= \frac{1}{n} \sum_{i=1}^n P_i, & \bar{P}^* &= \frac{1}{n} \sum_{i=1}^n P_i^*, \\ P_i' &= P_i - \bar{P}, & P_i'^* &= P_i^* - \bar{P}^*, \\ \tilde{\mathbf{M}} &= \sum_{i=1}^n P_i'^* P_i'^T, & [UDV^T] &= \text{svd}(\tilde{\mathbf{M}}), \\ \tilde{\mathbf{R}} &= V \cdot U^T, & \mathbf{T}^* &= \bar{P}^* - \tilde{\mathbf{R}} \bar{P}. \end{aligned} \quad (2)$$

The approach requires an initial guess for the values for λ_i for the first frame and it computes a guess for translation \mathbf{T}'^* and rotation $\tilde{\mathbf{R}}$ in an iterative process. In the initial step, it assumes the distances to the points in the new image are identical to the previous one $\lambda_i' = \lambda_i$ and, afterwards, it iteratively converges to the

true $\tilde{\mathbf{R}}$, \mathbf{T}'^* , and λ_i' . Details and simplifications of the algorithm are discussed in (Burschka and Hager, 2003). This algorithm requires only three corresponding points between both images to compute the pose difference between two camera frames $(\tilde{\mathbf{R}}, \mathbf{T}'^*)$, which makes it more suitable for the given application.

Eq. (2) updates the distance values λ_i' for all tracked points P_i' for the new frame. New points $P_x = \lambda_x n_x$ can easily be added to the system under a rigid body assumption. The parameters of P_x can be calculated by solving the following equation:

$$(\tilde{\mathbf{R}} n_x \quad -n_x^*) \begin{pmatrix} \lambda_x \\ \lambda_x^* \end{pmatrix} = \tilde{\mathbf{R}} \lambda_1 n_1 - \lambda_1^* n_1^*. \quad (3)$$

Alternatively, a more robust estimate can be computed from 3 frames by solving the following equation:

$$\begin{pmatrix} \tilde{\mathbf{R}}_1 n_x & -n_x^* & 0 \\ \tilde{\mathbf{R}}_2 \tilde{\mathbf{R}}_1 n_x & 0 & n_x^{**} \end{pmatrix} \begin{pmatrix} \lambda_x \\ \lambda_x^* \\ \lambda_x^{**} \end{pmatrix} = \begin{pmatrix} \tilde{\mathbf{R}}_1 \lambda_1 n_1 - \lambda_1^* n_1^* \\ \tilde{\mathbf{R}}_2 \tilde{\mathbf{R}}_1 \lambda_1 n_1 - \lambda_1^{**} n_1^{**} \end{pmatrix}. \quad (4)$$

The pose change from image $1 \rightarrow 2$ is annotated here as $(\tilde{\mathbf{R}}_1, \mathbf{T}_1)$ and the pose change between images $2 \rightarrow 3$ is annotated as $(\tilde{\mathbf{R}}_2, \mathbf{T}_2)$. This equation estimates the distance λ_x to a new point P_x in the scale of an already known point P_1 from the currently tracked set of points. This way the newly added points are still measured with the same scaling factor m and the resulting 3D model has a uniform scale.

2.2. Registration of the endoscope data to CT scan

The other important module in our system depicted in Fig. 4 is the registration of the scaled data extracted in the previous section. The extracted patch needs to be localized correctly within the CT scan and registered to it. This is a two-step process in which we recover the scale from a PCA based method and later use ICP method to align it correctly.

2.2.1. Scale recovery for 3D reconstruction

The scaling factor m in Section 2.1.3 depends on the scale of the λ_i -values for the initial set of points. In case of an unsupervised bootstrap using the *eight-point-algorithm* (Section 2.1.1) the resulting reconstruction has an arbitrary scale that depends on the scale of the translation vector \mathbf{T}' , which is usually assumed to be a unit vector.

The system usually has a rough estimate of the current camera position. We use this estimate to carve out part of the CT surface data that falls into the expected visibility cone of the camera. This cone is slightly enlarged in all directions to compensate for the unknown camera motion. The size of the extracted patch depends on the uncertainty about the camera position. Fig. 9 shows examples of CT scan and reconstructed data.

The visible regions are usually surfaces with two dominant directions of the cavity walls with a third dimension representing the surface structure or combinations of such walls. We assume for now that the CT data patch consists of a single surface with some surface structure on it. We will discuss later, how to split more complex structures into simple patches.

We use the property of two dominant surface directions for our scale recovery by calculating the covariance matrix \tilde{C}_k of the point cloud in the selected CT scan region and in the current camera reconstruction. The eigenvalues and eigenvectors of \tilde{C}_k define a new coordinate system with the two eigenvectors calculated from the larger eigenvalues defining the supporting

plane in the cloud and the third eigenvector describing the depth variation in the measurement.

In both cases, the smallest eigenvalue (E_{ct}, E_{3d}) represents a metric for the depth variations in the surface of the CT scan and in the reconstructed point cloud. The normalized eigen-vectors $\{V_{ctx}\}$ and $\{V_{3dx}\}$ and the eigenvalues allow us to calculate the scale m and the rotation \tilde{R}_{tot} between the two data sets to (5). The rotation matrix \tilde{R}_{tot} aligns both dominant surfaces along their normal vectors, which are represented by the eigenvector calculated from the smallest eigenvalue (last column in each of the rotation matrices in (5)). The rotation about the normal vector cannot be restored in this way

$$m = \frac{\sqrt{E_{ct}}}{\sqrt{E_{3d}}}, \quad V_{p \in \{ct, 3d\}} = (V_{px} V_{py} V_{pz})^T,$$

$$V_{n-p} = \frac{(0V_{pz} - V_{py})^T}{\|(0V_{pz} - V_{py})^T\|},$$

$$\tilde{R}_{ct} = ((V_{n-ct} \times V_{ct}) \quad V_{n-ct} \quad V_{ct}),$$

$$\tilde{R}_{3d} = ((V_{n-3d} \times V_{3d}) \quad V_{n-3d} \quad V_{3d}),$$

$$\tilde{R}_{tot} = \tilde{R}_{3d} \cdot \tilde{R}_{ct}^T.$$
(5)

We apply the scaling m and rotation \tilde{R}_{tot} to the zero mean point clouds that were used to calculate the covariance matrices above. This way, we obtain two point clouds with the same alignment, but the CT-scan represents a much larger area because of the unpredictable camera movement. Both clouds have a similar scale and alignment of the supporting plane. Fig. 10 depicts an alignment result of the scaled point cloud in the top right corner to the 3D surface from the CT scan. The transformed points may have a significant rotational error around the normal vector visible in the figure as crosses ('+'). We correct this error in the next step. Fig. 11 depicts the pseudo-image representations used to correct the rotational error.

We consider now both rotated point clouds as sparse “images,” where each “pixel” is represented by

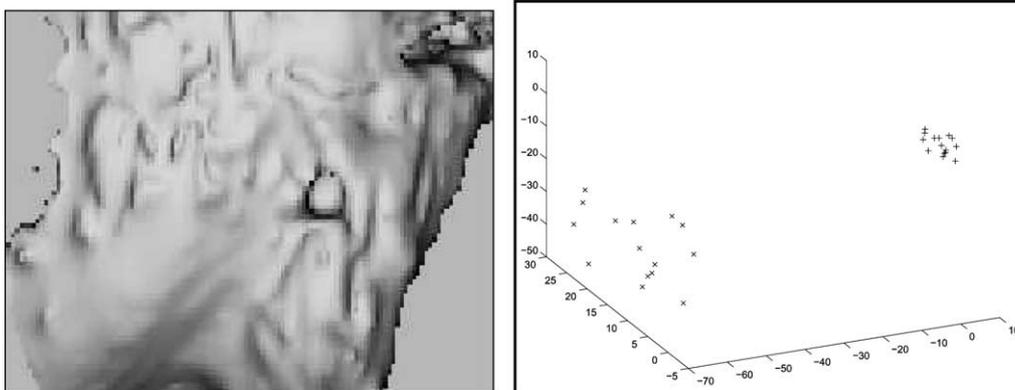


Fig. 9. Scaled reconstruction of surface points: (left) CT scan visualization of the area, (right) matched surface points with ICP {left point cloud}, scaled reconstructed points {right point cloud}.

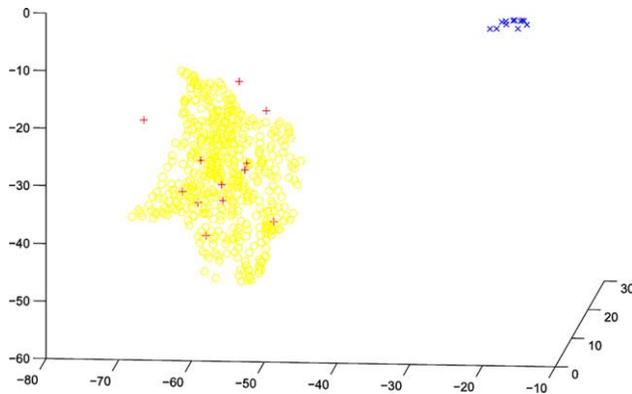


Fig. 10. After the alignment along the normal vector to the supporting plane the scale is roughly recovered, but rotation around the normal vector is possible.

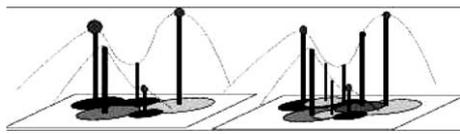


Fig. 11. Distance to the *supporting plane* calculated in (5) is used as a pseudo-image representation to match the sparse reconstruction (left) to the dense point cloud (right).

its distance to the supporting plane calculated from the covariance matrix. We use the reconstructed 3D structure from the current view as a template that is matched to the “image” constructed from the CT scan data using the SSD matching approach described in (Hager and Belhumeur, 1998). In this case we can allow only an optimization in in-the-image rotations to find the missing parameter for the alignment and we treat the missing information in the sparse image of the reconstruction as “occlusions.” This allows us to find the rotation between the two point sets and to localize the current view within the larger region of the CT scan to model the uncertainty in the registration of the camera relative to the CT scan. The “search window” can be adjusted dependent on the current knowledge about the relative position of the two coordinate frames to each other.

The physical positions of the sampling points, especially in the 3D reconstruction, do not necessarily correspond to the extreme values of the surface hull. Therefore, the above match trial can fail for a specific set of three points. The 3D reconstruction may not have reconstructed the peak value but some random value along the slope instead.

The resulting match is used to align the two data sets. The residual error is due to imperfect sampling and the coarse structure of the point sets, especially in the case of the reconstructed data from the phantom skull.

The 3D scaling step needs to be performed just in the initial bootstrap phase and in cases when the system was not able to maintain the minimum number of three features and needs to re-initialize the distance measurements.

The above scaling step relies on the fact that the data contains only one supporting plane with depth variations that are used to calculate the scaling factor m (5). In case, when the CT or the reconstructed data set contains two or more surfaces in a corner arrangement (Fig. 12), the structure needs to be split into single surfaces before we apply the above scaling. We use a modified version of a split-and-merge algorithm here. We establish, for the entire data set, the plane equation of the supporting plane A based on \tilde{C}_k and (5). For each point P_i of the data-set, we calculate the distance d_i from the plane

$$d_i = (P_i - \bar{P}) \cdot V_n \quad (6)$$

with V_n being the normal vector of the estimated plane A .

Points with the largest deviation from the original plane A and border points on both sides of the data set are used to split the original surface consecutively into sub-surfaces defined by these points. Each three points define one sub-surface. We split the points of the original surface into sub-surfaces depending on the distances to the resulting sub-planes. The above evaluation is repeated on the sub-surfaces until the maximum

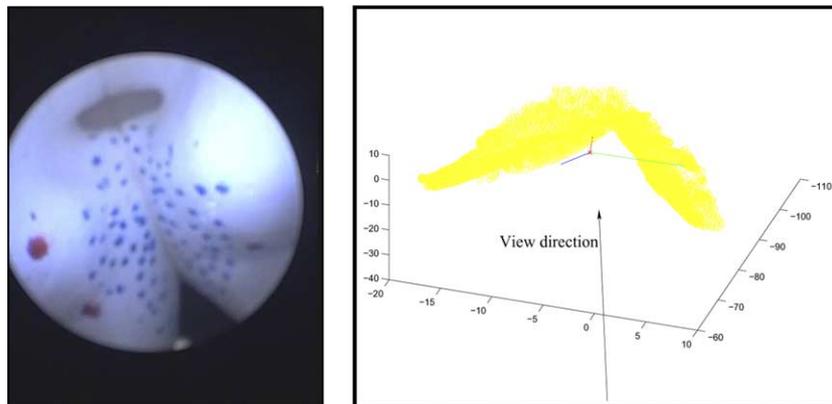


Fig. 12. The two surfaces need to be separated first, before our PCA-based scale recovery can be applied.

deviation $\max\{d_i\}$ is smaller than the expected depth structure in the surfaces.

2.2.2. ICP

Now, we have reconstructed and localized a 3D dataset with endoscopic images, which has right scale and similar orientation and translation in the coordinate frame of the CT scan.

We have chosen to perform the rigid registration between CT images and physical data reconstructed from endoscopic images using the ICP algorithm. For some applications in the endoscopic surgery, a deformable registration method can be further applied based on the results of the ICP. The anatomy of the nasal and sinus cavity is mostly bony tissue. To this end, ICP is considered a suitable registration solution for this application. An alternative approach for the alignment, that performs the point alignment and the registration with reduced sensitivity to deformation and outliers is described in (Chui and Rangarajan, 2000). This approach uses a thin-plate spline assumption to iteratively match sets of points and recover the alignment scale. The advantage in our case is that for the initial alignment and scale recovery only the region properties of the CT scan are considered. We have a strong mismatch in the number of points reconstructed by the camera and recovered from the CT scan that makes the application of the algorithm in (Chui and Rangarajan, 2000) difficult.

We use a covariance tree data structure to search for the closest point for ICP. A covariance tree is a variant of a k -dimensional binary tree (k -D tree). The traditional k -D tree structure partitions space recursively along principal coordinate axes. In our covariance tree each sub-space is defined in the orthogonal coordinate system of the eigenvectors centered at the center of mass of the point set, and is recursively partitioned along this local coordinate frame. An important advantage of covariance trees is that the bounding boxes tend to be much tighter than those found in conventional k -D trees and tend to align with surfaces, thus producing a more efficient search (Williams et al., 1997).

3. Results

The system is implemented in C++ using the image processing library XVision2 (Hager and Toyama, 1998) and the Vision-based simultaneous localization and mapping (SLAM) library VGPS (Burschka and Hager, 2004). It was tested on a laptop computer with a Pentium-M processor at 1.2 GHz running Linux OS.

3.1. Geometric investigations

In this section, we first review the essential geometric concepts necessary for our algorithms, and calculate the

theoretical accuracy with which geometry can be computed in several representative cases. We then describe initial registration results we have achieved using prototype algorithms on phantom data. This data validates the concept that 3D–3D surface registration from video data is possible given good image information. Finally, we describe preliminary results on video tracking and reconstruction from porcine cadaver data to demonstrate that the image-level matching problem can be solved on realistic data.

3.1.1. Point reconstruction accuracy

We performed a standard calibration procedure (Zhang, 2000; Bouget) on a Storz Telecam 20212113U NTSC equipped with a 0° lens. Using the resulting parameters, we then calculated the expected accuracy of point location reconstruction assuming image point localization to 1 pixel. This is generally considered conservative – many systems report point localization accuracy to 1/4 pixel or better (Brown et al., 2003).

Fig. 13 shows the expected reconstruction accuracy for several cases. The x axis is the distance to the target point, and the y axis is the reconstruction error due to a 1 pixel mislocalization. Three cases are considered: (1) a point directly in front of the endoscope and a motion parallel to the image plane of 1 mm, (2) a point 15° to the side and motion of 3 mm in the direction of the optical axis, and (3) a point 15° to the side and a motion consisting of 1 mm side to side motion with 3 mm forward motion. As can be seen from the figure, for targets within 2 cm of the endoscope, reconstruction accuracies on the order of 1.2 mm or less can be expected. Indeed, for targets 1 cm from the endoscope, accuracies are 0.3 mm or better.

3.1.2. Viewpoint registration accuracy

Consider a set of points $p_i = (x_i, y_i, z_i)^T$ lying on a surface, let c_i denote the optical center of the camera, and

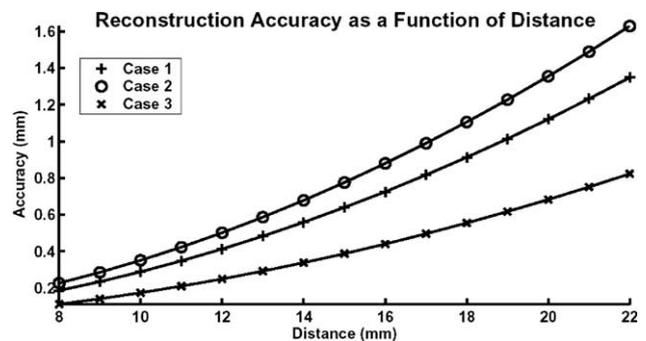


Fig. 13. Expected reconstruction accuracy for several cases. Case 1: a point located in different distances along the optical axis estimated using 1mm lateral motion. Case 2: a point located along a 15° off-axis estimated using a 3 mm forward motion. Case 3: a point located 15° off-axis estimated using combined 1 mm lateral and 3 mm forward motion.

let $q_i = (u_i, v_i)$ denote the image observations of the points. Our objective is to determine the accuracy with which the observer position c_i can be calculated from noisy observations. We note that we can consider the case when $c_i = 0$, as any other problem can be transformed to this case by a rigid body motion. Also, note that, in general, we need at least three points in general position to compute the full rigid body pose of the camera.

Under these assumptions, the first order term in a Taylor series expansion of the camera projection equations yields

$$\Delta q_i = \frac{1}{z_i} \begin{bmatrix} f & 0 & u_i \\ 0 & f & v_i \end{bmatrix} \Delta c = J_i \Delta c. \tag{7}$$

If we now assume that we have a covariance matrix $A_q = \sigma^2 I_2$ modeling the accuracy of each observation q_i , we can compute the expected covariance on c_i as

$$\lambda_c = \sigma^2 \left(\sum_i J_i^T J_i \right)^{-1}. \tag{8}$$

Table 1 shows the expected accuracy of registration as a function of distance and point spread for 4 points arrayed symmetrically about the optical center.

Table 1
Expected registration accuracy for several representative cases

Angle	5°		10°		15°	
	10 mm	20 mm	10 mm	20 mm	10 mm	20 mm
Accuracy parallel to image plane	0.0148	0.0296	0.0146	0.0292	0.0143	0.0287
Accuracy along optical axis	0.1195	0.2389	0.0586	0.1172	0.0378	0.0756

It is assumed that four points are observed. The points are located at the given distance from the image plane and their rays of projection form the stated angle with the optical axis.

This data is again for the previously described endoscope. It can be seen that the expected registration accuracy is more than an order of magnitude better than that stated for commonly used navigation systems (Fried et al., 1997; Metson et al., 1998).

3.2. Tracking

To verify the feasibility of video feature tracking on sinus tissue, we acquired data ex vivo from a porcine cadaver. The model was acquired within 24 h of slaughter. The nose was shortened to make the anatomy more consistent with the human sinuses. Video and tracking data was acquired in a manner identical to the skull phantom with the exception that the endoscope was manually (rather than robotically) manipulated.

In this case, tracking was performed using gradient optimization of a local correlation measure with respect to the four parameters: two locations, scale and orientation. Fig. 14 shows tracking results in two of five test cases where different types of features were tracked. The tracking uses a variety of anatomical structures in the images varying from significant vessel structures visible in the images to small structures on the actual surfaces that allow a robust tracking. A validation method

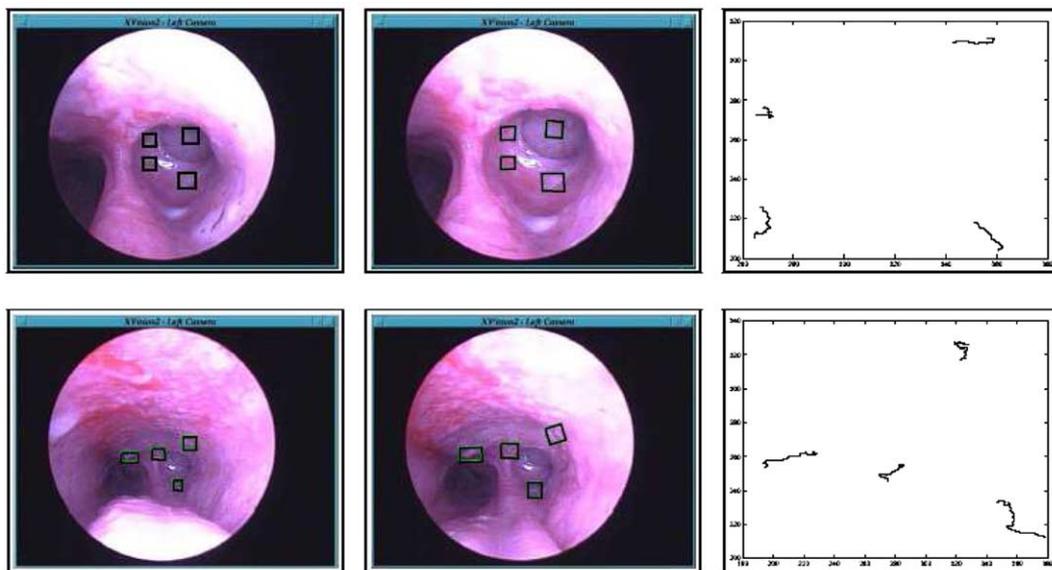


Fig. 14. The first column shows the originally chosen regions, the second the locations at a later point in the sequence and the final column shows the change in location of the features through the sequence.

described in (Burschka and Hager, 2004) allows a robust filtering of reliable features in the set of tracked landmarks. We need to ensure that the feature position is stable to ensure an accurate calculation of the camera ego-motion from the collected image data.

In this case, the features themselves were chosen by hand rather than automatically as it was the case for the phantom skull. Manual inspection of the tracking data shows it to be quite accurate, however, no external reference validating these tracking results is yet available.

The experiments on the porcine head reveal that the type of useful features for tracking depends on the region where the navigation is performed. Along the nasal area, useful features seem to be vascular structures that provide the necessary gradient information to uniquely define the tracking region over a sequence of images. In deeper regions of the skull, small surface structures visible in Fig. 14 provide a good gradient information for tracking. Finally, the sinus structures themselves define robust structures to be tracked in the actual sinus area.

3.3. Localization and mapping subsystem

The system is capable of tracking and reconstruction of 4–6 point features with 3D structure recovery in

frame-rate on a Pentium-M processor at 1.2 GHz running Linux OS. The actual iterative structure calculation with the reported accuracy of 0.5 mm takes ≈ 10 ms on the above CPU. This is part of the Ego-Motion Estimator depicted in Fig. 4.

3.4. Registration to the CT scan

The experimental validation of our approach is carried out on the setup depicted in Fig. 5. We track the position of the endoscope with the OptoTrak™ system in the background to verify the motion estimation results from our system.

Fig. 15 shows two reconstruction results from a camera motion of (4.8, 0.2, 5.2) (mm) with small and large rotation between the consecutive frames. The resulting reconstruction errors had a standard deviation of (0.62, 0.3382) for each of the cases. The minimal rotational error expressed as Rodrigues vector was $r = (0.0017, 0.0032, 0.0004)$, $(-0.0123, -0.0117, -0.0052)$. The error in the estimate of the translation vector was $\Delta T = (0.05, -0.398, 0.2172)^T$, $(-0.29, 0.423-0.4027)^T$ (mm).

We tested our registration with different reconstruction results (patches) that were registered to CT skull images. Because the 3D surface data reconstructed by a monocular camera may not cover the entire surface

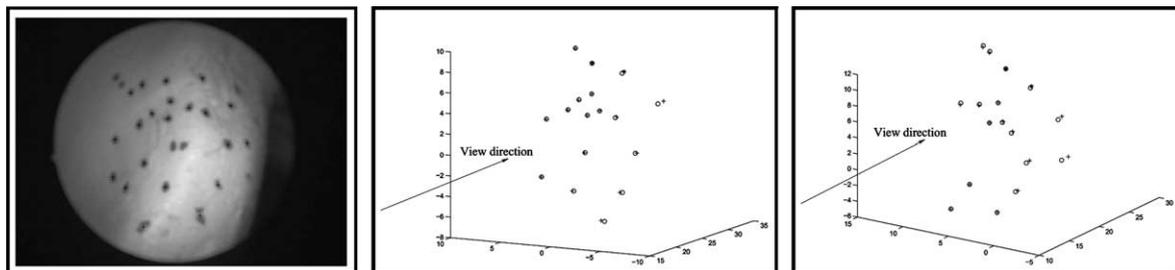


Fig. 15. 3D reconstruction results in camera coordinate frame from two consecutive reconstructions: (left) camera view (middle, right) reconstructed points '+', ground-truth from OptoTrak 'O'.

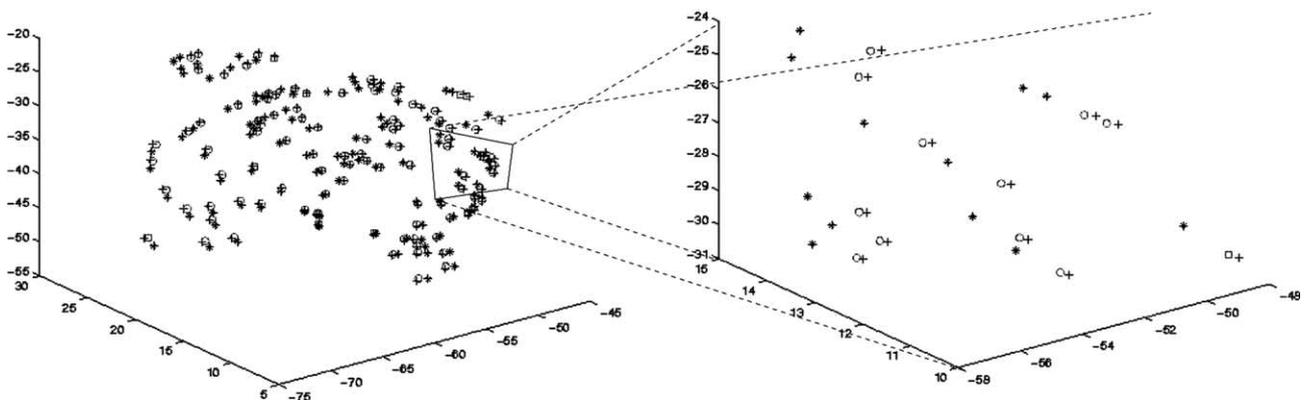


Fig. 16. The relative displacements of the sparse samples (+), their initial position recovered by VGPS (*) and their final position after alignment by ICP (O). Left is the global view of the sample data for a patch. Right is a closer look.

Table 2
The accuracy results of ICP

	Translation offset range (mm)	Rotation offset range (°)	Noise level (mm)	Average error (mm)
Patch 1	$\{\Delta X, \Delta Y, \Delta Z\} = \pm 10$	$\{\Delta \alpha_{k \in \{X, Y, Z\}}\} = \pm 10$	0.5	0.65
Patch 2	$\{\Delta X, \Delta Y, \Delta Z\} = \pm 10$	$\{\Delta \alpha_{k \in \{X, Y, Z\}}\} = \pm 10$	0.5	0.48

Patch 1 contains around 2500 triangles, and 100 sample points are used for test. Patch 2 contains around 900 triangles, and 50 sample points are used for test.

patch, we were interested in the sensitivity to drop-outs. We purposely removed parts of the data from the reconstructed patch. Our experiments with the phantom show that the ICP can accommodate noise levels in the data up to 0.6 mm, combined with translational offsets of up to 10 mm, and rotational offsets within 10°. The vision-based reconstruction gives us errors an order of magnitude below these limits.

After ICP alignment the average distance error for the sample points is around 0.65 mm. By comparison, the fiducial based registration residual error is around 0.40 mm for four fiducials that are attached to the surface of the skull. However, our method directly tells the registration error of the target region for the surgery (see Fig. 16 and Table 2).

4. Conclusions and future work

The presented system performs accurate reconstruction of 3D surface points based on images from an endoscopic camera. The points are successfully aligned with CT scans of our phantom skull in the sinus area. We plan to enhance the quality of the final registration using ICP by reconstructing dense surface models around the points reconstructed from the vision-based SLAM system. The reconstructed points will be used as seeds for a dense disparity reconstruction under a smoothness assumption in the local area around the point. This increased description of the surface properties around the reconstructed point promises additional matching criteria that can be used in the alignment process.

Our major goal is to more extensively test our system in different parts of the skull and on other range images to better evaluate the performance of the system. We are currently investigating the feature type that can be used for a robust estimation and tracking of our point features in real endonasal images obtained in a preliminary experiment on an animal cadaver.

Acknowledgments

Partial funding of this research was provided by the National Science Foundation under Grants EEC9731748 (CISST ERC), IIS9801684, IIS0099770, and IIS0205318. This work was also partially funded by a DARPA Mars grant.

References

- Adams, L., Gilsbach, J.M., Krybus, W., Meyer-Ebrecht, D., Moses, R., Schlondorff, G., 1990. Cas – a navigation support for surgery. In: 3D Imaging in Medicine. Springer, Berlin, pp. 411–423.
- Adams, L., Knepper, A., Krybus, W., Meyer-Ebrecht, D., Pfeiffer, G., Rueger, R., Witte, M., 1991. Navigation support for surgery by means of optical position detection and real-time 3D display. In: Proceedings Computer Aided Radiology. Springer, Berlin.
- Adams, L., Knepper, A., Krybus, W., Meyer-Ebrecht, D., Pfeiffer, G., Ruger, R., Witte, M., 1992. Orientation aid for head and neck surgeons. *Innovation et Technologie en Biologie et Medicine* 14 (4), 409–424.
- Banks, J., Corke, P., 2001. Quantitative evaluation of matching methods and validity measures for stereo vision. *International Journal of Robotics Research* 20 (7), 512–532.
- Barnett, G.H., Kormos, D.W., Piraino, D., Steiner, C.P., Weisenberger, J., Hajjar, F., Wood, C., McNally, J., 1993. Frameless stereotaxy using a sonic digitizing wand: development and adaptation to the picker vistar midical imaging system. In: Robert Maciunas, J. (Ed.), *Interactive Image-guided Neurosurgery*. American Association of Neurological Surgeons (Chapter 10).
- Bartz, D., Gurvit, O., Freudenstein, D., Schiffbauer, H., Hoffman, J., 2001. Integration of navigation, optical and virtual endoscopy in neurosurgery and oral and maxillofacial surgery. In: 3rd Caesarium – Computer Aided Medicine.
- Bouget, Jean-Yves. The matlab camera calibration toolkit Available from: <http://www.vision.caltech.edu/bougetj/calib_doc/>.
- Brown, C.M., Ballard, D.H., 1982. *Computer Vision*. Prentice-Hall, Englewood Cliffs, NJ.
- Brown, M.Z., Burschka, D., Hager, Gregory D., 2003. Advances in computational stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25 (8), 993–1008.
- Bucholz, R., Smith, K., 1993. A comparison of sonic digitizers versus light emitting diode-based localization. *Interactive Image-Guided Neurosurgery*, 179–200.
- Burghart, C., Krempien, R., Redlich, T., Pernozzoli, A., Grabowsky, H., Muncherberg, J., Hassfeld, S., Albers, J., Vahl, C., Rembold, U., Worn, H., 1999. Robot assisted craniofacial surgery: first clinical evaluation. *Computer Assisted Radiology and Surgery*, 828–833.
- Burschka, D., Hager, G.D., 2003. V-GPS – Image-based control for 3D guidance systems. In: Proceedings of IROS, pp. 1789–1795.
- Burschka, D., Hager, G.D., 2004. V-GPS (SLAM): – Vision-based inertial system for mobile robots. In: Proceedings of ICRA, pp. 409–415.
- Chui, H., Rangarajan, A., 2000. A new algorithm for non-rigid point matching. In: Proceedings of Computer Vision and Pattern Recognition, pp. 44–51.
- Cleary, K., 1999. Workshop report: technical requirements for image-guided spine procedures.
- Cutting, C.B., Grayson, B., Kim, H.C., 1990. Precision multi-segment bone positioning using computer aided methods in craniofacial surgery applications. In: 12th IEEE Engineering in Medicine and Biology Conference. IEEE Press, New York.

- Cutting, C.B., Bookstein, F.L., Taylor, R.H., 1997. Applications of simulation and morphometrics, robotics in craniofacial surgery. In: Taylor, R.H., Lavallee, S., Burdea, G., Mosges, R. (Eds.), *Computer-integrated Surgery*. MIT Press, Cambridge, MA, pp. 541–544.
- Cutting, C., Taylor, R., Bookstein, R., Khorramabadi, D., Haddad, B., Kalvin, A., Kim, H., Nox, M., 1992. Computer aided planning and execution of craniofacial surgical procedures. In: *Proceedings of IEEE Engineering in Medicine and Biology Conference*.
- Dey, D., Gobbi, D., Slomka, P., Surry, K., Peters, T., 2002. Mixed reality merging of endoscopic images and 3D surfaces. *IEEE Transactions on Medical Imaging* 21 (1), 23–30.
- DiGioia, A.M., Simon, D.A., Jaramaz, B., Morgan, F., Blackwell, M., O'Toole, R.V., Colgan, B., Kischell, E., 1996. Hipnav: pre-operative planning and intra-operative navigational guidance for acetabular implant placement in total hip replacement surgery. *Computer Assisted Orthopedic Surgery*.
- Fried, M., Kleefeld, J., Gopal, H., Reardon, E., Ho, B., Kuhn, F., 1997. Image-guided endoscopic surgery: results of accuracy and performance in a multicenter clinical study using an electromagnetic tracking system. *Laryngoscope* 107 (5), 594–601.
- Galloway, R.L., Edwards, C.A., Schreiner, S., Thomas, J.G., Maciunas, R.J., 1991. A new device for interactive, image guided surgery. In: *Proceedings of SPIE Medical Imaging V*.
- Hager, G.D., Belhumeur, P., 1998. Efficient region tracking with parametric models of geometry and illumination. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20 (10), 1125–1139.
- Hager, G., Toyama, K., 1998. The XVision system: a general-purpose substrate for portable real-time vision applications. *Computer Vision and Image Understanding* 69 (1), 23–37.
- Heilbrun, M.P., McDonald, P., Wiker, C., Koehler, S., Peters, W., 1991. Stereotactic localization and guidance using a machine vision technique. *Stereotactic and Functional Neurosurgery* 58, 94–98.
- Heilbrun, M.P., Koehler, S., McDonald, P., Semenov, V., Peters, W., Wiker, C., 1993. Implementation of a machine vision method for stereotactic localization and guidance. In: Maciunas, R. (Ed.), *Interactive Image-Guided Neurosurgery*, pp. 169–177. AANS.
- Herline, A.J., Stefansic, J.D., Hartmann, S.L., Debelak, J.P., Pinson, C.W., Galloway, R.L., Chapman, W.C., 1999. Image-guided surgery: preliminary feasibility studies of frameless stereotactic liver surgery. *Archives of Surgery* 134, 644–650.
- Hofstetter, R., Slomczynowski, M., Sati, M., Nolte, L.P., 1999. Principles of precise fluoroscopy based surgical navigation. In: *Proceedings of 4th International Symposium on CAOS*, p. 28.
- Kennedy, D.W., Bolger, W.E., Zinreich, S.J., Zinreich, J., 2001. *Diseases of the Sinuses: Diagnosis and Management*.
- Kosugi, Y., Watanabe, E., Goto, J., Watanabe, T., Yoshimoto, S., Takakura, K., Ikebe, J., 1988. An articulated neurosurgical navigation system using MRI and CT images. *IEEE Transactions on Biomedical Engineering*, 147–152.
- Kunz, M., Langlotz, F., Strauss, J., Ruther, W., Nolte, L.-P., 2001. Development and verification of a non-CT based total knee arthroplasty system for the lcs prosthesis. In: *First Annual Meeting of CAOS International, Davos*, p. 131.
- Lavallee, S., Sautot, P., Troccaz, J., Cinquin, P., Merloz, P., 1994. Computer assisted spine surgery: a technique for accurate transpedicular screw fixation using CT data and a 3-D optical localizer. *Medical Robotics and Computer-assisted Surgery* 2, 315–332.
- Lavallee, S., Troccaz, J., Sautot, P., Mazier, B., Cinquin, P., Merloz, P., Chirossel, J.-P., 1996. Computer-assisted spinal surgery using anatomy-based registration. In: *Computer-integrated Surgery*. MIT Press, Cambridge, MA, pp. 425–449.
- Maciunas, R.J., 1993. *Interactive image-guided neurosurgery*. American Association of Neurological Surgeons.
- Merloz, P., Tonetti, J., Eid, A., Faure, C., Pittet, L., Coulomb, M., Sautot, P., Raoult, O., 1997. Computer-assisted versus manual spine surgery: clinical report. In: Grimson, E., Mosges, R. (Eds.), *Proceedings of First Joint Conference of CVRMed and MRCAS*, vol. 1205. Springer, Berlin, pp. 541–544.
- Metson, R., Gliklich, R., Cosenza, M., 1998. A comparison of image guidance systems for sinus surgery. *Laryngoscope* 108 (8), 1164–1170.
- Nolte, L.P., Zamorano, J., Jiang, F., Langlotz, G., Arm, E., Visurius, H., 1994. A novel approach to computer assisted spine surgery. In: *First International Symposium on Medical Robotics and Computer Assisted Surgery (MRCAS 94)*. Shadyside Hospital, Pittsburgh, pp. 323–328.
- Nolte, L.P., Visarius, H., et al., 1996. *Computer-assisted Orthopaedic Surgery*. Hofgrefe & Huber.
- Picard, F., DiGioia, A., Sell, D., Jaramaz, B., Moody, J., Nikou, C., LaBarca, R., Levison, T., 2001. Computer-assisted navigation for knee arthroplasty: intra-operative measurements of alignment and soft tissue balancing. In: *First Annual Meeting of CAOS International, Davos*, p. 114.
- Poulin, F., Amiot, L., 2001. Electromagnetic tracking in the or: accuracy and sources of intervention. In: *Proceedings of CAOS USA*, pp. 233–235.
- Ramey, N.A., Jason Corso, J., William Lau, W., Burschka, D., Hager, G.D., 2004. Real time 3D surface tracking and its applications. In: *Proceedings of Workshop on Real-time 3D Sensors and Their Use (at CVPR 2004)*.
- Reinhardt, H.F., 1996. Neuronavigation: a ten years review. In: Taylor, R., Lavallee, S., Burdea, G., Moegses, R. (Eds.), *Computer-integrated Surgery*. MIT Press, Cambridge, MA, pp. 329–342.
- Reinhardt, H., Meyer, H., Amrein, A., 1988. A computer-assisted device for intraoperative CT-correlated localization of brain tumors. *European Surgical Research* 20, 51–58.
- Scharstein, D., Szeliski, R., 2002. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision* 47 (1–3), 7–42.
- Scholtz, M., Konen, W., Tombrock, S., Adams, L., Fricke, B., During, M.v., Hentsch, A., Heuser, L., Harders, A., 1998. Development of an endoscopic navigating system based on digital image processing. *Journal of Computer Aided Surgery* 3 (3), 134–143.
- Simon, D.A., Jaramaz, B., Blackwell, M., DiGioia, A.M., Morgan, F., Kischell, M.D.E., Colgan, B., Kanade, T., 1997. Development and validation of a navigational guidance system for acetabular implant placement. In: Troccaz, J., Grimson, E., Mosges, R. (Eds.), *Proceedings First Joint Conference of CVRMed and MRCAS*, vol. 1205. Springer, Berlin, pp. 583–592.
- Smith, K.R., Frank, K.J., Bucholz, R.D., 1994. The neurostation – a highly accurate minimally invasive solution to frameless stereotactic neurosurgery. *Computerized Medical Imaging and Graphics* 18, 247–256.
- Stefansic, J., Herline, A., Shyr, Y., Chapman, W., Fitzpatrick, J., Dawant, B., Galloway, R.J., 2000. Registration of physical space to laparoscopic image space for use in minimally invasive hepatic surgery. *IEEE Transaction in Medical Imaging* 19 (10), 1012–1023.
- Taylor, R.H., Paul, H.A., Mittelstadt, B., Cutting, C.B., Hanson, W., Kazanzides, P., Musits, B., Kalvin, A., Kim, Y.-Y., Haddad, B., Khorramabadi, D., Larose, D., 1992. Augmentation of human precision in computer-integrated surgery. *Innovation et Technologie en Biologie et Medicine* 13 (4), 450–459.
- van Hellenmond, G., deKleuver, M., Pavlov, P., 2001. Computer assisted pelvic osteotomies; clinical experience in 25 cases. In: *First Annual Meeting of CAOS International*, p. 123.
- VanderKolk, C., Zinreich, S., Carson, B., Bryan, N., Manson, P., 1992. An interactive 3d-ct surgical localizer for craniofacial surgery. In: Montoya, A. (Ed.), *Craniofacial Surgery*.
- Watanabe, E., Watanabe, T., Manka, S., et al., 1987. Three-dimensional digitizer (neuronavigator): new equipment for computed

- tomography-guided stereotaxic surgery. *Surgical Neurology* 27, 543–547.
- Williams, J.P., Taylor, R.H., Wolff, L.B., 1997. Augmented k -D techniques for accelerated registration and distance measurement of surfaces. In: *Computer Aided Surgery: Computer-Integrated Surgery of the Head and Spine*, pp. 1–21.
- Wu, X., Taylor, R., 2003. A direction space interpolation technique for calibration of electromagnetic surgical navigation systems. In: *Proceedings of the Sixth International Conference on Medical Image Computing and Computer Assisted Intervention*, vol. II. Springer, Berlin, pp. 215–222.
- Zhang, Z., 2000. A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22 (11), 1330–1334.
- Zinreich, S.J., Tebo, S.A., Long, D.M., Mattox, D.E., Brem, H., Loury, M.E., Vander Kolk, C.A., Kennedy, D.W., Koch, W.M., Bryan, R.N., 1993. Frameless stereotaxic intergration of CT imaging data: accuracy and initial applications. *Radiology*, 735–742.