Improving Inter-fragmentary Alignment for Virtual 3D Reconstruction of Highly Fragmented Bone Fractures

Beibei Zhou[†], Andrew Willis[†], Yunfeng Sui[†],
Donald Anderson[‡], Thaddeus Thomas[‡], Thomas Brown[‡]
[†]UNC - Charlotte, 9201 University City Blvd., Charlotte, NC 28223
[‡]The University of Iowa, 2181 Westlawn Building, Iowa City, IA 52242

ABSTRACT

This article describes two new algorithms that, when integrated into an existing semi-automatic virtual bone fragment reconstruction system, allow for more accurate anatomic restoration. Furthermore, they spare the user the painstaking task of positioning each fragment in 3D, which can be extremely time consuming and difficult. The virtual interactive environment gives the user capabilities to influence the reconstruction process and to examine multiple potential reconstruction scenarios. Coarse fracture surface correspondences specified by the user are refined by a new alignment functional that allows idiosyncratic geometric surface variations such as ridges and valleys to more heavily influence the final alignment solution. Integration of these algorithms into the system provides improved performance, stability, and, most importantly, reconstruction accuracy, which is critical for increasing the likelihood of satisfactory clinical outcome after the injury.

Keywords: shape registration, statistical methods, 3d fragment reconstruction

1. INTRODUCTION

Extremity injuries that involve highly comminuted bone fractures almost always occur as a result of high-energy trauma such as vehicular accidents or falls from a height. They also are a major concern in military conflicts. Treatment goals include achieving expeditious bony union in a position of acceptable limb alignment, and avoiding **P**ost-**T**raumatic **O**steo**A**rthritis (PTOA) when there is involvement of an articular joint such as the hip, knee, or ankle. As a point of reference, for axial "pilon" fractures of the distal tibial articular surface (Figure 1), the incidence of PTOA of the ankle is in the range of 60% to 80% [1, 2]. Accurate restoration of the articular surface is believed to be critical in avoiding PTOA, but in many comminuted articular fractures, this task can be quite challenging. Often, dozens of individual fragments are involved, displaced appreciably from their site of anatomic origin, and interspersed in a complex geometric pattern.

The trauma surgeon reconstructing a comminuted fracture faces a problem very much akin to puzzle solving, albeit with obvious additional complexities. Muscle forces, combined with complex displacements, intervening soft tissues, and interaction between fragments and fracture surfaces all serve to make it difficult to reposition the fracture fragments. To partially or completely restore the osseous anatomy, traditional surgical treatment (open fracture reduction) exposes the fragments by a surgical approach through the damaged soft tissue envelope, so that the surgeon can directly access and reposition the fragments. Soft tissues are dissected off the fragments, further contributing to the injury that the body must heal. The surgeon assesses fragment inter-digitations visually and by touch, effecting a sequential fragment reduction. This often requires considerable force, and involves trial and error. Unfortunately each "error" prolongs the procedure and adds yet more trauma to the fragments and the surrounding soft tissues. The fracture is considered to be "reduced" when the surgeon judges that an optimal fit has been obtained between all relevant fragments, completing the fracture puzzle solution.

The more extensive the surgical dissection, the better the fracture is visualized and the easier it is to execute an open fracture reduction. However, wide surgical exposure comes at a significant price, including increased risk for wound healing failure, infection, joint stiffness, delayed fracture healing, and damage to articular surfaces. These risks provide tremendous impetus for developing less invasive techniques to surgically restore displaced

Further author information: (Send correspondence to A. Willis)
A. Willis: E-mail: arwillis@uncc.edu, Telephone: 1 704 687 8420
D. Anderson: E-mail: don-anderson@uiowa.edu, Telephone: 1 319 335 8135

Figure 1: Radiographs of a range of comminuted tibial pilon fractures, varying in the number of bone fragments, the amount of fragment dispersal and the degree of intra-articular extension into the ankle.

comminuted fractures. To reduce fractures through less invasive approaches, however, the surgeon must utilize special techniques, and to substantially replace wide operative approaches new techniques must be developed. Fragment reduction accuracy is a particularly important consideration for peri-and intra-articular comminuted fractures. Residual geometric incongruity of an articular bearing surface strongly predisposes to painful and debilitating PTOA, a poor clinical outcome regardless of osseous union.

2. CONTRIBUTION OF THIS WORK

The literature regarding semi-automatic or automatic reconstruction of bone fragments is surprisingly sparse given the very large body of research on the topic of bone segmentation, medical image registration, and surface registration. Early efforts on this topic included a system for reconstructing a simple two-fragment bone fracture [3]. There has been increased research activity on this topic recently in terms of algorithm development for reconstruction systems [4, 5] and complete interactive reconstruction systems such as that by Harders *et. al.* [6]. The present work extends a previously proposed system [7] by introducing significant improvements to both the interactive aspects and alignment algorithms of the system.

While intra-operative navigation tools exist to aid in reconstruction of long-bone fractures [3, 6], these systems require a large degree of user interaction. This can be cumbersome, especially for the complex geometries which arise in comminuted bone fractures. In previous work [7], we proposed a largely automatic system that uses robust computational algorithms to solve important aspects of the reconstruction problem. This work develops new solutions for two critical aspects of this system: (1) a new interface that allows users to direct the search for fragment matches and (2) a new matching functional that boosts alignment accuracy for matched surfaces, especially for nearly flat surfaces and surfaces with subtle variations. It also introduces an experimental methodology that provides a platform for continuing algorithmic developments.

3. EXPERIMENTAL METHODOLOGY

This work was facilitated by fracture experiments performed utilizing test specimens of controlled geometry, machined from bone-surrogate material [8], which provide fracture surface data that enables fine tuning of alignment algorithms and optimization of system performance (see Figure 2 for details). This provides a well-ordered developmental platform that bypasses unnecessarily confounding difficulties inherent in biologic variability. An important goal of the work is to build evidence that the enabling technology of 3-D puzzle solving is applicable to fragmentation / dispersal patterns typical of comminuted orthopaedic fractures. Puzzle solution computations are performed for drop-tower-created fractures of anatomically realistic bony surrogates fabricated from a specialty high-density polymeric foam material that has two important attributes. First, it exhibits mechanical behavior nominally comparable to that of human cortical bone in terms of mechanical strength, fracture toughness (i.e., energy absorption capacity during impact), and propensity to shatter into bone-similar numbers and morphologies of discrete fragments. Second, when appropriately doped with a radio-opacification agent (20% anhydrous $BaSO_4$, by weight), it has bone-similar radiographic appearance on x-ray and CT. Anatomically accurate test specimens (distal tibia replicas) are custom-machined, using source data provided from clinical CT scans. Following fracturing of the specimens, CT scan data are obtained for subsequent fracture reconstruction work.

(a) (b) (c)

Figure 2: (a) Distal tibia replicas (3D laser scan shown, with photographic texture overlaid), machined from bone surrogate material, are marked to allow definitive post-fracture identification of the site of origins for each of the bone fragments. (b) A frame from a high-speed video of an impact event is shown here. A custom drop-tower is used to deliver a fracturing impact to the bone, embedded in a soft tissue simulant gelatin. (c) Bone fragments are laser scanned for virtual reconstruction.

4. SYSTEM OVERVIEW

Fracture reconstruction using the system from [7] proceeds via a series of five steps:

1. Segment the volumetric CT data via [9], with an emphasis on generating discrete closed surfaces, i.e., fragment surfaces that may later be aligned .
2. Generate a surface mesh for each segmented bone fragment, and map each point on the surface mesh to a Hounsfield intensity (for the 3D surface point $\mathbf{x}_i$ this map is referred to as $f(\mathbf{x}_i)$) from the 3D CT scan volume $\mathbf{I}(x, y, z)$ using a 3D $max()$ filter centered about the voxel containing back-projection of each 3D fragment surface point into the CT scan volume.
3. Segment fragment surfaces into intact- and fracture-surfaces using a two-class Bayesian classifier based on the Hounsfield intensities mapped onto the surface vertices in step 2.
4. Coarsely identify matching patches over contiguous fracture-surface regions, utilizing an interactive user interface.
5. Optimize the specified fragment alignments to restore the original bone anatomy from the bone fragments, using the interactively specified regions.

Figure 3 shows graphically an example of this process. Our work in this paper focuses on steps 4 and 5. In step 4, we provide new interactions for surface matching (see §5). For step 5, we use a new surface alignment criterion that constrains the alignment to match well in regions that contain significant geometric variation such as valleys and ridges (see §6).

5. USER-DIRECTED SEARCH FOR FRAGMENT MATCHES

The most time-consuming aspect of the reconstruction system is the user interaction needed in step 4. Here the user manually subdivides automatically identified fracture-surface regions into pairs of surfaces that are likely to, but need not exactly correspond. This interaction avoids a complex (and often unreliable) computational search necessary to identify fracture surface sub-regions on each fragment that match with *only one other fragment*. Such surface patches are required to perform a bottom-up reconstruction of the fracture fragments via a sequence of pairwise fragment alignments detailed in §6. Custom-built 3D interactive software allows the user to easily specify correspondences between fragment pairs by selecting regions from the apparent image of the surface as rendered to the viewing screen. Specified matches are appended to a running match-list that tracks all of the fragment

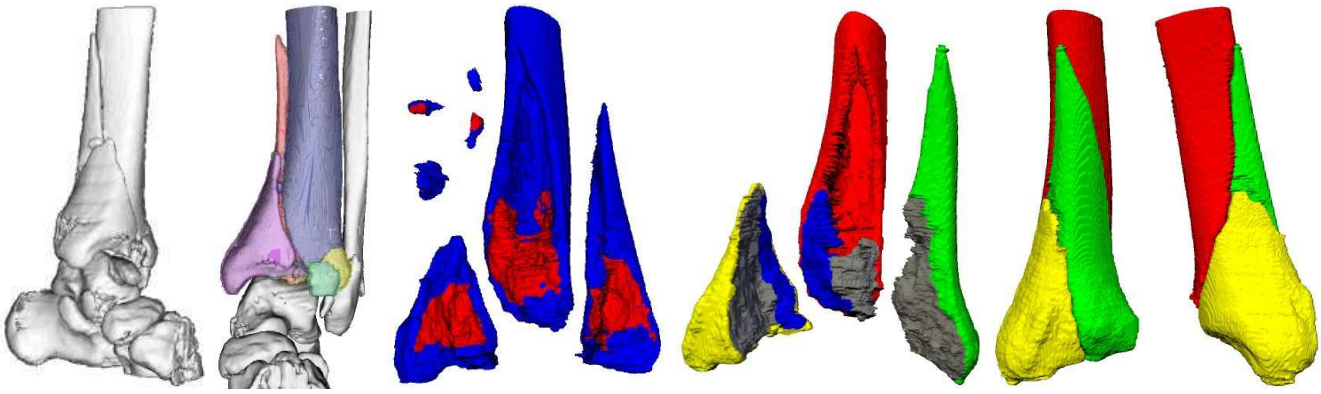

(a) (b) (c) (d) (e) (f)

Figure 3: An overview of the fracture fragment identification and reconstruction system: (a) an iso-surface of a fractured tibia, (b) segmentation of the CT scan data generates discrete closed fragment surfaces shown in different colors (c) fragment surface points are then divided into intact-surface points (blue) and fracture-surface points (red) (d) fracture surfaces are subdivided interactively into gray and blue regions denoting sub-regions of various fragments' fractures surfaces (e-f) two views of an aligned fracture case using the multi-fragment surface alignment algorithm described in §6.

matches for later reconstruction. At any point in the assembly process, the user can trigger an automatic bone reconstruction which is animated for the user, based upon the system's instantaneous computation of the most probable solution. These interactive tools also provide insights on the actual surgical reconstruction.

6. AUTOMATIC ALIGNMENT OF FRAGMENT SURFACES

The automatic solution to this alignment problem turns out to be non-trivial, since the user may specify several surface correspondences for a given fragment pair, and since a single fragment may include matches with many other fragments. The solution is obtained in two "planning" steps followed by a final step that aligns the surfaces in the match-list. The planning steps involve (i) dividing the match-list into groups such that each includes only matches for a single fragment pair and (ii) deciding on the sequence in which fracture fragments will be geometrically aligned to the global reconstruction solution. Fracture reconstruction is accomplished by iteratively merging individual fragments into a group of aligned fragments using the geometric data from the match-list and the alignment sequence computed in the planning steps. Following each fragment merge, the pose of each fragment in the reconstruction is adjusted to accommodate the newly added fragment data. The algorithmic steps for the reconstruction are listed below and closely follow those described in [10] (see Figure 4):

1. Compute a graph $G(M, F)$ where the i^{th} graph node, F_i, corresponds to fragment i in the fracture and each edge M_k represents a user-specified surface match (correspondence) between two fragments.
   Note that under this graph model two graph nodes may be connected by more than one and possibly many edges. These situations occur when the user has specified several different matches between the same two fragments.
2. Construct an active set and a dormant set of specified surface patch pairs. Initially, all scans are placed in the dormant set.
3. Find the fragment in the dormant set that has the highest number of edges to fragments from the active set. Initially, we simply choose the fragment with the highest number of edges.
4. For the selected fragment pair, we find all specified patches that include this fragment pair and merge the correspondences into a single correspondence set.
5. All of these patches are simultaneously aligned to find the best local alignment between the fragment pair for all specified correspondences.

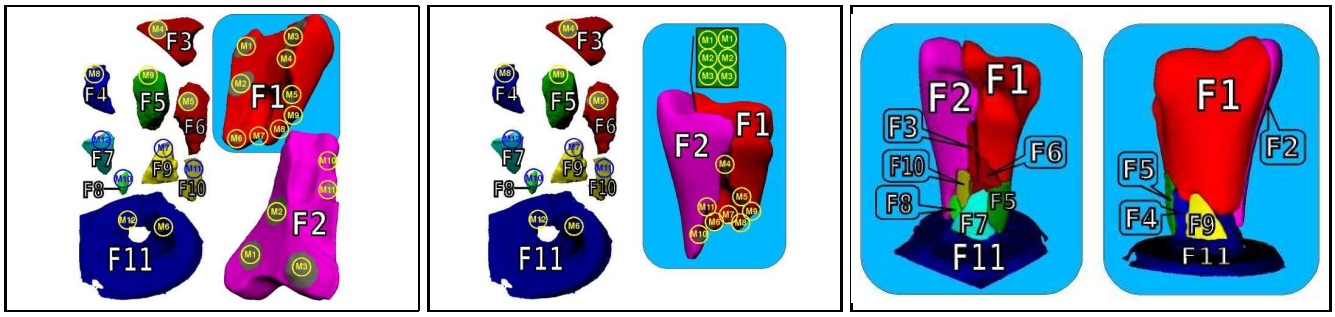

(a) (b) (c)

Figure 4: (a) shows eleven fragments, numbered $F_1 - F_{11}$, generated from droptower impacts of a distal tibia surrogate (Figure 2). A user interactively selects regions that coarsely correspond on fragment pairings. Each correspondence is referred to as a surface match and is denoted M_i for the i^{th} surface match. (a,b) show the first two steps of the fracture reconstruction algorithm. (a) The F_1 fragment is added to the active set, denoted as the region with blue background. (b) Fragment F_2 shares the most (3) matches with the fragment in the active set, hence, it is aligned with F_1 generating a fragment pair. (c) shows the result generated by subsequent alignment of fragments $F_3 - F_{11}$ to the pair in (b) resulting in a global reconstruction where all fragments are in the active set.

6. All aligned fragment matches within the active set are re-aligned to accommodate for the addition of the new fragment.
7. If the dormant set is not empty, return to step 3. Otherwise, we exit the alignment algorithm.

Steps 1-3 are initialization steps for the algorithm and the choice of the initial fragment in step 3 may be arbitrary. We find that the algorithm tends to converge more quickly by choosing the fragment with most matches first. However, there may be reason to select a different initial fragment. One such case exists when one of the measured "fragments" is still in its proper anatomic position, by virtue of belonging to that portion of the bone that was not fractured (e.g., the proximal portion of the tibia in pilon fractures). Here, one may choose this "fragment" as the initial fragment such that the reconstruction algorithm will align pieces to the intact portion of the bone.

Steps 5 and 6 from this algorithm require one or more corresponding fragment surfaces to be aligned. Alignments are computed by optimizing a new alignment functional which is a modified metric for the Iterative Closest Point (ICP) [11]. Our modifications incorporate several improvements to the original ICP algorithm: (1) a point-to-plane Euclidean distance metric [12], (2) an intensity matching metric [13, 14], and (3) a sub-sampling of the surface which selects points more densely in regions of significant geometric surface variation such as ridges and valleys [15] which are simultaneously enforced during surface alignment by solving for the rotation $\mathbf{R}$ and translation vector $\mathbf{t}$ that minimizes the functional (1).

$$e = \min_{\mathbf{R},\mathbf{t}} \sum_{(i,j)\, pairs} p(f(\mathbf{x}_i)|\omega_o) p(f(\mathbf{y}_j)|\omega_o) \left\| \mathbf{x}_i - proj_{\mathbf{n}_i}(\mathbf{R}\mathbf{y}_j + \mathbf{t}) \right\|^2 \quad (1)$$

The point-to-plane error metric is enforced by the term $proj_{\mathbf{n}_i}(\mathbf{R}\mathbf{y}_j + \mathbf{t})$ which denotes the point lying in the plane passing through $\mathbf{y}_j$ with normal indicated by the surface normal at $\mathbf{y}_j$ and also lying along the 3D line passing through $\mathbf{x}_i$ in the direction given by the surface normal at $\mathbf{x}_i$ denoted $\mathbf{n}_i$, i.e., the point that is the intersection of this 3D line and 3D plane. This metric has been found to improve both the accuracy and convergence behavior of 3D surface alignment algorithms [12].

The intensity matching metric is enforced by the term $p(f(\mathbf{x}_i)|\omega_o)p(f(\mathbf{y}_j)|\omega_o)$. This method for alignment was first described in [7] and is based on the likelihood that the surface point is from cortical bone. Cortical bone regions have high contrast with respect to the neighboring soft tissues in CT images, and cortical surfaces are therefore more reliably extracted by bone fragment segmentation algorithms. For our experimental fractures generated using the drop tower, the fragments are readily available. In such cases we user laser scanning equipment

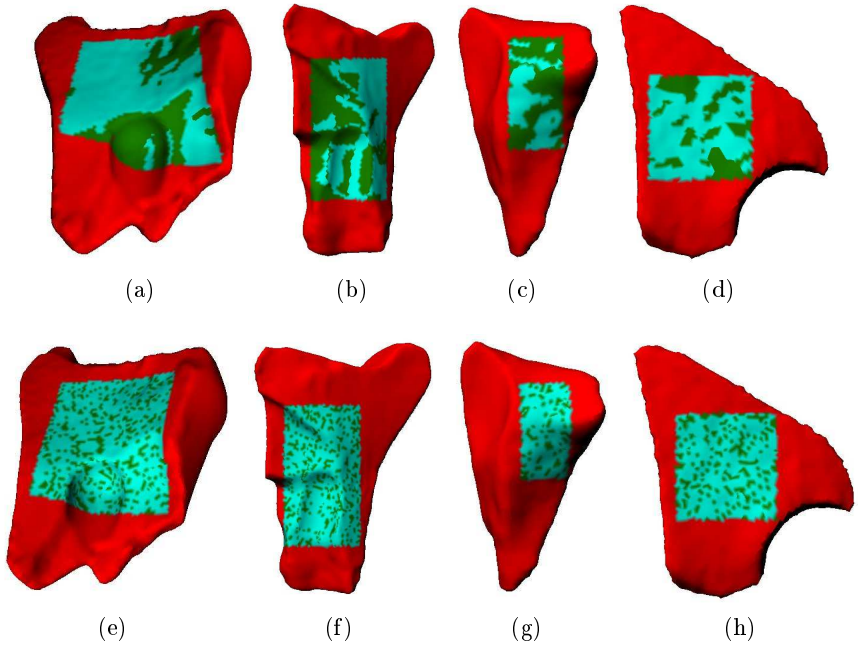

Figure 5: The top and bottom rows show fracture surfaces from four different fragments. A square surface region on each fracture surface (in cyan) has been interactively selected and different methods were used to select points for surface alignment. The alignment points are indicated by dark green locations within the each selected surface region. The top row shows the surface points automatically selected using geometrically stable sampling [15]. The bottom row shows surface points automatically selected using random sampling (see §6 for details). Note that geometrically-stable sampling tends to select points that lie on salient geometric structures on surfaces such as the ridges apparent in the surfaces shown in (a,b) and also more subtle surface variations such as the undulations on surfaces shown in (c,d).

to obtain surface measurements and the alignment term $p(f(\mathbf{x}_i)|\omega_o)p(f(\mathbf{y}_j)|\omega_o)$ is set to 1 since no CT intensities are available for each surface location.

For sparsely sampled surfaces, typically all points on each of the alignment surfaces are needed for a good alignment. However, more often the sampling of the alignment surface is dense and a subset of the complete set of points available for alignment are used. There are several potential approaches for selecting specific 3D points from each surface to align. Until recently, the typical approach was to select points uniformly distributed across the surface or by taking a random subset of the alignment points. Yet, work in [15, 16] propose new methods that identify specific points on the alignment surfaces that, when used for alignment, improve the accuracy and performance of geometric surface alignments. We use a modified version of the method proposed in [15], which computes a set of geometrically-stable points on the alignment surface. The geometrically-stable collection of points is determined by finding the collection of alignment points that cause the largest change in alignment error for am infinitessimal change in the fragment pose, i.e., it finds the subset of points that tend to dominate the alignment error when the fragment undergoes a change in it's 3D pose. Figure 5 shows the difference between geometrically-stable sampling and random sampling of an alignment surface for several objects.

For use in fracture reconstruction, we introduce a modification of the proposed alignment algorithm from [15]. Our modification uses geometrically-stable sampling for alignment surfaces that are being merged into the group of fragments that constitute the active set. In this way, we select the subset of points on the moving fragment that have the most influence on the final aligned fragment pose. These points are matched with all available points on the corresponding matching surface. This technique improves alignment results, especially for flat surfaces where the surface points selected using geometrically stable sampling may not closely correspond for a fragment pair. This geometrically-biased sampling of the alignment surfaces enhances the stability, convergence and accuracy of

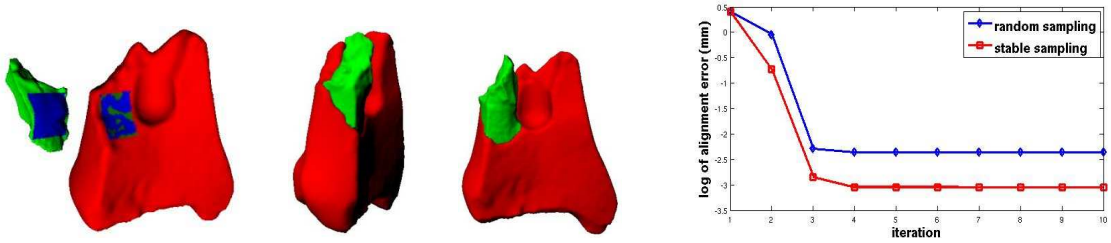

(a) (b) (c) (d)

Figure 6: (a) shows two fragments and their alignment surfaces (in blue). Geometrically stable sampling has been applied to the alignment surface of the red fragment, and the resulting subset of alignment points are shown in dark green. (b,c) show two views of the aligned fragment pair. (d) Shows the convergence behavior of the pairwise alignment using different sampling techniques to draw samples from the alignment surfaces. It compares the point-to-plane alignment error of 58 sampled points using the random sampling method (blue curve) and the geometrically-stable sampling method (red curve).

the alignment, particularly for surface regions whose curvature is zero in one or more principal directions, e.g., planes, cones, and cylinders, or surfaces having two equal principal curvatures, e.g., spheres. This is particularly important for bone fracture surfaces, which often include nominally cylindrical (endosteal cavity), planar (oblique fracture surfaces), conical (valleys / ridges with linear slope) and spherically-shaped surface patches.

Solving the minimization (1) provides the Euclidean transformation in terms of $\mathbf{R}$, an unknown 3x3 rotation matrix, and $\mathbf{t}$, a 3x1 translation vector which aligns two surfaces by minimizing the distance between the points on one surface to the plane defined by the nearest point and its associated normal on the other surface. For each matched point pair, the point-to-plane distance is weighted by the joint-likelihood that the surfaces at that location consist of cortical bone tissue. This weighting has two beneficial effects: (1) it encourages matched surfaces to join smoothly by emphasizing good matches along the surface perimeter and (2) it weights surface vertices derived from cortical tissue more than those of cancellous tissue, whose relatively low CT intensity contrast with respect to surrounding tissues yields more inherent segmentation noise (step 1).

7. RESULTS

Figure 3(a-e) provide a detailed overview of the reconstruction process for a clinical tibial pilon fracture case. Many fracture cases involve numerous (>15) fragments, e.g., the fracture in (a) has 18 fragments. Yet, many fragments are too small to be considered for reconstruction. For example, the fracture in Figure 3 contains 7 fragments of sufficient size to be clinically important for restoring alignment. In practice, smaller fragments are often left in place, as they aid in the bone healing process. However, their exact position and orientation are not considered to be clinically significant for successful fracture reduction.

Plots from Figure 6 show the convergence behavior of the alignment for a fragment pair using different sampling of the selected alignment surface. It shows that the alignment error using geometrically-stable sampling is lower than the alignment error when using random sampling.

Figure 7 shows reconstructions of fractures generated by the experimental drop tower. Images in the top row (a,b) show the reconstruction result for the surrogate distal tibia replicas in Figures (2 and 4). Images in the middle row (c,d,e) show an alignment result for fragments generated from a high-energy fracture of an artificial tibia thick-walled cylinder. Images in the bottom row (f,g,h) show an alignment result for a cadaveric bovine tibia fractured via the drop tower.

8. CONCLUSIONS

We have described algorithms that, when integrated with our semi-automatic fragment reconstruction system, allow for more accurate reconstructions without requiring painstaking positioning of each fragment in 3D, which can be extremely time consuming and difficult. The interactive environment gives the user abilities to influence the reconstruction process and to examine multiple potential reconstruction alternatives. Coarse fracture surface

correspondences specified by the user are refined by an augmented version of ICP which allows surface variations such as ridges and valleys to more heavily influence the final alignment solution. Integration of these algorithms into the system provides improved reconstruction accuracy, toward the goal of providing improved prognosis for these cases that are particularly difficult to surgically reconstruct.

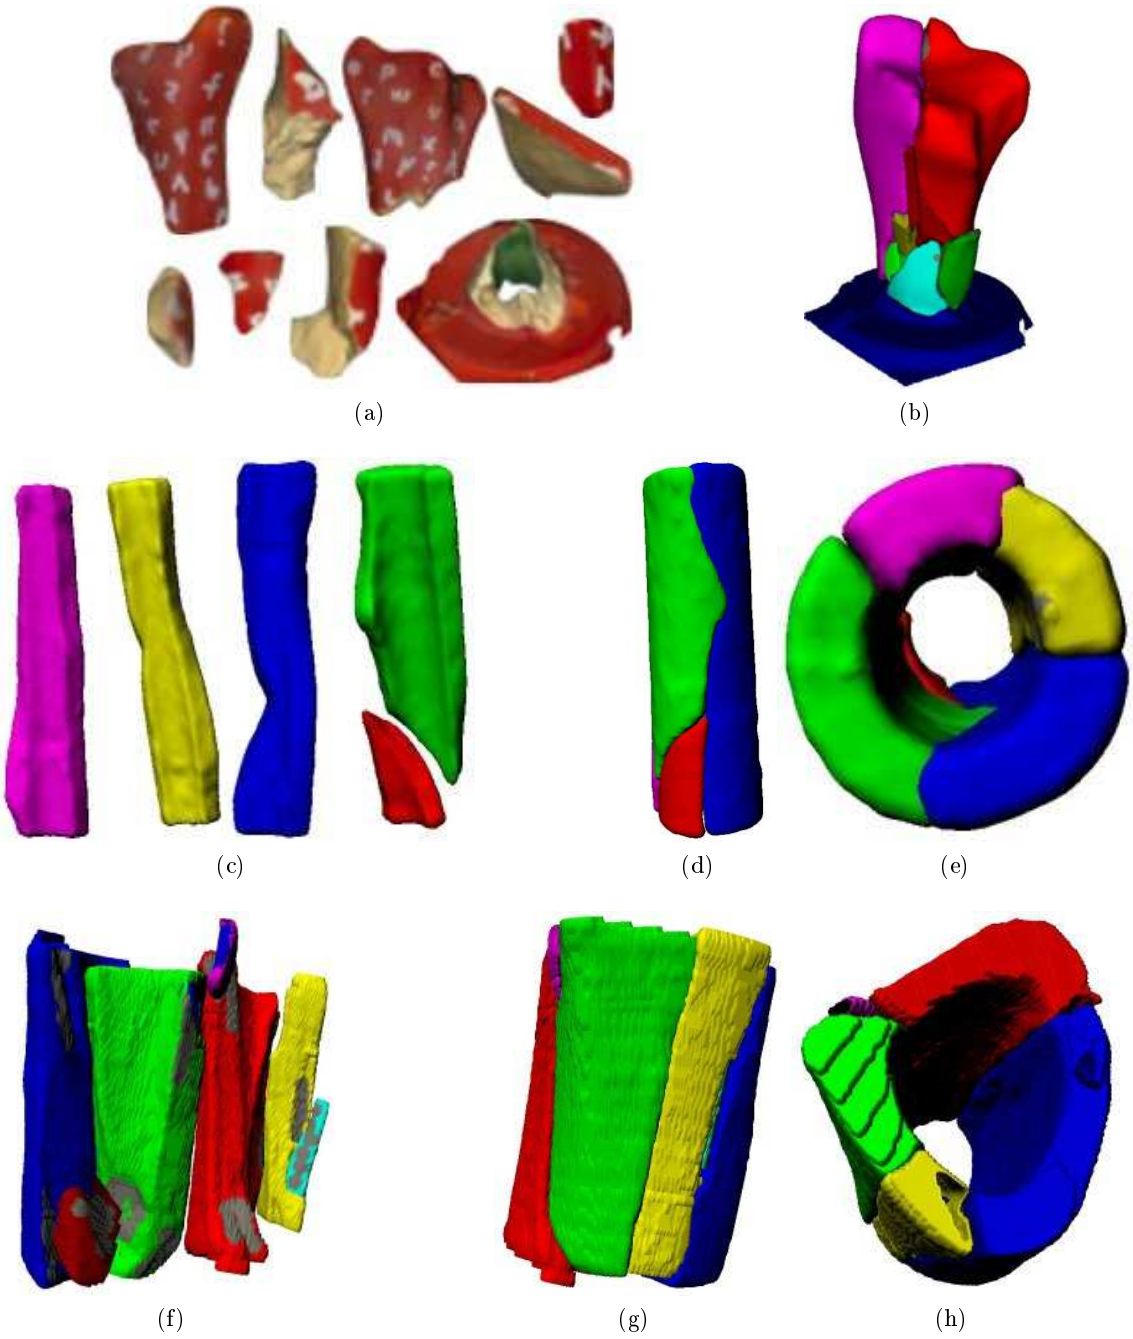

Figure 7: (a,b), and (c,d) show two views of three fracture cases which have been reconstructed using the system (see §7 Results for details).

9. ACKNOWLEDGEMENTS

This work was supported by NIH grants P50AR055533 and 1R21AR054015 and by a Medical Research Initiative Grant from the Roy J. Carver Charitable Trust at The University of Iowa.

REFERENCES

1. J. L. Marsh, S. Bonar, J. V. Nepola, T. A. Decoster, and S. R. Hurwitz, "The use of an articulated external fixator for fractures o fthe tibial plafond," *Journal of Bone and Joint Surgery*, vol. 77A, pp. 1498–1509, 1995.
2. J. L. Marsh, D. P. Weigel, and D. R. Dirshl, "Tibial plafond fractures. how do these ankles function over time?," *Journal of Bone and Joint Surgery*, vol. 85A, pp. 287–295, 2003.
3. O. Ron, L. Joskowicz, A. Simkin, and C. Milgrom, "Computer-based periaxial rotation measurement for aligning fractured femur fragments: Method and preliminary results," *Lecture Notes in Computer Science*, no. 2208, pp. 17–23, 2001.
4. M. Scheuering, C. Rezk-Salama, C. Eckstein, and G. G. K. Hormann, "Interactive repositioning of bone fracture segments," in *Proceedings of Vision, Modeling and Visualization (VMV)*, pp. 499–505, 2001.
5. S. Winkelbach and F. M. Wahl, "Pairwise matching of 3d fragments using cluster trees," *International Journal of Computer Vision*, vol. 78, pp. 1–13, June 2008.
6. M. Harders, A. Barlit, C. Gerber, J. Hodler, and G. Székely, "An optimized surgical planning environment for complex proximal humerus fractures," in *MICCAI Workshop on Interaction in Medical Image Analysis and Visualization*, 2007.
7. A. Willis, D. Anderson, T. Thomas, T.and Brown, and J. Marsh, "3d reconstruction of highly fragmented bone fractures," in *Proceedings of the SPIE Conference on Medical Imaging 2007: Image Processing*, vol. 6512, pp. 65121P1–65121P10, 2007.
8. C. L. Beardsley, A. D. Heiner, E. A. Brandser, J. L. Marsh, and T. D. Brown, "High density polyetherurethane foam as a fragmentation and radiographic surrogate for cortical bone," *The Iowa Orthopaedic Journal*, vol. 20, pp. 24–30, 2000.
9. D. D. Anderson, V. Muehling, J. Marsh, and T. Brown, "Precise identification of bone fragment boundaries to assist in reduction of highly comminuted fractures," *Computer-Aided Surgery*, vol. 9, no. 3, p. 116, 2004.
10. K. Pulli, "Multiview registration for large data sets," in *Proc. Int. Conf. on 3D Digital Imaging and Modeling*, pp. 160–168, 1999.
11. P. Besl and N. McKay, "A Method for Registration of 3-D Shapes," *PAMI*, vol. 14, no. 2, pp. 239–256, 1992.
12. Y. Chen and G. Medioni, "Object modelling by registration of multiple range images," *Image and Vision Computing*, vol. 10, pp. 145–155, 1992.
13. A. Johnson and S. Kang, "Registration and integration of textured 3-d data," in *Proc. Int. Conf. on 3D Digital Imaging and Modeling*, pp. 279–290, 1994.
14. G. Godin, M. Rioux, and R. Baribeau, "Three-dimensional registration using range and intensity information," in *SPIE*, vol. Vol. 2350: Videometrics III, pp. 279–290, 1994.
15. N. Gelfand, L. Ikemoto, S. Rusinkiewicz, and M. Levoy, "Geometrically stable sampling for the icp algorithm," in *Proc. International Conference on 3D Digital Imaging and Modeling (3DIM)*, pp. 260–267, 2003.
16. S. Rusinkiewicz and M. Levoy, "Efficient Variants of the ICP Algorithm," in *Intl. Conf. on 3D Digital Imaging and Modeling*, pp. 145–152, 2001.