

Registration of Challenging Image Pairs: Initialization, Estimation, and Decision

Gehua Yang, *Student Member, IEEE*, Charles V. Stewart, *Member, IEEE*,
Michal Sofka, *Student Member, IEEE*, and Chia-Ling Tsai, *Member, IEEE*

Abstract—Our goal is an automated 2D-image-pair registration algorithm capable of aligning images taken of a wide variety of natural and man-made scenes as well as many medical images. The algorithm should handle low overlap, substantial orientation and scale differences, large illumination variations, and physical changes in the scene. An important component of this is the ability to automatically reject pairs that have no overlap or have too many differences to be aligned well. We propose a complete algorithm including techniques for initialization, for estimating transformation parameters, and for automatically deciding if an estimate is correct. Keypoints extracted and matched between images are used to generate initial similarity transform estimates, each accurate over a small region. These initial estimates are rank-ordered and tested individually in succession. Each estimate is refined using the Dual-Bootstrap ICP algorithm, driven by matching of multiscale features. A three-part decision criteria, combining measurements of alignment accuracy, stability in the estimate, and consistency in the constraints, determines whether the refined transformation estimate is accepted as correct. Experimental results on a data set of 22 challenging image pairs show that the algorithm effectively aligns 19 of the 22 pairs and rejects 99.8 percent of the misalignments that occur when all possible pairs are tried. The algorithm substantially out-performs algorithms based on keypoint matching alone.

Index Terms—Image registration, feature extraction, iterative closest point, radial lens distortion, decision criteria, keypoint.

1 INTRODUCTION

THIS paper addresses the problem of developing an automated image-pair registration algorithm that can work on a wide variety of image types, scenes, and illumination conditions. Much of this variety is captured in a data set of 22 image pairs we have gathered, some of which are shown in Fig. 1.¹ The set includes image pairs taken of indoor and outdoor scenes, in natural and man-made environments, at different times of day, during different seasons of the year, and using different imaging modalities. It includes image pairs with low overlap (e.g., 2 percent), substantial differences in orientation (90 degrees), and large changes in scale (up to a factor of 6.4). A general-purpose registration algorithm should be able to align each of these image pairs with high accuracy. Moreover, such an algorithm should be able to indicate that two images *cannot* be aligned, either when the images truly do not overlap or when there is insufficient information to determine an accurate, reliable transformation between images. Such a registration algorithm will have

numerous applications, ranging from mosaic construction to change detection and visualization.

Three primary technical challenges must be addressed to solve this problem: initialization, estimation, and decision.

- While automatic initialization is not a significant problem for aligning images in a video sequence or for multimodal registration of images taken from roughly pre-aligned sensors, it is a major concern for more general-purpose registration.
- In combination with initialization, the estimation process must tolerate position, scale, orientation, and illumination differences. Moreover, estimation must accommodate the possibility that there is no relationship between the intensities for a large fraction of the pixels in the two images. For example, in the Summer-Winter pair from Fig. 1, snow on the roofs in winter produces homogeneous intensity regions, whereas these roofs appear as dark, textured regions in the summer image (Fig. 4). Because of this, an effective estimation technique should automatically and adaptively exploit what is consistent between the images.
- Decision criteria are required not only to choose among different estimates obtained from different starting conditions, but also to decide when the images may not be aligned at all. The need for effective decision criteria is particularly acute when handling low overlap and large changes in orientation, illumination, and scale due to the extremely large search space of initial estimates.

1.1 Related Techniques in Registration

The literature on image registration is large; see [6], [25], [54] for reviews. We focus our attention here on four classes of methods that appear most appropriate for the general-purpose registration problem being addressed.

1. Both the data set and our software are available at <http://www.vision.cs.rpi.edu/gdbicp/>.

- G. Yang, C.V. Stewart, and M. Sofka are with the Department of Computer Science, Rensselaer Polytechnic Institute, 110 8th St., Troy, NY 12180. E-mail: {yangg2, stewart, sofka}@cs.rpi.edu.
- C.-L. Tsai is with the Department of Computer Science and Information Engineering, National Chung Cheng University 168, University Rd., Min-Hsiung Chia-Yi, Taiwan, ROC. E-mail: tsaic@cs.ccu.edu.tw.

Manuscript received 15 Aug. 2005; revised 25 May 2006; accepted 23 Jan. 2007; published online 14 Feb. 2007.

Recommended for acceptance by M. Hebert.

For information on obtaining reprints of this article, please send e-mail to tpami@computer.org, and reference IEEECS Log Number TPAMI-0442-0805. Digital Object Identifier no. 10.1109/TPAMI.2007.1116.




Fig. 1. Some images from our 22-image pair data set. The three “Downtown” images produce three of our test pairs.

1.1.1 Feature-Based Methods and ICP

Our first class of prior techniques is feature-based methods and the Iterative Closest Point (ICP) algorithm [5], [12], [13]. Starting from an initial estimate, the ICP algorithm iteratively 1) maps points (features) from the moving image to the fixed image, 2) finds the closest fixed image point for each mapped point, and 3) re-estimates the transformation based on these temporary correspondences. In registration of range data, the sensor data points (together with estimated

normals) are the features. In alignment of intensity images, feature extraction is applied to obtain the points matched in ICP [11], [22]. There are two common problems with feature-based methods and ICP: 1) ICP has a narrow domain of convergence and, therefore, must be initialized relatively accurately and 2) feature extraction can be unreliable and overly sensitive to the choice of parameters and the image content. (We will address both issues.) Finally, several papers [14], [21], [32] have proposed Expectation Maximization (EM) algorithms where multiple correspondences per

feature are simultaneously considered. While this approach improves the domain of convergence of ICP in many cases, it has not been shown to handle the low overlap and high outlier fractions that occur in our challenging data set. We discuss this further in Section 7.

1.1.2 Keypoint Indexing Methods

While ICP correspondences are based on geometric distances, keypoint correspondences are based on appearance similarity measures. Keypoint methods have received growing attention because of their demonstrated ability to tolerate low image overlap and image scale changes. These methods start with keypoint detection and localization, followed by computation of the descriptor that summarizes the image in a region around the keypoint. Existing extraction algorithms are based on approaches ranging from the Laplacian-of-Gaussian operator [31], information theory [28], Harris corners [36], and intensity region stability measures [34]. A detailed comparison can be found in [38]. Region descriptors are based on steerable filters [19], moment invariants [20], shape contexts [3], image gradients [31], and Haar wavelets [8]. They are usually invariant to similarity or affine transformations of the image, as well as linear changes in intensity. A comparison of descriptors can be found in [37].

Keypoint matching has been applied to fundamental matrix estimation [34], multiview matching [43], [52], registration [7], and recognition [16], [17]. Often these techniques use minimal-subset random-sampling methods [7], [34]. (In essence, this combines our initialization and estimation steps.) Statistical tests of randomness based on the number of overall and consistent keypoint matches are then used to decide if the transformation is correct [7] (the decision criteria). While this approach has all of the steps needed for a general-purpose registration system, it has important limitations for the types of image pairs considered in this paper. Most importantly, experiments on our data set show that keypoint indexing and matching methods only produce a small number of correct matches, occasionally none, and sometimes fewer than 10 out of the top 50 matches. Finding an accurate estimate in this case is either impossible or requires an expensive, combinatorial search in candidate match sets. It is possible that more sophisticated keypoint detection, description, and matching methods will make this overall approach viable for such challenging image pairs, but this is not the direction taken in this paper.

1.1.3 Mutual Information

The third class of methods is based on mutual information (MI) [33], [53]. MI registration is primarily designed for multimodal alignment and has been used extensively in medical imaging applications [41], but also in other areas as well [18]. For the current problem, it is less useful than it might first appear. First, algorithms based on MI [18], [40] are only estimation techniques and include neither initialization nor decision techniques. The former is usually not a problem in medical applications and the latter is often left to human judgment. Moreover, MI minimization procedures are quite sensitive to local minima, a crucial concern when aligning low-overlap image pairs.

1.1.4 Direct Methods

A fourth class of techniques involves the direct minimization of intensity differences between images [1]. These methods work by using intensity differences and image gradients to

compute an update to the estimate of the aligning transformation and then applying this update to warp one image on top of the other. Multiresolution techniques are used to accommodate larger initial misalignments [4]. The approach has been extended to handle multisensor data by computing and normalizing derivatives in four directions as the “intensity” at each pixel [27]. These techniques require good initialization, although coarse searches of the translation-only parameter space have sometimes proven effective [42]. Moreover, direct methods have not yet been demonstrated on registration problems involving substantial scale and orientation differences between images. Finally, no decision criteria are associated with these methods.

1.2 Approach

We propose an end-to-end registration algorithm—actually a system of algorithms—for aligning pairs of images using parametric global transformation models, building heavily on existing techniques. We use our data set both to carefully evaluate the main components of our algorithm and to demonstrate its overall effectiveness. The primary novelties of our work are in the construction of the overall algorithm, the design and interrelationship of the components, and the success of the system and the components. As an illustration of this, the alignment process is driven by image features extracted using auto-correlation matrices [44], which have been widely used in Harris corner detectors and its variants [2], [23], [36]. Our implementation produces edge-like as well as corner-like features distributed throughout the images, even in extremely low-contrast regions. By stressing dense coverage of features in images, we ensure that features will nearly always be available, trusting the robustness of the matching and estimation algorithms, as well as the decision criteria to avoid the effects of inconsistencies between images. Thus, the novelty is in tailoring feature extraction for general-purpose registration rather than feature extraction per se.

As mentioned earlier, the three primary components of the system are the initialization algorithm, the estimation technique, and the decision criteria:

- The initialization method uses extraction and matching of keypoints [31]. But, unlike current techniques, it does not attempt to combine matches in any way. Instead, each match is used to generate an initial similarity transformation that is accurate only in a small region surrounding the matched keypoints from each image. There are two reasons for this: first, as suggested above, for challenging image pairs, a relatively small fraction of keypoint matches is correct—too small for effective use of minimal-subset random sampling search techniques. Second, for the keypoint matches that are correct, our growth-and-refinement-based estimation procedure usually aligns the images accurately. Both reasons will be illustrated experimentally.
- The estimation technique starts with the initial local regions and associated transformation estimates, treating each one individually and in succession. The goal for each region and estimate is to expand the region to cover the entire overlap between images while refining the estimate each time the region changes. This can be thought of as keeping the estimate close to the optimum as the problem grows

in complexity. At the same time, as more information is incorporated, it may be possible to switch to a higher-order transformation model that more accurately describes the mapping of larger image regions. These intuitions are realized by generalizing the Dual-Bootstrap ICP algorithm, which we proposed for retinal image registration [48]. Several generalizations are needed to make this work. The most important is the use of generic features, as discussed above, in place of detected blood vessels to drive registration. Other enhanced techniques include bidirectional, across-scale matching between images, region growth in both images, use of a more general model selection technique, and an estimation technique that accounts for variations in feature location uncertainty. It is interesting to note that these generalizations allow the new algorithm to handle some retinal image pairs that the original algorithm could not.

- The decision technique determines if an estimate generated by the Dual-Bootstrap algorithm is indeed a correct alignment of the two images. The technique combines novel decision criteria that measure accuracy, consistency, and stability in the alignment.

The overall procedure takes an ordered sequence of initial estimates (generated from keypoint matching) and applies Dual-Bootstrap ICP to each in succession. Following the convergence of each Dual-Bootstrap estimate, the decision criteria are applied. If the estimate passes these tests, it is accepted as correct and the two images are considered registered. Otherwise, the process continues to the next initial keypoint estimate. It stops with a failure if none succeed. The complete algorithm will be referred to as the Generalized Dual-Bootstrap ICP—GDB-ICP for short.

The remainder of this paper describes the details of GDB-ICP and then evaluates it on our data set of challenging image pairs (Fig. 1). The experimental analysis demonstrates the overall performance of the algorithm and explores the significance of many of the individual components and design decisions. Importantly, this includes testing not only on the image pairs that overlap and therefore should be registered, but also on all possible image pairs, including those with no overlap. The paper concludes with a discussion of the strengths and limitations of GDB-ICP and its relationship to other techniques.

2 INITIALIZATION

Our initialization method is based on Lowe's multiscale keypoint detector and SIFT descriptor [31]. This has proven to be the most effective in the experimental evaluation of [37]. We have used both our own implementation and the one used in the above evaluation with approximately equal success. We have also used affine-invariant multiscale Harris corners [36] and again the differences in results were minor. The results reported here used the implementation from [37].

Here is a summary of the initialization method. Each Lowe keypoint is a local maximum of the magnitude of the Laplacian-of-Gaussian operator in both spatial and scale dimensions. A neighborhood is established at each keypoint, with size determined by the keypoint scale and orientation determined by the local gradient direction. The intensity gradient vectors within this neighborhood are collected in histograms to form a 128-component SIFT descriptor vector. All SIFT vectors are stored in a spatial data structure, one for

each image. Keypoint matching between images occurs by taking the keypoints from one image and using their descriptors to index into the data structure for the other image. The descriptor distance, measured as the Euclidean distance between vectors, is computed for each candidate match. The two closest matches for each descriptor are found and the ratio of the distances to the closest and second closest is calculated. Lowe establishes an upper bound threshold on this ratio of 0.8 and considers only matches below this threshold. Instead, we sort the matches by the ratio and test the top N in order, stopping when GDB-ICP produces an estimate that the decision criteria accept or when all N matches have been tested unsuccessfully. We have found $N = 50$ to be a good, conservative choice. The initial similarity transformation is established from the positions of the two keypoints, the orientations of their dominant intensity gradients, and their scales (Fig. 3). The initial bootstrap region is a square centered at each keypoint location, with half width $30 + 3s_i$, where s_i is the scale of the i th keypoint out of all detected ones in the image. This width setting ensures that there are enough constraints for Dual-Bootstrap to start successfully for keypoints of various scales. One initial region is generated on each of the two images, centered at the keypoint location. Together with the initial transformation, the initial regions are provided as input to the Dual-Bootstrap.

One important consideration is why keypoint descriptors designed only for invariance to linear changes in intensity should be useful in multimodal image registration. Indeed, as we will show experimentally, keypoint matching is the least successful component of our algorithm. On the other hand, our algorithm relies on finding only one correct keypoint match and, in generating this match, the descriptors do not have to match exactly. Instead, they must only be distinct from other matches. Furthermore, *local* intensity differences between analogous keypoints in different modalities are sometimes well-approximated by linear transformations. These two observations explain why matching of Lowe keypoints successfully generates at least a few correct matches on a variety of multimodal pairs.

3 FEATURE EXTRACTION

The Dual-Bootstrap procedure is feature-driven for two reasons. First, matching image features provides direct measurement of the geometric alignment error. This is needed to compute the statistics that drive the growth, model selection, and decision processes. The second reason is motivated by the changes seen between images that must be aligned. Much of the image texture and details may change between images—e.g., snow covering rooftops, leaves missing, textured regions falling into deep shadows—but structural outlines usually remain unchanged (Fig. 4). These outlines, large or small, can be captured by properly extracted features.

As discussed in the introduction, the primary consideration in designing the feature extraction method is not obtaining a complete or a perceptually-significant set of features. It is solely to extract features useful for driving the alignment process. Because of this, our goal is a feature-extraction method that 1) does not depend on thresholds or parameters that must be tuned to individual image content, 2) produces repeatable features, and 3) distributes features throughout the image rather than concentrating them in the image regions of highest contrast.

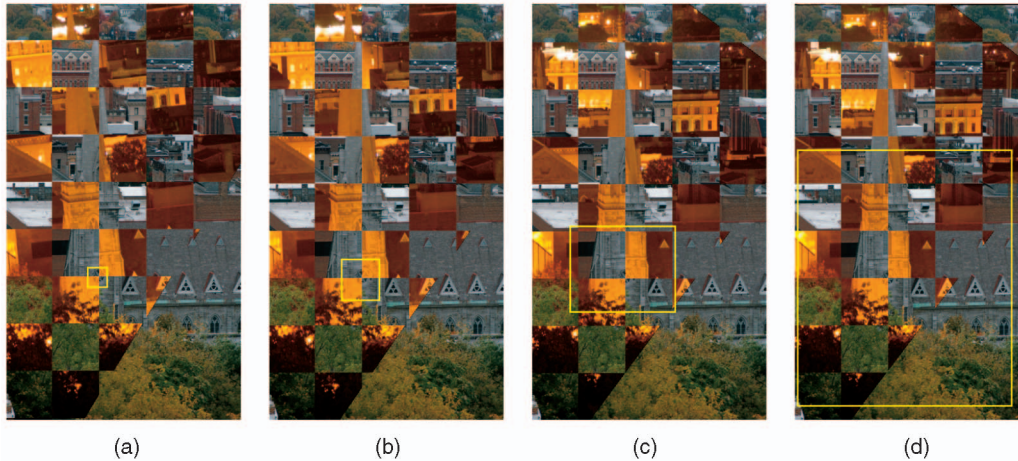


Fig. 2. Example steps of the Dual-Bootstrap growth and refinement process on the Day-Night Summer pair. The checkerboard images above show the alignment results and bootstrap region for iterations 0, 3, 6, and 9. The yellow rectangle outlines the bootstrap region in one image. Within this region, where the computation is concentrated, the alignment is generally accurate. Outside the region, especially for the small regions early in the computation, the alignment tends to be inaccurate. As the bootstrap region expands, more and more of the images are accurately aligned. The final alignment is shown in Fig. 10. (a) Initialization. (b) Third iteration. (c) Sixth iteration. (d) Ninth iteration.

Two different types of features are extracted—corner points and face points. Corner points provide two constraints on registration, whereas face points, because of tangential position ambiguity, provide only one. On the other hand, face points are more densely distributed. We will evaluate the contribution of each feature type to the registration process in the experimental section. Both corners and face points are extracted in scale-space with scale increasing in half-octave steps (e.g., scales $1, \sqrt{2}, 2, 2\sqrt{2}, \dots$), with no combination across scales. Features at all scales are used simultaneously during registration. The following details of the feature-extraction algorithm are presented for a single scale.

We use a single response measure for both feature types. At each pixel location \mathbf{x} at scale σ , the intensity gradient, $\nabla I(\mathbf{x})$, is computed. A weighted neighborhood outer product (auto-correlation) matrix is then computed,

$$\mathbf{M}(\mathbf{x}) = \sum_{\mathbf{y} \in \mathcal{N}(\mathbf{x})} w(\mathbf{x} - \mathbf{y}) (\nabla I(\mathbf{y})) (\nabla I(\mathbf{y}))^\top, \quad (1)$$

where w is a Gaussian weight function with standard deviation σ and the neighborhood size is approximately 3σ . Next, the eigen-decomposition is computed: $\mathbf{M}(\mathbf{x}) = \sum_{i=1,2} \lambda_i(\mathbf{x}) \mathbf{\Gamma}_i(\mathbf{x}) \mathbf{\Gamma}_i(\mathbf{x})^\top$, with $\lambda_1(\mathbf{x}) \leq \lambda_2(\mathbf{x})$. Potential *corners* are at pixels where $\lambda_1(\mathbf{x})/\lambda_2(\mathbf{x}) > t_a$. This criterion is similar to the Harris and Stephens corner detector [23]. Potential *face points* are at pixels for which $\lambda_1(\mathbf{x})/\lambda_2(\mathbf{x}) \leq t_a$. Decision value t_a has been experimentally set to 0.1, although the choice of values is not crucial. A strength is assigned to each point: $m(\mathbf{x}) = \text{trace}(\mathbf{M}(\mathbf{x}))$.

The next step is designed to make the final selection of feature points adaptive to local image contrast. First, a very low threshold, $t_m = 1$, is applied to the strength to eliminate points that are obviously noise. The result is illustrated in Fig. 5a. The next step, local pruning, starts by computing the median μ_m and median absolute deviation (MAD) [49] σ_m of the strength values $m(\mathbf{x})$ in a coarse set of overlapping neighborhoods (30×30 pixels) throughout the image. Then, pixels with $m(\mathbf{x}) < \mu_m + 0.5\sigma_m$ are eliminated from further consideration. As the final preliminary step, nonmaximum suppression is applied at each point—in 2D for potential corner points and in 1D, along the direction of the eigenvector

$\mathbf{\Gamma}_2$, for face points—and surviving point locations are interpolated to subpixel accuracy. Fig. 5b shows an example result.

The final steps are to extract the actual features, generating corners and face points independently. Points surviving the previous step are sorted by $m(\mathbf{x})$ values. The highest strength point is labeled as a feature. Remaining points within a small neighborhood are eliminated. This continues until a maximum number of features—determined from the size of the images—is extracted or the list is exhausted. This produces a set of features which we call the *matchable features* (Fig. 5c). The procedure is repeated starting from the noise thresholding step with the elimination neighborhood doubled in size, half the allowed number of features, and a threshold requiring $m(\mathbf{x}) > 2t_m$. This produces a set of *driving features* (Fig. 5d). As will be explained later, driving features are transformed and matched against matchable features, similarly to [46]. Since driving features must pass stricter criteria than matchable features, it is less likely that a driving feature in one image will be missed as a matchable feature in the other due to random effects.

An example of extracted driving and matchable feature sets at different scales is shown in Fig. 6. Features are spread throughout the image and summarize the local image structure: A corner is placed in regions containing substantial intensity variations in all directions; a face point occurs where the region contains variation in one direction.

4 GDB-ICP ESTIMATION

The estimation step of the Generalized Dual-Bootstrap ICP (GDB-ICP) algorithm starts with an initial similarity transformation generated from keypoint matching, together with the initial bootstrap region surrounding the keypoint location in each image. The algorithm iterates steps of 1) refining the current transformation inside the current bootstrap region R , 2) applying model selection to determine if a more sophisticated model may be used, and 3) expanding the region, growing inversely proportional to the uncertainty of the mapping on the region boundary (Fig. 2). The entire algorithm is outlined in Fig. 7.



Fig. 3. Initial keypoint match and side-by-side alignment for a summer-winter pair. The image region on the right has been scaled by a factor of 2.25 and there are substantial illumination, blurring, and physical changes (snow) between the regions.



Fig. 4. Examples of substantial variations between image regions due to illumination differences, scale differences, and physical changes (snow).

4.1 Notation

The two images are I_p and I_q . The matchable corner and face points are $\mathcal{P}_c = \{\mathbf{p}_c\}$ and $\mathcal{P}_f = \{\mathbf{p}_f\}$ from I_p and $\mathcal{Q}_c = \{\mathbf{q}_c\}$ and $\mathcal{Q}_f = \{\mathbf{q}_f\}$ from I_q . Driving feature sets are subsets of \mathcal{P}_c , \mathcal{P}_f , \mathcal{Q}_c , and \mathcal{Q}_f . Points from all scales are combined to form these sets. Abusing notation, \mathbf{p} and \mathbf{q} represent both the feature and its location. Each feature has associated with it the scale, s , at which it was detected and each face point has a normal direction, η .

The *forward* transformation of point location \mathbf{x} from I_p onto I_q is $\mathbf{T}(\mathbf{x}; \theta^{pq})$, where θ^{pq} is the parameter vector to be




Fig. 5. Intermediate results of the feature extraction process: (a) Initial corners (shown with circles) and face points (shown with line segments along the eigenvector corresponding to the larger eigenvalue) after applying the threshold t_m (only one of every three is shown in this display because they are so dense), (b) corners and face points after adaptive local pruning, (c) “matchable” corners and face points, and (d) “driving” corners and face points.

estimated. An estimate is $\hat{\theta}^{pq}$ and its covariance matrix is $\hat{\Sigma}_{\theta^{pq}}$. The *backward* transformation from I_q onto I_p is $\mathbf{T}(\mathbf{x}; \theta^{qp})$, with an estimate $\hat{\theta}^{qp}$ and covariance estimate $\hat{\Sigma}_{\theta^{qp}}$. Finally, the regions over which the transformation is being estimated are called the “bootstrap” regions and are denoted by R_p on image I_p and R_q on image I_q . Each is defined as an axis-aligned rectangular box on its own image.

4.2 Matching within the Bootstrap Region

The transformation is refined within current bootstrap regions R_p and R_q , ignoring everything else in the two images. Recall that, in standard ICP, the current transformation is used to generate a new set of correspondences; in turn, these correspondences are used to generate a new transformation and this process iterates. By contrast, GDB-ICP proceeds to model selection and region growing before selecting a new set of matches.

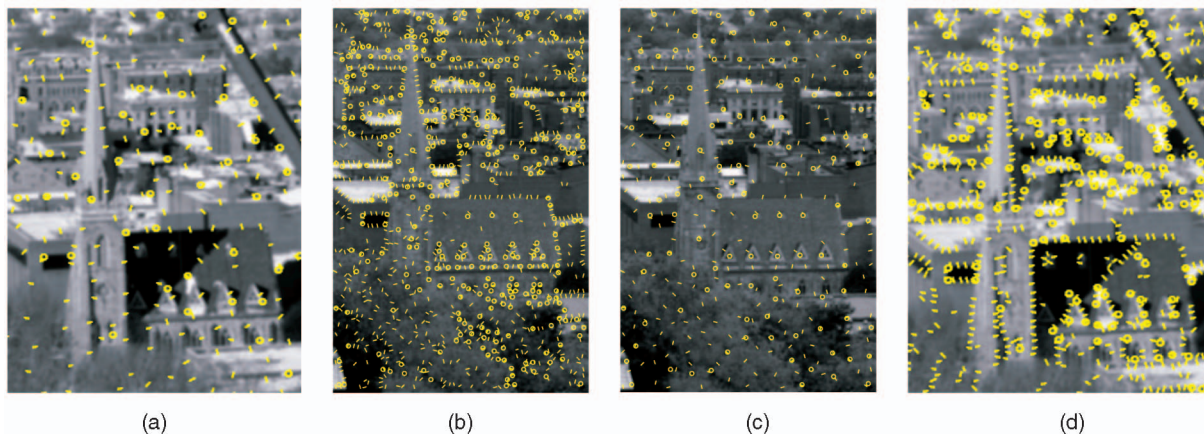


Fig. 6. Example of “matchable” and “driving” features and bidirectional matching. In matching, “driving” features from the (a) (cropped) winter image are mapped onto the (b) summer image and matched against the “matchable” features. Symmetrically, “driving” features from the (c) summer image are mapped onto the (d) winter image and matched against the “matchable” features. Although not illustrated in this figure, matching may occur between features at different scale-space scales.

- 1) Extract keypoints and features from images I_p and I_q (Section 2 and 3).
- 2) Match keypoints from I_p to I_q and rank-order the matches (Section 2).
- 3) **Do**
 - a) Choose next keypoint match and generate the initial transformations and bootstrap regions (Section 2).
 - b) **Repeat**
 - i) Use the current model and parameter estimates $\hat{\theta}^{pq}$ and $\hat{\theta}^{qp}$ to generate the forward and backward match sets \mathcal{C}_c^{pq} , \mathcal{C}_f^{pq} , \mathcal{C}_c^{qp} and \mathcal{C}_f^{qp} (Section 4.2).
 - ii) Use the match sets to re-estimate the forward and backward transformation parameter estimates, $\hat{\theta}^{pq}$ and $\hat{\theta}^{qp}$, and their covariance matrices, $\hat{\Sigma}_{\theta^{pq}}$ and $\hat{\Sigma}_{\theta^{qp}}$, for the current model and remaining higher order models (Section 4.3).
 - iii) Use the estimates and covariance matrices to select the model for the next iteration (Section 4.4).
 - iv) Grow the bootstrap regions in each image separately using the selected model, parameter estimates, and covariance matrices (Section 4.5).
 - Until** region growth, model selection and parameter estimation have converged.
 - c) Apply the decision criteria using the final model, parameter estimates, covariance matrices, and face correspondence sets. If all pass, accept parameter estimates $\hat{\theta}^{pq}$ and $\hat{\theta}^{qp}$ and terminate with success (Section 5).
- 4) **While** fewer than N keypoint matches have been tried.
- 5) Terminate with failure.

Fig. 7. Outline of the generalized dual-bootstrap algorithm.

GDB-ICP uses bi-directional matching. This provides more constraints and helps to produce more numerically stable estimates, especially for small bootstrap regions. A driving feature \mathbf{p} is mapped from I_p to I_q , producing $\mathbf{p}' = \mathbf{T}(\mathbf{p}; \hat{\theta}^{pq})$. The three closest matchable features (of the same type) to \mathbf{p}' are found in I_q and the best matching feature, \mathbf{q} , is chosen from among these three based on a similarity measure described below. The correspondence pair (\mathbf{p}, \mathbf{q}) is added to match sets \mathcal{C}_c^{pq} (corners) or \mathcal{C}_f^{pq} (faces). Reversing order, the pair (\mathbf{q}, \mathbf{p}) is also added to either \mathcal{C}_c^{qp} or \mathcal{C}_f^{qp} . The same procedure is used in the opposite direction to generate correspondences between driving features from I_q and matchable features from I_p . Once these matches are generated, the forward and backward transformation estimates $\hat{\theta}^{pq}$ and $\hat{\theta}^{qp}$ are calculated. Since these use the same set of constraints, just with the feature roles reversed, the two transformation estimates are close to being inverses of each other (typical root mean square error of 0.1 pixels or less). The remainder of the discussion focuses on the calculation of $\hat{\theta}^{pq}$ using \mathcal{C}_c^{pq} and \mathcal{C}_f^{pq} .

A similarity measure is used both in choosing between matches for a feature and in weighting the chosen match during estimation. For corners, it depends only on the feature scales, but, for face points, it depends on orientations as well. For a feature point \mathbf{p} , let s_p be the feature scale after the transformation is applied. For a face point, let η_p be the transformed normal direction. The similarity measures for a prospective match with \mathbf{q} , with scale s_q and (for a face point) normal η_q , are

$$w_c = \min(s_{p'}/s_q, s_q/s_{p'}) \quad \text{and} \\ w_f = \min(s_{p'}/s_q, s_q/s_{p'}) \cdot |\eta_p^\top \eta_q|$$

for corners and face points, respectively. This biases the selection toward features at similar (mapped) scales and orientations and allows for contrast reversals in face point matches as well.

4.3 Estimation

Estimation is applied to the current model and, as the basis for model section (Section 4.4), to higher-order models under consideration. This section describes estimation for a single model.

Before defining the transformation estimate objective function, we need to define the error distances. For corner points, these are Euclidean distances, whereas, for face points, these are normal distances:

$$d_c(\mathbf{p}, \mathbf{q}; \theta^{pq}) = \|\mathbf{T}(\mathbf{p}; \theta^{pq}) - \mathbf{q}\|/s_q \quad \text{and} \\ d_f(\mathbf{p}, \mathbf{q}; \theta^{pq}) = |(\mathbf{T}(\mathbf{p}; \theta^{pq}) - \mathbf{q})^\top \eta_q|/s_q.$$

In each case, the distance is normalized by the scale at which the feature is detected, reflecting the fact that feature location uncertainty increases with increasing scale. This makes distances of features at different scales approximately comparable.

Combining the foregoing, the objective function for estimating transformation parameters θ^{pq} from a fixed set of correspondences is

$$E(\theta^{pq}; \mathcal{C}_c^{pq}, \mathcal{C}_f^{pq}) = \sum_{(\mathbf{p}_i, \mathbf{q}_i) \in \mathcal{C}_c^{pq}} w_{c;i} \rho(d_c(\mathbf{p}_i, \mathbf{q}_i; \theta^{pq})/\sigma_c) \\ + \sum_{(\mathbf{p}_i, \mathbf{q}_i) \in \mathcal{C}_f^{pq}} w_{f;i} \rho(d_f(\mathbf{p}_i, \mathbf{q}_i; \theta^{pq})/\sigma_f), \quad (2)$$

where $\rho(\cdot)$ is the Beaton-Tukey [26], [35], [49] robust loss function

$$\rho(u) = \begin{cases} \frac{a^2}{6} \left[1 - \left(1 - \left(\frac{u}{a} \right)^2 \right)^3 \right], & |u| \leq a \\ \frac{a^2}{6}, & |u| > a. \end{cases} \quad (3)$$

Following standard usage, the constant a is set to 4, which means that normalized alignment error distances beyond 4σ have a fixed cost. The parameters σ_c and σ_f are the robust alignment error scales (standard deviations) for the normalized distances of corners and faces.

Objective function (2) is minimized using the Iteratively-Reweighted Least-Squares (IRLS) technique from the robust statistics literature [35], [49], which alternates computation of 1) the distance-based weight $w_{d,i}$ for each correspondence, i , based on fixed transformation parameters with 2) weighted least-squares reestimation of the parameters from

$$F(\theta^{pq}; \mathcal{C}_c^{pq}, \mathcal{C}_f^{pq}) = \sum_{(\mathbf{p}_i, \mathbf{q}_i) \in \mathcal{C}_c^{pq}} w_{d,i} w_{c,i} d_c^2(\mathbf{p}_i, \mathbf{q}_i; \theta^{pq}) + \sum_{(\mathbf{p}_i, \mathbf{q}_i) \in \mathcal{C}_f^{pq}} w_{d,i} w_{f,i} d_f^2(\mathbf{p}_i, \mathbf{q}_i; \theta^{pq}). \quad (4)$$

The distance-based robust weight factor for corners is $w_{d,i} = w(d_c(\mathbf{p}_i, \mathbf{q}_i; \theta^{pq})/\sigma_c)/\sigma_c^2$, where $w(\cdot)$, derived from the Beaton-Tukey robust loss function, is

$$w(u) = \begin{cases} 1 - \left(\frac{u}{a} \right)^2, & |u| \leq a \\ 0, & |u| > a. \end{cases}$$

An analogous computation produces the weights for face points. Normalization factors $1/\sigma_c^2$ (for corners) and $1/\sigma_f^2$ (for face points) make corners and face points comparable.

The robust standard deviation, σ_c , for corner matches is recomputed once (per correspondence set and Dual-Bootstrap iteration) from the weights and current transformation estimate as

$$(\sigma_c)^2 = \sum_{(\mathbf{p}_i, \mathbf{q}_i) \in \mathcal{C}_c^{pq}} w_{d,i} w_{c,i} d_c^2(\mathbf{p}_i, \mathbf{q}_i; \hat{\theta}^{pq}) / \sum_{(\mathbf{p}_i, \mathbf{q}_i) \in \mathcal{C}_c^{pq}} w_{d,i} w_{c,i},$$

with a similar computation for face matches. At the start of the Dual-Bootstrap procedure for a given initial transformation, the MUSE algorithm [39] is used to estimate σ_c and σ_f from the first set of matches, since weights are unavailable.

Finally, the computation of the weighted least-squares estimates from (4) and the associated covariance matrix of the parameter estimates, which is needed for the region growth and the decision criteria, use standard techniques. In particular, for the transformation models that are unconstrained and linear in their parameters, $\hat{\theta}$ is obtained in closed-form using linear weighted least-squares. The covariance matrix, $\hat{\Sigma}_{\theta^{pq}}$, of the estimate is the inverse of the Hessian of (4). Usually, this must be multiplied by a noise variance term, but this is already built into the distance-based robust weights $w_{d,i}$. For constrained or nonlinear models, such as planar homography or planar homography plus radial-lens distortion terms, Levenberg-Marquardt is used, with the pseudo-inverse of the Hessian giving the covariance matrix, $\hat{\Sigma}_{\theta^{pq}}$ (see [24, Chapter 4]).

4.4 Model Hierarchy and Model Selection

The goal of model selection is to select the model from a set (a hierarchy) of transformation models that best describes the

alignment in the current bootstrap region. As the region expands, model selection is applied to choose between the model used for the previous bootstrap region and the remaining, high-order models. Model selection must be done carefully. Switching to a higher-order model too early, especially when the region is small and there are insufficient constraints, may lead to overfittings and distortions in the estimate. Switching too late causes an increase in mapping errors and results in an increase in mismatches. In either case, incorrect model selection may drive the estimate into a local minimum representing an incorrect alignment.

Two different model hierarchies are used in GDB-ICP. One, used for retinal images, is a hierarchy moving from similarity to a reduced quadratic to a quadratic model (see [48] for details). The second, used for natural images, is a hierarchy of similarity, affine, homography, and homography plus radial lens distortion (HRD). The HRD model is defined as

$$\mathbf{T}(\mathbf{p}; \theta) = \mathbf{D}(\mathbf{T}_H(\mathbf{D}(\mathbf{p}; k_p); \mathbf{h}); k_q),$$

where $\theta^\top = (\mathbf{h}^\top, k_p, k_q)$, $\mathbf{T}_H(\mathbf{x}; \mathbf{h})$ is the usual planar homography (\mathbf{h} is a 9-component vector formed from the 3×3 homography matrix) and $\mathbf{D}(\mathbf{x}; k) = (1 + k\|\mathbf{x} - \mathbf{x}_0\|^2)\mathbf{x}$ is the radial distortion function, given image center \mathbf{x}_0 , assumed to be the center of the pixel array. This model is important for accurate alignment of digital photographs taken with off-the-shelf cameras.

Model selection techniques have been studied extensively in the literature [9], [10], [29], [50]. In our earlier work [48], we used a Bayesian technique derived in [9] that depends on computing the determinant of the parameter-estimate covariance matrix, $\hat{\Sigma}_{\theta^{pq}}$. For homographies, this is problematic because $\hat{\Sigma}_{\theta^{pq}}$ is not full rank. Rather than developing an appropriate projection onto a full-rank covariance, we have replaced the Bayesian criteria with a modified version of Akaike Information Criteria (AIC), derived from the Kullback-Leibler measure, and found it to be quite effective [10, Chapter 2]. Using our robust objective function (2) and taking advantage of having transformation estimates and match sets available in both directions, a second-order Akaike Information Criteria may be written as

$$\begin{aligned} \mathcal{I} = & -2 \left[|\mathcal{C}_c^{pq}| \log(\sigma_c^{pq}) + |\mathcal{C}_f^{pq}| \log(\sigma_f^{pq}) + E(\hat{\theta}^{pq}; \mathcal{C}_c^{pq}, \mathcal{C}_f^{pq}) \right] \\ & -2 \left[|\mathcal{C}_c^{qp}| \log(\sigma_c^{qp}) + |\mathcal{C}_f^{qp}| \log(\sigma_f^{qp}) + E(\hat{\theta}^{qp}; \mathcal{C}_c^{qp}, \mathcal{C}_f^{qp}) \right] \\ & + 2nl/(n-l-1), \end{aligned} \quad (5)$$

where l is the degrees of freedom in the current model, $n = 2|\mathcal{C}_c^{pq}| + 2|\mathcal{C}_c^{qp}| + |\mathcal{C}_f^{pq}| + |\mathcal{C}_f^{qp}|$ is the effective number of constraints (each corner match provides two constraints, while each face point match provides one), and the term $2nl/(n-l-1)$ adjusts for small sample bias [10, p. 51].

Expression (5) is evaluated for each candidate model using a fixed set of matches found using the transformation estimate of the best model from the previous Dual-Bootstrap iteration. The final objective function value of (2) (after IRLS converges) is used for each model to evaluate expression (5). The model that minimizes (5) is chosen as the current model and its estimated parameters become the current parameters. Model selection is turned off once the selection procedure reaches the highest-order model.

4.5 Region Growth

Region growth, illustrated in Fig. 2, is based on the uncertainty in the transformation estimate, represented by the covariance matrix $\hat{\Sigma}_{\theta^{pq}}$. Expansion of the axis-aligned rectangle representing the current bootstrap region is inversely proportional to the transfer error—the error in applying the estimated transformation to points on the boundary of the bootstrap region. The following is a summary of the details of this procedure taken from [48]. Subsequently, a simple modification is given to make the algorithm more effective in registering image pairs with large scale variations.

Let the center of the bootstrap region be \mathbf{y}_0 , let a point location centered on one of the four sides of the region be \mathbf{y} , and let $\boldsymbol{\eta}_y = (\mathbf{y} - \mathbf{y}_0) / \|\mathbf{y} - \mathbf{y}_0\|$ be the outward-pointing normal to the rectangle. The mapping error covariance at the mapped point $\mathbf{y}' = \mathbf{T}(\mathbf{y}; \hat{\theta}^{pq})$ is computed from the Jacobian of the mapping, $\mathbf{J}_y = \partial \mathbf{T}(\mathbf{x}; \boldsymbol{\theta}) / \partial \boldsymbol{\theta}$ evaluated at $\mathbf{x} = \mathbf{y}$ and $\boldsymbol{\theta} = \hat{\theta}^{pq}$, together with the covariance of the transformation parameters

$$\Sigma_{\mathbf{y}'} = \mathbf{J}_y \hat{\Sigma}_{\theta^{pq}} \mathbf{J}_y^\top. \quad (6)$$

The outward growth rate is inversely proportional to the error variance in the mapped outward normal direction, $\boldsymbol{\eta}_{\mathbf{y}'}$,

$$\delta_y = \beta \frac{(\mathbf{y} - \mathbf{y}_0)^\top \boldsymbol{\eta}_y}{\max(1, \boldsymbol{\eta}_{\mathbf{y}'}^\top \Sigma_{\mathbf{y}'} \boldsymbol{\eta}_{\mathbf{y}'})},$$

where $\beta = 2.0$ is a fixed constant. The new center of the side is given by $\hat{\mathbf{y}} = \mathbf{y}_0 + (1 + \delta_y)(\mathbf{y} - \mathbf{y}_0)$. The new region is obtained after all of the side centers are updated with the above method.

The extension made here is to form and grow bootstrap regions independently in each image, I_p and I_q . Each of the two initial regions is determined from the corresponding keypoint locations and scales, as described in Section 2. The above procedure is applied separately for the two regions at each Dual Bootstrap iteration separately. These regions are implicitly kept relatively consistent through the use of bi-directional matching, which keeps the transformations close to being inverses of each other and the covariance matrices commensurate with each other. Keeping separate regions in the two images is important for handling large scale differences between images.

5 DECISION CRITERIA

Once the GDB-ICP refinement procedure just described expands to cover the apparent overlap between images (based on the estimated transformation) and the refinement process has converged, the final alignment is tested for correctness. If this confirms that the transformation is correct, the images are considered to be aligned and the algorithm stops. Otherwise, GDB-ICP is restarted on the next keypoint match.

Three tests—accuracy, stability, and consistency—form the decision criteria. The tests are applied in each direction using the final match sets. A transformation that passes all three tests in both directions is accepted as correct.

Accuracy is measured as the weighted average error $\zeta_e(\hat{\theta}, \mathcal{C}_f)$, computed on the final face matches, \mathcal{C}_f . Face points are used because their positions (along the normal direction)




Fig. 8. Example histograms of the orientation difference of a set of face point matches from (a) correct and (b) incorrect alignments using the Summer-Winter pair.

are more accurate than corner points. Using the measures introduced above, accuracy is

$$\zeta_e(\hat{\theta}^{pq}, \mathcal{C}_f) = \left(\sum_{(\mathbf{p}_i, \mathbf{q}_i) \in \mathcal{C}_f} w_{f,i} w_{d,i} d_f(\mathbf{p}_i, \mathbf{q}_i; \hat{\theta}^{pq}) \right) / \left(\sum_{(\mathbf{p}_i, \mathbf{q}_i) \in \mathcal{C}_f} w_{f,i} w_{d,i} \right). \quad (7)$$

Stability is measured by the error covariance—the mapping transfer error introduced in the context of region growth in Section 4.5. To check this, points are uniformly sampled in the overlap area between aligned images. For each sample point \mathbf{y}_i , the mapping error covariance $\Sigma_{\mathbf{y}_i}$ is computed from (6). The overall measure is $\zeta_t(\hat{\theta}, \hat{\Sigma}_{\theta}) = \max_i \text{trace}(\Sigma_{\mathbf{y}_i})$. This is particular effective at avoiding incorrect low-overlap transformations.

The consistency measure is derived from the orientation differences of the face point match set \mathcal{C}_f after the application of the transformation estimate $\hat{\theta}$. These differences, measured in absolute angle difference, are put into a histogram $h(\hat{\theta}, \mathcal{C}_f)$ of the range $[0, \pi/2]$. (Absolute angle differences greater than $\pi/2$ are subtracted from π , effectively accommodating intensity reversals.) If the transformation is incorrect, this angle difference will tend toward being uniformly distributed, whereas, if the images are well-aligned, the histogram will tend to have a strong peak near 0 degrees (Fig. 8). The consistency measure is based on the Bhattacharyya measure against an exponential distribution. The probability density function of an exponential distribution is $f(x; \lambda) = \lambda e^{-\lambda x}$ for $x \geq 0$. We use $\lambda = 4.7$, which dictates that 70 percent of the face point matches have orientation differences no greater than 10 degrees. This exponential distribution, denoted as e , is represented as a second histogram. Then, the consistency measure $\zeta_c(\hat{\theta}, \mathcal{C}_f)$ is

$$\zeta_c(\hat{\theta}, \mathcal{C}_f) = 1 - \sum_i \sqrt{h_i(\hat{\theta}^{pq}, \mathcal{C}_f) e_i}, \quad (8)$$

where i indexes the histogram bins.

To make a decision with these measures— ζ_e , ζ_t , and ζ_c —lower and upper thresholds are introduced for each: $Z_e^L \leq Z_e^H$ for ζ_e , $Z_t^L \leq Z_t^H$ for ζ_t , and $Z_c^L \leq Z_c^H$ for ζ_c . When $\zeta_e \leq Z_e^L$, $\zeta_t \leq Z_t^L$, and $\zeta_c \leq Z_c^L$, the transformation estimate is accepted as correct. When $\zeta_e > Z_e^H$, $\zeta_t > Z_t^H$, or $\zeta_c > Z_c^H$, the transformation is rejected. Otherwise, the transformation is saved. If all initial transformations have been tested and none have been accepted, the saved transformation with the minimum value of alignment error ζ_e is accepted. If there are no saved transformations, the algorithm rejects the image




Fig. 9. (a) shows a checker mosaic of the Brugge Square pair using the homography model in which there are places, notably inside the boxes, that have misalignments. Column (b) shows zoom-in of the boxes. Column (c) shows the same area after using the Homography with Radial lens Distortion (HRD) model. The window frames and letters are now well aligned.

pair, indicating that it cannot be aligned. Transformations tend to fall into the “saved” category for image pairs that involve significant changes or that cannot be precisely aligned using the final transformation model.

These thresholds are fixed at $Z_e^L = 1$, $Z_e^H = 2$, $Z_t^L = 0.3$, $Z_t^H = 1$, $Z_c^L = 0.09$, and $Z_c^H = 0.2$ for all experiments here. For efficiency, the algorithm also applies a set of higher thresholds, starting after the third Dual Bootstrap iteration when the estimate has begun to stabilize, in order to identify and terminate estimates that are clearly wrong. We refer to this step as the *early termination criteria*.

6 EXPERIMENTS

This section presents experiments designed to illustrate the overall performance of GDB-ICP (Section 6.2), compare it to minimal-subset random sampling methods (Section 6.3), and analyze in detail the most important aspects of the algorithm. The focus of the latter is on the success of growth and refinement (Section 6.4), the choices of features and matching criteria (Section 6.5), and the effectiveness of the final decision criteria (Section 6.6)—the newest aspects of the algorithm. See [48] for analysis that shows the significance of the Dual-Bootstrap refinement, growth, and model selection procedures in the context of retinal image registration.

6.1 Data Set

All experiments use the data set of 22 image pairs discussed in Section 1. This set was constructed from our own digital photographs, from pairs found on the Web, and from challenging pairs suggested by colleagues. Many easier pairs have been left out in order to keep the tests manageable. As an example of this, we included one pair, with 2 percent overlap (the “Dashpoint” pair here), from the test suite of [8] (GDB-ICP registers all overlapping pairs from this suite). On the other hand, some types of pairs, such as PET-CT images, which have no common geometric structure, have been purposely left out. We discuss this more in Section 7. The results are clearly conditioned on the data set, but the range of challenging pairs shown should be suggestive of the broad effectiveness of our algorithm. In order to allow the community to test GDB-ICP beyond the experiments presented here, an executable version of the software has been posted on the Web.

The images range in size from 676×280 to $2,500 \times 2,500$. Image pairs overlap as little as 2 percent, differ in scale by a factor as high as 6.4, and differ in orientation by as much as

90 degrees. Five pairs are multimodal (retina angiogram versus red-free photograph, two infrared versus video airport scenes, pan-chromatic and infrared satellite images, and proton density versus T1 weighted brain MRI slices). Four pairs involve substantial illumination changes and two other pairs are of different seasons. The selection of scenes includes aerial, urban, landscape, indoor, and medical. On the Melanoma and EO-IR 1 pairs, the intensity of one image is negated before keypoint generation since SIFT is not invariant to intensity reversal. Finally, the retinal images involve quadratic transformations, whereas the others involve the use of the homography or the homography plus radial lens distortion (HRD) models. The choice of final model is specified by a command-line argument. All other parameter settings are fixed for these experiments.

6.2 Overall Results

GDB-ICP successfully aligned 19 of all 22 image pairs in our data set with alignment error less than a pixel. Success is defined here as no visible misalignments between homologous structures following application of the transformation, as judged independently by a graduate student who is not one of the authors. The successful transformations, one for each pair, are labeled as “verified” transformations to be used in subsequent experiments. Example alignments are shown in Fig. 10; complete results are posted at our Web site, including animations. Interestingly, for pairs “Brugge,” “Brugge Square,” and “Brussels,” the 10 degree-of-freedom “Homography plus Radial lens Distortion” (HRD) model eliminated small, but visible misalignments produced by using only a homography (see Fig. 9 for details).

Table 2 shows, for each pair, the index number of the first keypoint match in the rank ordering for which the algorithm succeeded, the index of the same successful keypoint match among only those consistent with the verified transformation, the final alignment error, and the chosen transformation model. A consistent keypoint match is somewhat arbitrarily defined to have a location error of less than six pixels, a scale ratio within the interval 0.8 to 1.25 (one step in scale space), and an orientation difference of 15 degrees, all computed following application of the verified transformation. Intuitively, these are matches that appear to be geometrically close to correct. The remaining keypoint matches are labeled “inconsistent.” As can be seen from the table, in most cases, a consistent match appears among the first five in the rank ordering and, in 15 cases, GDB-ICP successfully refined this initial transformation to a verified final transformation.

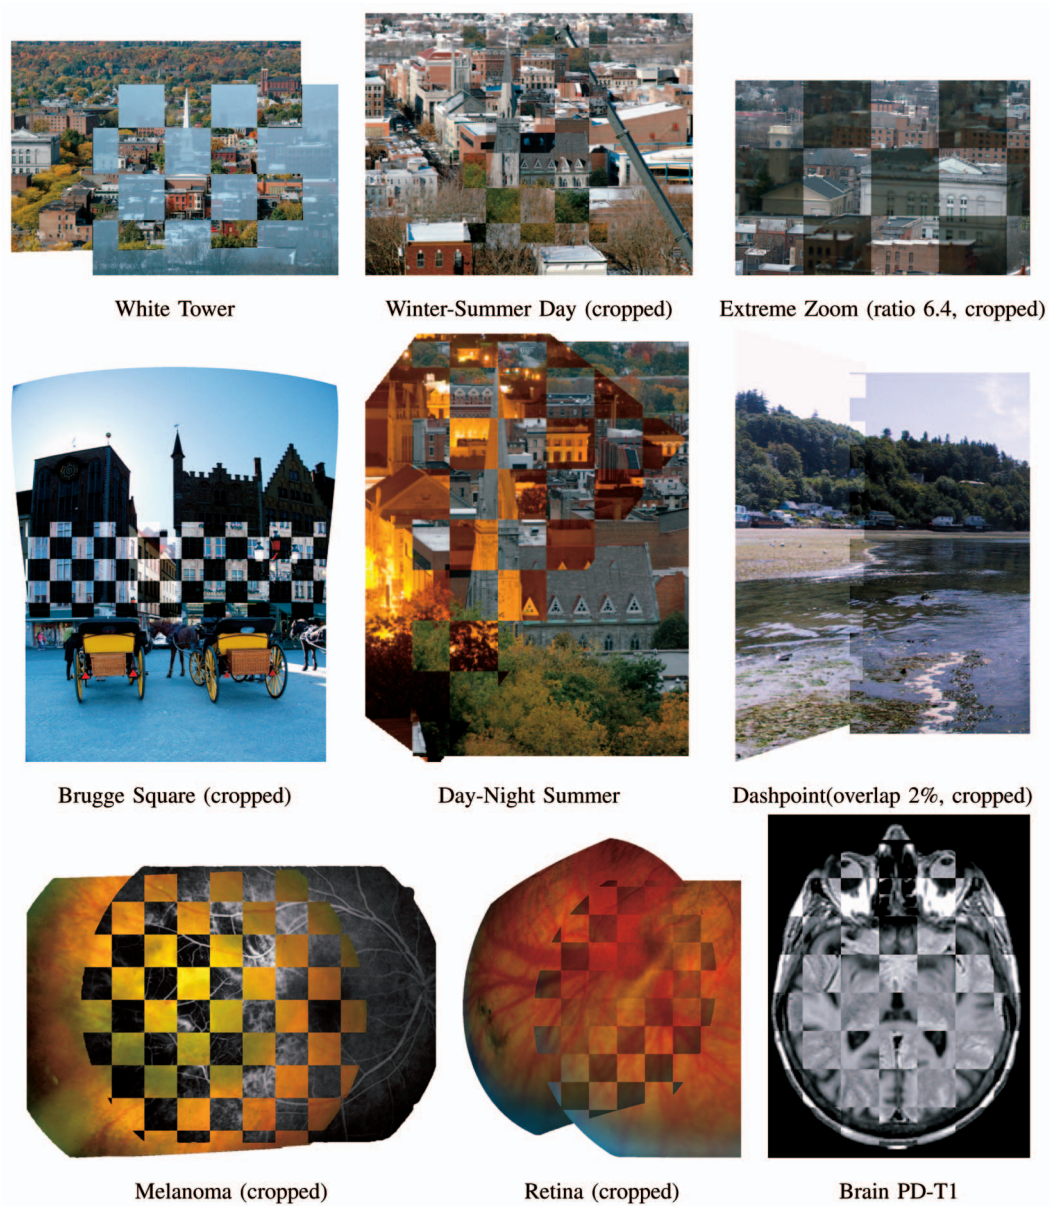


Fig. 10. Final alignment checkerboard images.

GDB-ICP failed for three pairs. In each case, manual specification of three initial correspondences in a small initial region and computation of an initial estimate of an affine transformation followed by application of the Dual-Bootstrap growth and refinement procedure led to a verified transformation. This indicates that the failures are caused by keypoint detection and matching, by keypoint-based initialization, or by the early stages in the Dual-Bootstrap growth and refinement procedures. In one case in particular—Capital Region—the projective distortions are too severe to be handled starting from a local similarity transformation.

Image sizes and timing results are summarized in Table 1. Clearly, algorithm speed is mostly affected by image size and matching difficulty. The failures and the image pairs which require testing of all 50 keypoint matches (because no match provided results below the lower decision thresholds) are the only ones other than the huge “Satellite” pair requiring more than a minute.

6.3 Comparison to Keypoint Matching Algorithms

As one indication of the significance of these results, the publicly available code for the Autostitch keypoint matching algorithm [7] (with default parameters) produced five alignments (“Boston,” “Boston Library,” “Eiffel,” “Brugge Square,” and “Brussels”). The latter three have visible misalignments, partly due to the fact that the homography is insufficient for these pairs. On the other 17 pairs, Autostitch failed altogether. We obtained slightly better results with our own implementation using RANSAC and other random-sampling-based algorithms [39], [51], registering Grand Canyon 1 and White board pairs. The failures are due to both the small number and the small fraction of consistent keypoint matches, as shown in the last three columns of Table 2.

6.4 Success of the Growth and Refinement Procedure

The following experiment shows the effectiveness of starting from individual keypoint matches rather than combining

TABLE 1
Timing Results in Seconds

Image Pair Name	Image Dimension	On the successful keypoint match		Grand Total	
		# of Iterations	Time in sec.	# of Keypoint matches tried	Time in sec.
Boston	1712×1368	14	12	1	12
Boston Library	1504×1000	13	6.5	1	6.5
Brain T1 to PD	512×512	12	2.0	5	8
Brugge Square	1712×1368	17	58	1	58
Brugge Tower	1712×1368	20	28	1	28
Brussels	1712×1368	17	31	1	31
Capital Region	1712×1368	NA	NA	50	1225
Day-Night	1000×1504	24	30	2	43
Dashpoint	2048×1536	12	10	1	10
Eiffel	1712×1368	16	22	1	22
EO- IR 1	300×236	15	6.8	50	107
EO- IR 2	676×280	NA	NA	50	51
Extreme Zoom	1504×1000	25	34	50	761
Grand Canyon 1	1184×780	15	11	1	11
Grand Canyon 2	900×568	16	26	50	592
Melanoma	1156×880	14	13	50	126
Retina	1600×1200	17	9.9	5	23
Satellite	2878×2878	18	67	1	67
Whiteboard	1504×1000	19	12	1	12
White Tower	1504×1000	18	21	2	56
Winter-Summer	1504×1000	19	37	2	51
Winter Day- Summer Night	1504×1000	NA	NA	50	300

The first two columns are the image-pair name and the dimension of the larger image. The next two columns are the number of iterations and the time that the Dual Bootstrap growth-and-refinement procedure spent on the keypoint match that led to the successful alignment. The last two columns are the total number of keypoint matches tried and the total time used before GDB-ICP terminates. The performance is measured on a Pentium 4 3.2 GHz PC with 2 GB memory.

them, as in a random-sampling approach. We use the top 50 keypoint matches of the 19 pairs that GDB-ICP aligned. The GDB-ICP estimation process is applied to each of these keypoint matches, without any decision criteria. The resulting “test” transformation estimate is then compared to the verified transformation. Those that agree to within an average distance of less than two pixels in the overlap region between images are considered correct.

Among the 19 pairs of images, there are 781 keypoint matches in total (some pairs have fewer than 50 matches), 489 are “consistent” and 292 “inconsistent.” Among the 489 consistent ones, 397 led to correct final transformations, while 21 of 292 labeled “inconsistent” led to correct final transformations, resulting in a total of 418 correct alignments. Examination of the 21 shows that the estimation procedure recovered from initial location errors as high as 12 pixels and orientation differences as much as 18 degrees.

To interpret the significance of these results, based on a probability of $P_a = 397/489 \approx 0.81$ of succeeding from a consistent keypoint match, the overall probability of GDB-ICP producing a correct alignment given n consistent matches is $1 - (1 - P_a)^n$, which is 99.3 percent when there are just $n = 3$ consistent matches. By contrast, minimal subset random sampling techniques require four matches just to instantiate a transformation. Clearly, GDB-ICP can succeed despite an extremely small number of keypoint matches.

6.5 Choice of Features and Matching

The next set of experiments evaluates several variations on the choice of features, the scale of the features, and the directionality of matching. This is important to show the influence of these design decisions on the performance of the overall algorithm. These experiments show that face points and bidirectional matching are important and indispensable.

TABLE 2
Summary Statistics on All Pairs from Our Data Set

Image pair name	With GDB-ICP				keypoint statistics		
	1st success in all keypoint matches	1st success in only the consistent matches	Final alignment error	Final model	Total # keypoint matches	position consistent within 6 pixels	ratio
Boston	1	1	0.36	HRD	641	492	76%
Boston Library	1	1	0.32	HRD	334	144	35%
Brain T1 to PD	5	1	0.98	H	18	9	50%
Brugge Square	1	1	0.43	HRD	421	272	65%
Brugge Tower	1	1	0.49	HRD	558	369	66%
Brussels	1	1	0.35	HRD	1435	1180	82%
Capital Region	NA	NA	NA	HRD	142	6	4%
Day-Night (Summer)	2	1	0.62	HRD	99	30	30%
Dashpoint	1	1	0.30	HRD	54	17	31%
Eiffel	1	1	0.30	HRD	456	216	47%
EO - IR 1	3	2	1.59	HRD	13	4	31%
EO - IR 2	NA	NA	NA	H	5	0	0%
Extreme zoom	1	1	0.33	HRD	181	41	23%
Grand Canyon 1	1	1	0.88	HRD	52	45	86%
Grand Canyon 2	1	1	1.20	HRD	26	12	46%
Melanoma	7	1	1.88	Q	12	3	25%
Retina	5	0 ¹	0.58	Q	27	12	44%
Satellite	1	1	0.26	HRD	145	83	57%
White board	1	1	0.20	HRD	173	95	55%
White Tower	2	2	1.14	HRD	35	6	17%
Winter-Summer (Day)	2	1	0.60	HRD	95	21	22%
Winter Day-Summer Night	NA	NA	NA	HRD	NA	NA	NA

The second to fifth columns show the performance of GDB-ICP: The index (starting from 1) of the first successful keypoint match in the rank-ordered list, the index of the same successful keypoint match among only the **consistent** keypoint matches, the final alignment error, and the choice of final transformation model—Homography (H), Homography with Radial lens Distortion (HRD), and Quadratic (Q). The last three columns show keypoint statistics which help to explain why RANSAC-based algorithms are less successful: the total number of keypoint matches (Lowe’s similarity ratio < 0.8), the number whose positions are within 6.0 pixels following application of the verified transformation, and the ratio between the two. On the “Retina” pair, GDB-ICP succeeded on an “inconsistent” keypoint match—one with 8.0 pixels of position error.

Just as in the previous test, we evaluate all 50 keypoint matches from the 19 pairs GDB-ICP succeeds upon. We study the change in the aforementioned 418 successes with changes in the feature extraction and matching. We also determine whether one of these changes causes the entire GDB-ICP to fail on a pair on which it originally succeeded.

The tests are summarized in Table 3, which shows several important results. First, using corner points alone without face points results in a 39 percent drop in the number of successful initial keypoint matches and a loss of eight successful pairs. Apparently, corners are not widely enough and densely enough distributed for GDB-ICP to succeed consistently on our challenging data set. Using faces alone, the drop is only 4 percent and no pairs are lost.

Interestingly, using forward matching alone instead of bidirectional matching causes the loss of 15 percent of the successful keypoint initializations and two of the most difficult pairs—EO-IR 1 and Melanoma. In a related result, not shown in Table 3, the percentage of driving features that are mapped to within two standard deviations of their corresponding matchable features, thereby creating “inlier”

correspondences, ranges from 58 percent to 83 percent. This indirectly justifies 1) the ability of GDB-ICP to adapt to substantial differences between images and 2) the decision to

TABLE 3
GDB-ICP Success Numbers Based on Varying the Feature Set and the Matching

variations	succ. init.	succ. pairs	variations	succ. init.	succ. pairs
corners and faces	418	19	forward matching	356	17
corners only	256	11	scale 1.0	397	16
faces only	401	19	scale 1.4	404	15
driving only	415	17	scale 2.0	362	14
matchable only	402	19	combined across scales	390	19

The second column is the number of initializations (keypoint matches) that led to successful alignments, while the third column is the pairs for which at least one initialization succeeded.

push feature extraction toward covering as much of an image as possible, trusting the rest of the algorithm to automatically determine which features are consistent between images.

The final test, shown on the right in Table 3, explores multiscale feature extraction. Using scale 1.0 (standard deviation of Gaussian smoothing) results in a loss of 5 percent initializations and three image-pairs—EO-IR 1, Melanoma, and Grand Canyon 2. When the single scale at which the features are extracted is increased, the success rate drops slowly. Finally, when using features combined across scales, similarly to the scale-space detection technique of many keypoint matching algorithms [30], [36], there is a 7 percent drop in the number of successful initializations, but no loss of any pairs.

6.6 Decision Criteria

To analyze the strength of the three-component decision criteria, we compared them with simplified versions. The results show that all three components of the decision criteria are necessary, that bi-directional decision increases robustness, and that the complete decision criteria are effective in distinguishing correct alignments from incorrect ones, even in the presence of low overlap, scale differences and physical changes.

One of the simplified versions is the use of alignment error alone, a natural measure for registration based on geometric constraints, and the one used in [11], [48] for retinal image registration. We then considered the importance of the three criteria by leaving each out in turn. Finally, we considered the effect of several other aspects of the decision criteria. In these experiments, all $42 \cdot 41 = 1,722$ possible ordered pairs of images are used, with both orderings used for each pair because each can produce different initial keypoint matches (see Section 2 for details) and therefore different initializations. The final model is always the homography (for speed considerations). Alignments passed by the modified decision criteria were examined by a graduate student (not one of the authors) to determine correctness. This turned out to be crucial because this test discovered some small overlaps in our image set that we did not realize existed. Based on this human judgment and based on our verified results, the decisions made by GDB-ICP under various decision criteria could be classified as True Positives (TP), False Positives (FP), True Negatives (TN), and False Negatives (FN). No changes were made in the parameter values of the decision criteria throughout the experiment.

The results are summarized in Table 4. The first observation concerns the effectiveness of the full decision criteria. The six false negatives come from the three pairs which GDB-ICP is known to fail on. The three false positives are image pairs that appeared locally consistent, with one of the regions having very low contrast. On the other hand, most such pairs are rejected. In fact, 99.8 percent of the incorrect pairs are rejected. Stated even more strongly, among the 1,671 rejected pairs, there are $1,671 \times 50 = 83,550$ incorrect initializations, including many with low overlap, all of which are rejected.

As seen in the next four rows of the table, when using only part of the criteria, the number of false positives increases significantly—jumping to 158 for use of the accuracy measure only, but with fewer when two of the three measures are used. These false positives are due to locally consistent structures, especially near image boundaries (producing apparent low-overlap between images). These experiments show convincingly that all three decision criteria are important. Finally, when the decision criteria are applied

TABLE 4
Effects of Varying the Decision Criteria
When Applying GDB-ICP to All Possible Pairs

	TP	TN	FP	FN
Full	42	1671	3	6
Accuracy alone	41	1515	158	8
No Stability	42	1614	60	6
No Orientation	42	1646	28	6
No Accuracy	42	1638	36	6
Only forward	42	1641	33	6

TP stands for true positives, TN for true negatives, FP for false positives, and FN for false negatives. See text in Section 6.6 for details.

only in the forward direction—from the moving image to the fixed image—the false positive rate increases substantially.

7 DISCUSSION

The experiments on our challenging data set have demonstrated that GDB-ICP is a powerful registration algorithm, capable of aligning a wide variety of image pairs. Overall, our experience with this data set and with other pairs shows that GDB-ICP succeeds when keypoint matching produces a small number of consistent matches, when there is sufficient common structure between the images to drive the dual-bootstrap process and the decision criteria, and when the similarity transformation is a reasonable initial model. In this section, we examine this success, using the experiments to show how the design of the algorithm allows it to handle the image registration challenges outlined in Section 1. We also discuss some limitations of the algorithm. We conclude this section by reexamining several individual components of the algorithm.

Our experiments have shown that GDB-ICP can succeed with as few as one consistent keypoint match and with as little as 58 percent correct matches between the driving features of one image and the matchable features of the other. Remembering that matchable features must pass less stringent tests than driving features and recalling that contrast reversals are ignored by the matching process, this result explains why the algorithm does so well with substantial changes in image illumination and structure and even changes in modality. The tolerance for differences in feature extraction and matching is allowed in the algorithm because the decision criteria can be trusted to reject nearly all incorrect alignments. The effectiveness of the decision criteria is also crucial to the algorithm's success in handling low-overlap pairs. Using the criteria, GDB-ICP accepts a small number of correct, low-overlap alignments while rejecting the extremely large number of low-overlap alignments generated by incorrect keypoint matches. Finally, the ability to generate matches across scales is crucial to handling substantial differences in scale.² While another recent algorithm has shown the ability to handle large scale variations [15], it has not been demonstrated in as challenging a context as our data set.

2. After these tests were completed, we tested a pair with a zoom factor of 9.5 and GDB-ICP succeeded.

Despite the demonstrated success, GDB-ICP does have limitations:

- It cannot handle extreme appearance differences between image pairs. In our data set, this is mostly due to initial keypoint matching, but we anticipate that the algorithm will fail on other multimodal pairs such as PET-CT pairs, where it is unlikely that the features will capture enough structural similarity. Still, GDB-ICP did succeed on all but one of the multimodal pairs in our data set because multimodal images often do have sufficient common structure. Intuitively, the structural and textural differences between the color image and the fluorescein angiogram in the “Melanoma” pair put it near the limit of what GDB-ICP can handle.
- Currently, the decision criteria do not eliminate incorrect alignments of an image-wide repetitive structure such as a checkerboard. On the other hand, if only a moderate fraction of a scene involves repetitive structure, the decision criteria make the right decision: Although incorrect alignments in the repetitive region will appear accurate, these produce inconsistent matches image-wide. An example of this occurs in the “Brussels” pair of our data set.
- As a 2D registration algorithm, GDB-ICP currently only tolerates a small amount of parallax. In a similar vein, it is currently limited to global transformation models.
- There is no convergence proof in the overall algorithm, just as there is no convergence proof in the ICP using anything but Euclidean match distances. In practice, however, in thousands of tests, GDB-ICP has always converged. One reason for this is that both region growth and model selection are monotonic.
- Finally, GDB-ICP, while consistently running in less than a minute for two mega-pixel image pairs, is still somewhat slow.

All of these issues are topics of our ongoing work.

Finally, we make a few observations about the individual components of GDB-ICP:

- The experiments show the importance of using multiscale face point (edge-like) features during the alignment process instead of more sparse features such as corners, even when corners are detected at multiple scales.
- While it is straightforward to replace Lowe’s LoG keypoint detector and SIFT descriptor [31] with other current keypoint techniques [37], [38], it would be surprising if this would improve keypoint matching substantially on our data set. Still, a thorough exploration of this question is worthy of future study.
- As shown in our earlier work [48] and reinforced by our results here, the combination of re-estimation and model selection in the bootstrap region keeps the estimate close to the local minimum while gradually increasing the problem complexity through region growth. This allows a technique as simple and efficient as our robust version of ICP to succeed as the core refinement procedure on the vast majority of the initializations. More sophisticated procedures, such as EM-ICP and RPM [14], [21], which simultaneously consider multiple per feature matches during

registration, might be considered in place of ICP. In the context of the Dual-Bootstrap approach, however, our informal experiments with the central idea of these methods—multiple matches per feature—showed that, because of the large number of outliers, even in the bootstrap regions, refinement must heavily rely on the distance to the nearest matching feature. These observations pushed us back toward robust ICP instead of EM-ICP or RPM. Still, a definitive answer to the question of the optimal core refinement procedure working within the context of the overall Dual-Bootstrap algorithm is beyond the scope of this paper.

- The axis-aligned, rectangular shape of the region is a simple, efficient representation. The new use of two independent regions introduced here is important for handling large differences in scale. Region models more sophisticated than our rectangular model could be developed—e.g., ones encompassing only the points where the trace of the transfer error covariance matrix is below a threshold—but the rectangular shape has proven sufficient for all our experiments and has not been the cause of an algorithm failure. Note that other region growth techniques have recently been proposed in the literature [17], [47]. Ours differs in that its growth is monotonic and is controlled by the uncertainty in the transformation estimate.
- The three-part decision criteria have proven to be essential for handling the challenging image pairs studied here. Other techniques include Brown and Lowe’s combinatorial analysis of keypoint matches [7] and Belongie et al.’s use of distance, brightness variation, and bending energy for recognition [3]. Clearly, keypoint-based measures alone are insufficient. Measures based on intensity variation or gradient magnitude differences [45] are not appropriate for the range of appearance variation seen here. Finally, although our three-part criteria have proven highly successful, they are not perfect and further improvements are possible.

8 SUMMARY AND CONCLUSION

This paper has presented the fully-automatic Generalized Dual-Bootstrap ICP (GDB-ICP) image registration algorithm designed to handle a wide variety of image pairs, including those showing scale changes, orientation differences, low overlap, illumination differences, physical changes, and different modalities. Building extensively on existing work, the algorithm is in fact a series of algorithms designed to work together to solve the problem. Extensive experiments on a 22 image-pair data set representative of these challenges have shown the effectiveness of the design and demonstrated that a broadly applicable, fully automatic image registration is possible. The experiments have also highlighted areas of potential improvement. The most important of these is initialization, especially when there are large appearance variations between images caused by physical or illumination changes or differences in image modalities. Despite this, the experiments reported here and experience by both our group and others who have tested the GDB-ICP executable demonstrate that it is effective enough for widespread use.

ACKNOWLEDGMENTS

This work was supported in part by the US National Science Foundation (NSF) Experimental Partnerships under Grant EIA-0000417, in part by the Center for Subsurface Sensing and Imaging Systems through the NSF Engineering Research Centers Program under Award EEC-9986821, in part by the US Army Intelligence and Security Command under Grant ABCDE, and in part by Rensselaer Polytechnic Institute.

REFERENCES

- [1] S. Baker and I. Matthews, "Lucas-Kanade 20 Years On: A Unifying Framework," *Int'l J. Computer Vision*, vol. 56, no. 3, pp. 221-255, 2004.
- [2] S. Belongie, C. Carson, H. Greenspan, and J. Malik, "Color- and Texture-Based Image Segmentation Using EM and Its Application to Content-Based Image Retrieval," *Proc. IEEE Int'l Conf. Computer Vision*, pp. 675-682, 1998.
- [3] S. Belongie, J. Malik, and J. Puzicha, "Shape Matching and Object Recognition Using Shape Contexts," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 24, no. 4, pp. 509-522, Apr. 2002.
- [4] J.R. Bergen, P. Anandan, K.J. Hanna, and R. Hingorani, "Hierarchical Model-Based Motion Estimation," *Proc. Second European Conf. Computer Vision*, pp. 237-252, 1992.
- [5] P. Besl and N. McKay, "A Method for Registration of 3D Shapes," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239-256, Feb. 1992.
- [6] L.G. Brown, "A Survey of Image Registration Techniques," *ACM Computing Surveys*, vol. 24, no. 4, pp. 325-376, Dec. 1992.
- [7] M. Brown and D. Lowe, "Recognising Panoramas," *Proc. IEEE Int'l Conf. Computer Vision*, 2003.
- [8] M. Brown, R. Szeliski, and S. Winder, "Multi-Image Matching Using Multi-Scale Oriented Patches," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, vol. 1, pp. 510-517, 2005.
- [9] K. Bubna and C.V. Stewart, "Model Selection Techniques and Merging Rules for Range Data Segmentation Algorithms," *Computer Vision and Image Understanding*, vol. 80, pp. 215-245, 2000.
- [10] K.P. Burnham and D.R. Anderson, *Model Selection and Inference: A Practical Information-Theoretic Approach*, first ed. Springer, 1998.
- [11] A. Can, C. Stewart, B. Roysam, and H. Tanenbaum, "A Feature-Based, Robust, Hierarchical Algorithm for Registering Pairs of Images of the Curved Human Retina," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 24, no. 3, pp. 347-364, Mar. 2002.
- [12] G. Champelboux, S. Lavallee, R. Szeliski, and L. Brunie, "From Accurate Range Imaging Sensor Calibration to Accurate Model-Based 3-D Object Localization," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 83-89, 1992.
- [13] Y. Chen and G. Medioni, "Object Modeling by Registration of Multiple Range Images," *Image and Vision Computing*, vol. 10, no. 3, pp. 145-155, 1992.
- [14] H. Chui, A. Rangarajan, J. Zhang, and C.M. Leonard, "Unsupervised Learning of an Atlas from Unlabeled Point-Sets," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 26, no. 2, pp. 160-172, Feb. 2004.
- [15] Y. Dufournaud, C. Schmid, and R. Horaud, "Image Matching with Scale Adjustment," *Computer Vision and Image Understanding*, vol. 93, pp. 175-194, 2004.
- [16] R. Fergus, P. Perona, and A. Zisserman, "Object Class Recognition by Unsupervised Scale-Invariant Learning," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2003.
- [17] V. Ferrari, T. Tuytelaars, and L.V. Gool, "Simultaneous Object Recognition and Segmentation by Image Exploration," *Proc. Eighth European Conf. Computer Vision*, 2004.
- [18] R. Fransens, C. Strecha, and L.V. Gool, "Multimodal and Multi-band Image Registration Using Mutual Information," *Proc. Theory and Applications of Knowledge-Driven Image Information Mining with Focus on Earth Observation (ESA-EUSC)*, 2004.
- [19] W.T. Freeman and E.H. Adelson, "The Design and Use of Steerable Filters," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 9, pp. 891-906, Sept. 1991.
- [20] L.V. Gool, T. Moons, and D. Ungureanu, "Affine/Photometric Invariants for Planar Intensity Patterns," *Proc. Fourth European Conf. Computer Vision*, 1996.
- [21] S. Granger and X. Pennec, "Multi-Scale EM-ICP: A Fast and Robust Approach for Surface Registration," *Proc. Seventh European Conf. Computer Vision*, pp. 418-432, 2002.
- [22] W. Grimson, T. Lozano-Perez, W. Wells, G. Ettinger, and S. White, "An Automatic Registration Method for Frameless Stereotaxy, Image, Guided Surgery and Enhanced Reality Visualization," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 430-436, 1994.
- [23] C. Harris and M. Stephens, "A Combined Corner and Edge Detector," *Proc. Fourth Alvey Vision Conf.*, pp. 147-151, 1988.
- [24] R. Hartley and A. Zisserman, *Multiple View Geometry*. Cambridge Univ. Press, 2000.
- [25] D.L.G. Hill, P.G. Batchelor, M. Holden, and D.J. Hawkes, "Medical Image Registration," *Physics in Medicine and Biology*, vol. 46, no. 3, 2001.
- [26] P.W. Holland and R.E. Welsch, "Robust Regression Using Iteratively Reweighted Least-Squares," *Comm. Statistics: Theory and Methods*, vol. A6, pp. 813-827, 1977.
- [27] M. Irani and P. Anandan, "Robust Multisensor Image Alignment," *Proc. IEEE Int'l Conf. Computer Vision*, pp. 959-966, 1998.
- [28] T. Kadir, A. Zisserman, and M. Brady, "An Affine Invariant Salient Region Detector," *Proc. Eighth European Conf. Computer Vision*, 2004.
- [29] K. Kanatani, *Statistical Optimization for Geometric Computation: Theory and Practice*. Elsevier, 1996.
- [30] T. Lindeberg, *Scale-Space Theory in Computer Vision*. Kluwer Academic, 1994.
- [31] D.G. Lowe, "Distinctive Image Features from Scale-Invariant Keypoints," *Int'l J. Computer Vision*, vol. 60, no. 2, pp. 91-110, Nov. 2004.
- [32] B. Luo and E.R. Hancock, "Iterative Procrustes Alignment with the EM Algorithm," *Image and Vision Computing*, vol. 20, nos. 5-6, pp. 377-396, Apr. 2002.
- [33] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality Image Registration by Maximization of Mutual Information," *IEEE Trans. Medical Imaging*, vol. 16, no. 2, pp. 87-198, 1997.
- [34] J. Matas, O. Chum, M. Urban, and T. Pajdla, "Robust Wide-Baseline Stereo from Maximally Stable Extremal Regions," *Image and Vision Computing*, vol. 22, no. 10, pp. 761-767, Sept. 2004.
- [35] P. Meer, "Robust Techniques for Computer Vision," *Emerging Topics in Computer Vision*, G. Medioni and S.B. Kang, eds., Prentice Hall, 2004.
- [36] K. Mikolajczyk and C. Schmid, "Scale and Affine Invariant Interest Point Detectors," *Int'l J. Computer Vision*, vol. 60, no. 1, pp. 63-86, 2004.
- [37] K. Mikolajczyk and C. Schmid, "A Performance Evaluation of Local Descriptors," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 27, no. 10, pp. 1615-1630, Oct. 2005.
- [38] K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L.V. Gool, "A Comparison of Affine Region Detectors," *Int'l J. Computer Vision*, vol. 65, nos. 1-2, pp. 43-72, 2005.
- [39] J.V. Miller, "Regression-Base Surface Reconstruction: Coping with Noise, Outliers, and Discontinuities," PhD thesis, Rensselaer Polytechnic Inst., Aug. 1997.
- [40] J.P.W. Pluim, J.B.A. Maintz, and M.A. Viergever, "Image Registration by Maximization of Combined Mutual Information and Gradient Information," *IEEE Trans. Medical Imaging*, vol. 19, no. 8, pp. 809-814, Aug. 2000.
- [41] J.P.W. Pluim, J.B.A. Maintz, and M.A. Viergever, "Mutual-Information-Based Registration of Medical Images: A Survey," *IEEE Trans. Medical Imaging*, vol. 22, no. 8, pp. 986-1004, 2003.
- [42] H. Sawhney, S. Hsu, and R. Kumar, "Robust Video Mosaicing through Topology Inference and Local to Global Alignment," *Proc. Fifth European Conf. Computer Vision*, vol. II, pp. 103-119, 1998.
- [43] F. Schaffalitzky and A. Zisserman, "Multi-View Matching for Unordered Image Sets, or How Do I Organize My Holiday Snaps," *Proc. Seventh European Conf. Computer Vision*, vol. 1, pp. 414-431, 2002.
- [44] C. Schmid, R. Mohr, and C. Bauckhage, "Evaluation of Interest Point Detectors," *Int'l J. Computer Vision*, vol. 37, no. 2, pp. 151-172, 2000.
- [45] Y. Shan, H.S. Sawhney, and R. Kumar, "Unsupervised Learning of Discriminative Edge Measures for Vehicle Matching between Non-Overlapping Cameras," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2005.

- [46] D. Shen and C. Davatzikos, "Hammer: Hierarchical Attribute Matching Mechanism for Elastic Registration," *IEEE Trans. Medical Imaging*, vol. 21, no. 11, pp. 1421-1439, 2002.
- [47] K.L. Steele and P.K. Egbert, "Correspondence Expansion for Wide Baseline Stereo," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2005.
- [48] C. Stewart, C.-L. Tsai, and B. Roysam, "The Dual-Bootstrap Iterative Closest Point Algorithm with Application to Retinal Image Registration," *IEEE Trans. Medical Imaging*, vol. 22, no. 11, pp. 1379-1394, 2003.
- [49] C.V. Stewart, "Robust Parameter Estimation in Computer Vision," *SIAM Rev.*, vol. 41, no. 3, pp. 513-537, 1999.
- [50] P. Torr, "An Assessment of Information Criteria for Motion Model Selection," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 47-52, 1997.
- [51] P. Torr and A. Zisserman, "MLESAC: A New Robust Estimator with Application to Estimating Image Geometry," *Computer Vision and Image Understanding*, vol. 78, no. 1, pp. 138-156, Apr. 2000.
- [52] T. Tuytelaars and L.V. Gool, "Matching Widely Separated Views Based on Affine Invariant Regions," *Int'l J. Computer Vision*, vol. 1, no. 59, pp. 61-85, 2004.
- [53] P. Viola and W.M. Wells III, "Alignment by Maximization of Mutual Information," *Int'l J. Computer Vision*, vol. 24, no. 2, pp. 137-154, 1997.
- [54] B. Zitova and J. Flusser, "Image Registration Methods: A Survey," *Image and Vision Computing*, vol. 21, pp. 977-1000, 2003.



Gehua Yang received the BS degree in applied mathematics from South China University of Technology in 2000 and the MS degree in applied mathematics from Rensselaer Polytechnic Institute in 2002. Currently, he is pursuing the PhD degree in computer science at Rensselaer Polytechnic Institute and expects to graduate in summer 2007. His research interests include computer vision and medical image analysis. He is a student member of the IEEE.



Charles V. Stewart received the BA degree in mathematical sciences from Williams College in 1982 and the MS and PhD degrees in computer science from the University of Wisconsin in 1985 and 1988, respectively. Currently, he is a professor in the Department of Computer Science, Rensselaer Polytechnic Institute, Troy, New York. He has done sabbaticals at the GE Center for Research and Development in Niskayuna, New York, and at the Johns Hopkins University. In 1999, together with Ali Can and Badrinath Roysam, he received the Best Paper Award at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). He was local arrangements chair in 2003 and workshops chair in 2006 for CVPR. He is on the editorial board of *Image and Vision Computing*. His research interests include computer vision, medical image analysis, and graphics, emphasizing image registration and three-dimensional modeling. He is a member of the IEEE.



Michal Sofka did his undergraduate work at the Czech Technical University. He received the MS degree in electrical engineering from Union College in 2001 and the MS degree in computer science from Rensselaer Polytechnic Institute in 2006, where he is currently pursuing the PhD degree. In 2004, he was a technical employee in the Department of Intelligent Vision & Knowledge Based Reasoning at Siemens Corporate Research. His interests include feature extraction in medical imaging, initialization for image registration, uncertainty and constraints in image registration, range data registration, and deformable registration. He is a student member of the IEEE.



Chia-Ling Tsai received the BSc and BSc Hons degrees from the University of the Witwatersrand, South Africa, in 1994 and 1995, respectively, and the PhD degree in computer science from Rensselaer Polytechnic Institute, Troy, New York, in 2003. She has been at National Chung Cheng University, Taiwan, since 2005, where she is an assistant professor in the Department of Computer Science and Information Engineering. Her research interests include biomedical image analysis, computer vision, and image processing. Dr. Tsai is a member of the IEEE.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.