

Automatic Fracture Reduction

Thomas Albrecht and Thomas Vetter

University of Basel

Abstract. We present a method to automatically reposition the fragments of a broken bone based on surface meshes segmented from CT scans. The result of this virtual fracture reduction is intended to be used as an operation plan for a medical procedure. Particularly in minimally invasive surgery like intramedullary nailing, the correct repositioning of bone fragments is not always apparent or visible without an operation plan. We propose to achieve automatic fracture reduction by fitting the bone fragments to an intact reference bone mesh with a modified Iterative Closest Point (ICP) algorithm. A suitable reference could be the same patient's contra-lateral bone. In the absence of a CT scan of this bone, we propose to use a statistical shape model as a reference. The shape model is automatically adapted to match the anatomy of the broken bone, apart from the bone's length, which has to be correctly initialized. Our experiments show that we can limit the rotational alignment error to below 5 degrees, compared to 15 degrees in current medical practice.

1 Introduction

Fracture reduction, i.e. the task of repositioning the fragments of a broken bone into their original position is a common task in everyday medical practice. For many fractures, the correct repositioning is apparent and straight-forward to carry out in practice. For some fractures however, an accurate reduction is difficult to achieve, because the desired position of the fragments is difficult to deduce from the available medical images. The most widely researched fracture in this area, which has become somewhat of a model problem, is the femoral shaft fracture. Figure 1 shows an example, along with its minimally invasive treatment by intramedullary nailing. In this procedure, a long nail is inserted into the bone via a small incision at the hip or knee. The fracture site is not directly visible, and the surgeon has to rely on radiographs to align the fragments. This allows a fairly accurate repositioning in the image planes of the radiographs. The rotational alignment around the longitudinal axis of the bone poses a much greater challenge, as it cannot be observed in these radiographs. A recent retrospective clinical study has found a rotational malalignment of over 15 degrees in 28% of patients [5]. In our experiments we were able to limit the malalignment to 5 degrees. While we focus on this model problem in this paper, our method has the advantage that it can be applied to virtually any bone and fracture.



Fig. 1. A broken femur (left) treated with an intramedullary nail (right). Even in views from two perspectives, it is hard to judge the rotational alignment of the fragments.

1.1 Reduction Method

In our reduction method, the fragments of the broken bone are represented as surface meshes generated from segmented CT scans. The main idea is to align the main and functionally important fragments to an intact reference bone. This ensures an anatomically correct repositioning of the main fragments, independent of the geometry of the fracture surfaces and possible small additional fragments. Figure 2 shows an example of our proposed reduction method. The repositioning of the smaller fragments seen in Figure 2(d) will be postponed until Section 2.4.

In order to apply our method we need a mesh representation of a reference bone that is as similar to the broken bone as possible. If a CT scan of the broken bone from before the fracture, or a scan of the contra-lateral bone are available, these can be used to generate a reference mesh. In many cases however, such scans will not be available and for these we propose to use a statistical shape model [2] as reference. The mesh of the statistical shape model is automatically adapted to the characteristics of the fragments during the reduction in order to reflect the shape of the specific broken bone as accurately as possible. We performed experiments on real and simulated fractures. They show the importance of correctly adapting the shape model. We achieve to limit the rotational alignment error to 5 degrees when using a statistical shape model and to 0 degrees when using the ground truth bone as reference.

Prior Work. A few groups have worked on automatic bone fracture reduction. In [7], Moghari and Abolmaesumi align fragments of the proximal femur to the mean of a statistical shape model built from only 5 bones. However, they do not adapt the shape model to the individual anatomy of the broken bone. In [4], Gong et al. perform a reduction of distal radius fractures that uses a 3D statistical shape model to help align the fragments in calibrated 2D x-ray images. In [9] Westphal, Winkelbach et al. address femoral shaft fractures in terms of a “3D puzzle problem”. The individual fragments are aligned based on the fracture surfaces, i.e. the surfaces of the fracture site which should be rejoined by the fracture reduction. We argue however that such a method may depend heavily

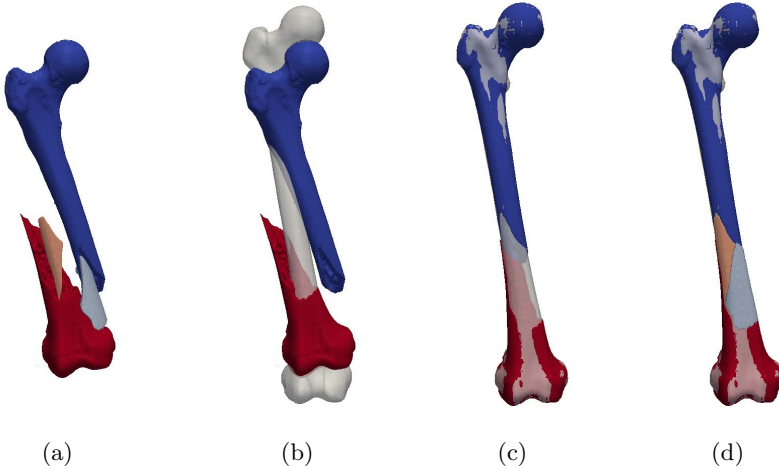


Fig. 2. Our proposed reduction algorithm. Instead of directly trying to puzzle the fragments of a broken bone (a) together, we align the main fragments to a reference bone (b), ensuring their functionally correct placement (c). If desired, the remaining fragments can later be fitted to the remaining parts of the reference (d).

on the accurate scanning, segmentation, and alignment of these fracture surfaces. A very small error in any of these steps, a shattering or compression of the bone structure could cause a large malalignment at the functionally important ends of the bones.

Any reduction plan, will have to be used in conjunction with a navigation or robot system to provide the surgeon with visual and/or haptic feedback about the current and desired position of the fragments. For such systems, we refer to [9] and references therein.

2 The Reduction Algorithm

Our method consists of two steps: 1. A rigid alignment of the fragments to the reference (Section 2.1). 2. The adaption of the statistical shape model to the given bone fragments’ individual anatomy (Section 2.2). The adaption is only possible if the fragments are already aligned to the model. The rigid alignment in turn depends on a good model adaption. We propose to solve these two steps simultaneously in an iteration scheme, alternating step 1 and 2. If the ground truth or contra-lateral bone is used as reference, step 2 is omitted.

2.1 Rigid Alignment

For the rigid alignment of each of the fragments, we use the well known Iterative Closest Point (ICP) mesh alignment algorithm as implemented in VTK [1]. For each “source point”, i.e. a point on the fragment, the ICP algorithm identifies

a “target point”, which is the closest points on the reference bone. Then, it computes the unique rigid transformation that minimizes the distance between the source and target points. This method is only guaranteed to give a correct alignment if the source and target points are *corresponding* points, which is in general not the case initially. But if the method is reasonably initialized, i.e. the shapes are roughly pre-aligned, and is iterated several times, the ICP algorithm has shown to provide a fast and accurate alignment of the shapes.

2.2 Model Mesh Adaption

If a statistical shape model is used as the reference bone, we adapt its shape to agree as well as possible with the fractured bone. After the rigid alignment, the individual fragments are already approximately aligned with the current model instance. Therefore, we can assume that source-target point pairs from the last ICP iteration, are not only pairs of *closest* points but also pairs of *corresponding* points, at least approximately. With a set of aligned and corresponding points, it is straight-forward to adapt the model shape to the fragments:

A (linear) statistical shape model represents instances of an object class (e.g. the human femur bone), as a sum $\mathbf{T}(\boldsymbol{\alpha}) = \boldsymbol{\mu} + \mathbf{U}\boldsymbol{\alpha}$ of a mean mesh $\boldsymbol{\mu}$ and a linear combination of shape deformations, expressed as the product of a model matrix \mathbf{U} and coefficients $\boldsymbol{\alpha}$. For more information see [2,6].

Now let us denote the source points on fragment i as \mathbf{S}_i , and the corresponding target points on the reference model as \mathbf{T}_i . Being points of a model instance, the latter can be represented as $\mathbf{T}_i(\boldsymbol{\alpha}) = \boldsymbol{\mu}_i + \mathbf{U}_i\boldsymbol{\alpha}$, where $\boldsymbol{\mu}_i$ and \mathbf{U}_i are the appropriate sub-vector and sub-matrix of $\boldsymbol{\mu}$ and \mathbf{U} . Then, to adapt the model mesh, we wish to find a set of shape model coefficients $\boldsymbol{\alpha}$ such that the mean square difference between all point pairs given by $(\mathbf{S}_i, \mathbf{T}_i(\boldsymbol{\alpha}))$ is minimized:

$$\boldsymbol{\alpha} = \operatorname{argmin}_{\boldsymbol{\alpha}} \sum_i \|\mathbf{T}_i(\boldsymbol{\alpha}) - \mathbf{S}_i\|^2 = \operatorname{argmin}_{\boldsymbol{\alpha}} \sum_i \|\boldsymbol{\mu}_i + \mathbf{U}_i\boldsymbol{\alpha} - \mathbf{S}_i\|^2. \quad (1)$$

By $\|\cdot\|^2$ we denote the sum of squared distance norm for point sets, i.e. $\|\mathbf{P}\|^2 := \sum_{p \in \mathbf{P}} \|p\|_2^2$ for $\mathbf{P} \subset \mathbb{R}^3$. Equation (1) is a convex optimization problem which admits a unique minimum $\boldsymbol{\alpha}$. It can furthermore be regularized in order to permit only parameters $\boldsymbol{\alpha}$ that represent plausible instances from the shape model. [2,6]

The model adaption depends on the accuracy of finding corresponding points by the rigid alignment and vice versa. Iterating both steps alternatively improves them in turn and is repeated until a given number of iterations or a convergence criterion is attained.

2.3 Initialization

As most iterative algorithms, our reduction method relies on a reasonable initialization. Regarding the rigid alignment, our experiments have shown that it suffices to place the reference and the fractured bone roughly in the same area and orientation. This amounts to placing the reference in the center of the CT scan.




Fig. 3. Difficulty in bone length estimation. The main fragments can be fitted equally well to a reference model with incorrect length (a) (b), because bone length and shape are relatively independent. Only with the middle fragment do we realize the correct length (c). This seems less feasible in a complex fracture like (d).

Regarding the shape adaption, the experiments have shown that for long bones it suffices to start the algorithm with the shape model’s mean bone, with one important exception: The length of the bone should be initialized as accurately as possible.

The main principle of our algorithm is that only a correctly adapted shape will minimize the model adaption, which in turn causes a correct alignment. However, Figure 3 illustrates how this can fail if the length is not correctly initialized. In long bones like femur, tibia, humerus, etc., the length of the shaft is relatively independent of the shape and size of the joint regions. Therefore, the joint regions do not carry enough information for determining the length during model adaption. The shaft is rather featureless and can be fitted equally well to bones of almost any length.

We therefore define the length as a mandatory user parameter. It can be estimated by the user/physician for instance by measuring the length of the patient’s contra-lateral bone or summing up the length of the fragments. Relying on the length of the individual fragments to determine the bone length automatically seems unfeasible for complicated fractures like Figure 3(d).

2.4 Implementational Details

We have improved the accuracy of the basic algorithm with a few enhancements.

Finding Closest Points. Both the ICP algorithm and the model adaption rely on finding closest points as an approximation to *corresponding* points in the source and target meshes. Typically, this approximation gets better with each iteration and in a perfect fit the two meshes and the corresponding points should

almost coincide with each other. In reality however, the ICP can converge to a local minimum where the matched closest points are not necessarily corresponding points. This is a well-known shortcoming of the ICP algorithm and has been discussed in the literature. We use a method proposed by Feldmar and Ayache in [3] to take the curvature and the normal vectors of the two meshes into account when searching for closest points, in order to have a higher chance of finding corresponding points. This is relatively straight-forward to include into the ICP algorithm [8] by re-implementing the sub-routine that searches for the closest point on the target mesh for a given source point. Even though this method still is not guaranteed to find corresponding points, it improves the results significantly.

Fitting Remaining Fragments. Finally, as we stated above, our main goal is to align the main fragments of a broken bone to their anatomically correct position. We do not aim at solving the “3D puzzle problem” [9], nor do we wish to rely on the shape of the fracture surfaces for bone alignment. Nevertheless, it is unsatisfactory to leave small fragments that were left out of the original fitting unaligned. Therefore, for easy cases, we propose an ad-hoc way to align the remaining fragments after the main fragments have already been aligned and the reference model has been adapted to their shape. At this stage, the aligned fragments and the reference model almost coincide, except in those parts where the remaining fragments have been left out, see Figure 2(c) for an example. We can therefore align the remaining fragments to these remaining parts of the model mesh with an additional ICP alignment. We do not claim that complex fractures like that in Figure 3(d) can be reduced in this way.

3 Evaluation

Figures 2 and 3 show successful reductions of real femur fractures. Because no ground truth data is available for these real fractures, we have additionally evaluated our method on artificial fractures. For this purpose, we separated a database of 145 human femora into a training set of 120 bones, from which we built our statistical model, and a test set of 25 bones. In order to test the reduction method on these intact bones, we cut the bone meshes in two parts and randomly displaced these two fragments. Figure 4(a) shows a few of these simulated fractures. The fracture site, angle, rotation, and translation were drawn from a uniform distribution on varying intervals. We used an interval of $[-100, 100]$ mm for the offset of the fracture from the middle of the bone, $[-20, 20]^\circ$ for each component of the rotation represented by Euler angles, and $[-30, 30]$ mm for the translation. The algorithm was initialized with the correct bone length.

Ideally, the reduction algorithm should compute the inverse of the random transforms, putting the fragments back into place. The composition of the random transforms with the reduction results should be the identity map. By measuring the deviation from the identity, we can evaluate the accuracy of the reduction. Figure 4(b), visualizes the results of our experiment. We measured the

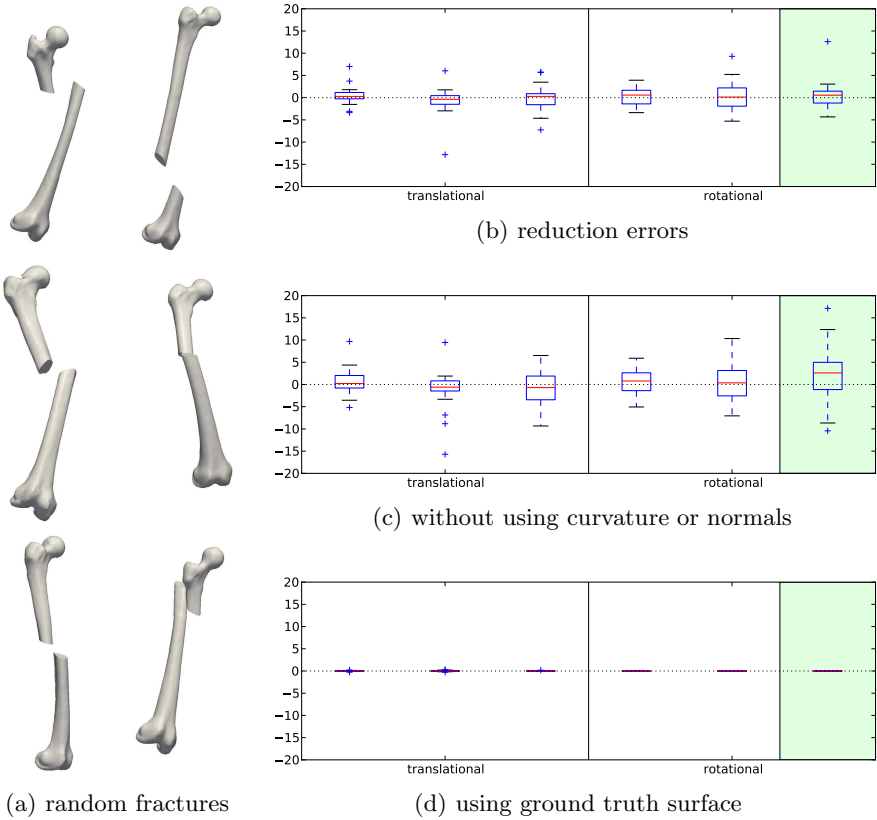


Fig. 4. Evaluation: Box plots of the alignment errors of the reduction algorithm

translation error along and the rotation error around the left-right (x), anterior-posterior (y) and superior-inferior (z) axis of the bone, in that order. The errors from the two bone fragments were added to get a measure of the overall reduction accuracy. The most important plot for our application scenario is therefore the very rightmost plot. It represents the rotational malalignment of the two fragments along the long axis of the bone and corresponds to the value measured in [5]. We see that, besides one outlier, we have successfully limited the amount of malrotation to within $\pm 5^\circ$, with a mean of 0.44° and a standard deviation of $\sigma = 3.13^\circ$. If these results can be carried through in the operating room, this will mean a massive improvement over the results reported in [5].

A closer look at the single outlier revealed that it exhibits a bone deformity known as *coxa retrotorta*, in which the relative rotation of the distal and proximal joint area, known as the *antetorsion* of the femur, is close to 0° . The statistical shape model on the other hand favors more common bone shapes with normal antetorsion angles. Extreme cases as this outlier will always have to be individually evaluated by a medical professional.

Figure 4(c) reveals that without using the curvature and surface normals in the fitting, the error is higher ($1.54 \pm 6.26^\circ$), because the model mesh is not as well adapted to the individual anatomy. If the original ground truth bone is used as reference, Figure 4(d), the reduction error is almost zero, ($-0.002 \pm 0.03^\circ$).

4 Conclusion

We have presented a method for automatic fracture reduction by aligning bone fragments to an intact reference bone. In the absence of a ground truth or contralateral bone as reference, we use a statistical shape model which is adapted to the broken bone's individual anatomy. While automatic bone length estimation remains a challenge, we have shown visually plausible reductions of real bone fractures and experimental errors that are far smaller than the current target error of 15 degrees rotational malalignment in medical practice. Future work will need to address the implementation of these reduction plans in the operating room using a navigation system.

References

1. Besl, P., McKay, H.: A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14(2), 239–256 (1992)
2. Blanz, V., Vetter, T.: Reconstructing the complete 3d shape of faces from partial information. *Informationstechnik und Technische Informatik* 44(6), 1–8 (2002)
3. Feldmar, J., Ayache, N.: Rigid, affine and locally affine registration of free-form surfaces. *International Journal of Computer Vision* 18(2), 99–119 (1996)
4. Gong, R., Stewart, J., Abolmaesumi, P.: Reduction of multi-fragment fractures of the distal radius using atlas-based 2d/3d registration. In: *Proceedings of the SPIE, Medical Imaging*, p. 726137–1 (2009)
5. Jaarsma, R., Pakvis, D., Verdonchot, N., Biert, J., van Kampen, A.: Rotational malalignment after intramedullary nailing of femoral fractures. *Journal of Orthopaedic Trauma* 18(7), 403 (2004)
6. Lüthi, M., Albrecht, T., Vetter, T.: Probabilistic Modeling and Visualization of the Flexibility in Morphable Models. In: Hancock, E.R., Martin, R.R., Sabin, M.A. (eds.) *Mathematics of Surfaces*. LNCS, vol. 5654, pp. 251–264. Springer, Heidelberg (2009)
7. Moghari, M.H., Abolmaesumi, P.: Global registration of multiple bone fragments using statistical atlas models: feasibility experiments. In: *30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS 2008*, pp. 5374–5377. IEEE (2008)
8. Schroeder, W., Martin, K., Lorensen, B.: *The visualization toolkit*. Prentice Hall PTR (1998)
9. Westphal, R., Winkelbach, S., Gössling, T., Oszwald, M., Hüfner, T., Krettek, C., Wahl, F.: Automated robot assisted fracture reduction. In: *Advances in Robotics Research*, pp. 251–262 (2009)