Asef Islam

4/16/20

CIS 2: Spring 2020

Group 10: 3D Imaging, Analysis, and Guidance for Robotic-Assisted Ankle Fracture / Dislocation Surgery

Critical Review: Application of Active Shape Models to Automatic Segmentation of Ankle Joint

Project Overview

The surgical repair process for a fibular ankle fracture involves two major steps: repair of the proximal fracture and syndesmotic reduction. The repair of the immediate fracture itself can be done relatively simply using screw-based fixation to promote re-fusion of the bone. However, any fracture of this nature can result of displacement of the lateral malleolus of the fibula in the ankle joint. Thus, it is necessary to reposition the fibula to restore syndesmosis, or natural alignment of bones within the ankle joint complex. Failure to do so adequately can result in a high incidence of osteoarthritis, which is highly problematic in the current standard of care for the reduction, which is visual estimation and manual handling by the surgeon.

The aim of this project is thus to introduce an image-guided approach to repositioning the shifted fibula after it is displaced during a high ankle fracture in order to restore syndesmosis. The envisioned workflow would be to take a scan of the patient's ankle, automatically segment it to identify where the bones are positioned and model them as surfaces in space, compute registration from the injured fibula to a healthy one, and use that to instruct a robot to move the fibula. The most crucial step is the automatic segmentation, as this will often define the accuracy of the registration and ultimate reduction. One approach to automatic segmentation is active shapes. There has been a lot of past work over the last few years in developing active shape frameworks for automatic segmentation of the ankle, and this project was seen as an application of such.

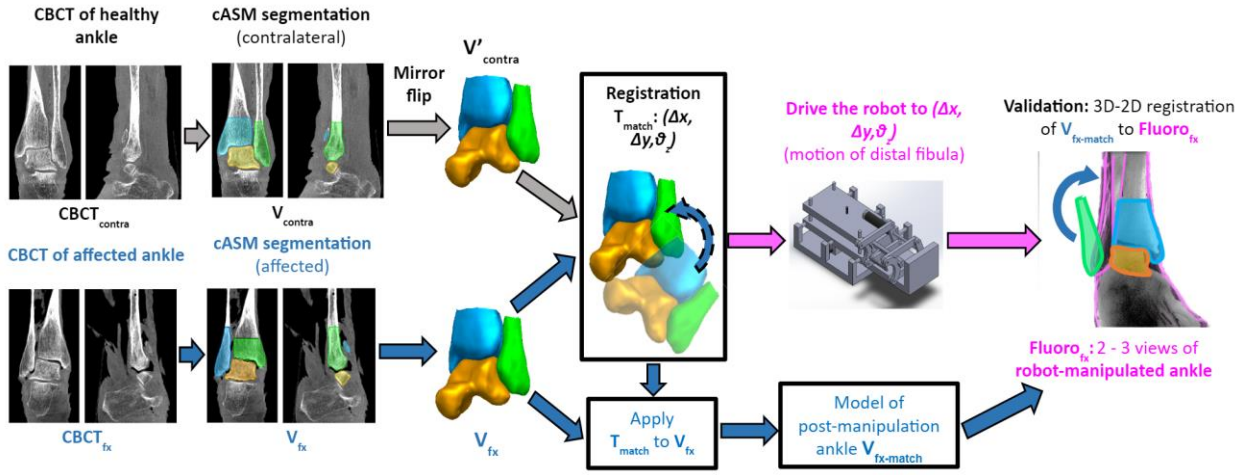

Papers Selected

[1] Görres, J., Brehler, M., Franke, J. *et al.* Articular surface segmentation using active shape models for intraoperative implant assessment. *Int J CARS* **11,** 1661–1672 (2016).

This was the first major attempt to apply active shape models (ASMs) to automatic segmentation of ankle bones.

[2] Brehler M, Islam A, Vogelsang L, et al. Coupled Active Shape Models for Automated Segmentation and Landmark Localization in High-Resolution CT of the Foot and Ankle. *Proc SPIE Int Soc Opt Eng*. 2019.

This work improved upon the traditional ASMs by introducing Coupled Active Shape Models (cASMs) with proximity constraints.

Paper 1

Purpose and Significance

This paper sought to apply ASMs to automatic segmentation of ankle bones and envisioned it to be useful in any surgical guidance application, but in particular were exploring its clinical usefulness in the application of locating surgical intra-articular implant screws that are inserted through the calcaneus into the articular region between the calcaneus and talus. Thus, in this case accurate segmentation is essential for proper localization of the edges of the two bones, and thus the articular space between them. The paper sought to prove that the segmentations were accurate, and also that it would be useful for this implant localization application.

Background

Active Shape Models should be introduced. As with most models for artificial intelligence or computer vision applications, these models are trained on a training set, which in this case is a set of manually created and verified segmentations of the bones of interest. After creating the segmentations, the positions of many corresponding landmark points (on the order of the minimum number of vertices on the segmentations on the training set) is also annotated on them.

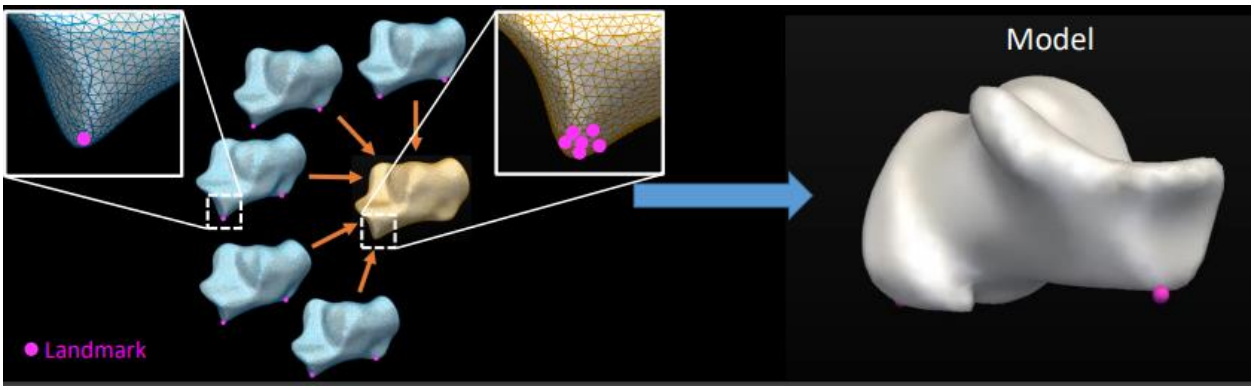

These points are represented in 3D space and are fed into the model to summarize the morphological variance in the training population. Each entire training segmentation can thus be represented as a point in 3n space, where n is the number of landmark points. Then, Principal

Component Analysis (PCA), a common algorithm for variance-preserving dimensionality reduction, is used to project this representation into the t-dimensional subspace that preserves the most variance of the data. Hence, the dimensionality of the training data has been reduced from 3n to t principal modes of variance, where t is specified for the particular application. Any example x in the training set can then be represented as the mean plus the deviation from the mean, or $\bar{x} + Pb$, where $\bar{x}$ is the average shape of the whole training set, P is the projection matrix from 3n to t produced by PCA, and b is a t-vector representing the deviation of the particular example from the mean along each mode. Thus, this can also be thought of as representing a deformable template that can be morphologically transformed by adjusting the variation from the mean along each mode in the b vector to encompass all of the shape variation of the bone within the population.

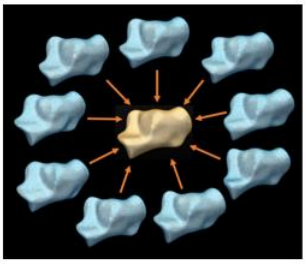

n landmark points
t principal components

$$\mathbf{x} \approx \bar{\mathbf{x}} + \mathbf{Pb}$$

x and $\bar{x}$: 3n x 1
P: 3n x t
B: t x 1

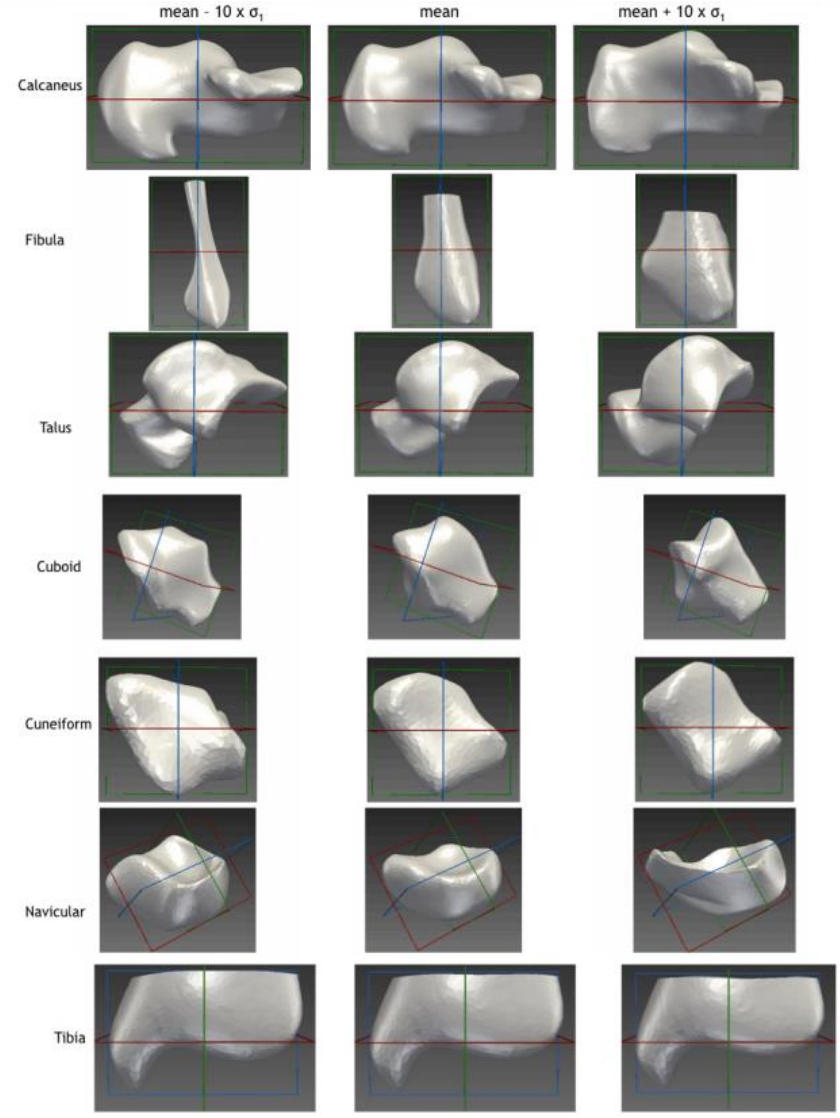

The above example shows the mean shape, as well as the two shapes representing ± 10 standard deviations from the mean for the 7 bones of the ankle. Although some differences are notable, for many of these bones the shapes at 10 standard deviations appear more visually similar to the mean than one would expect for that many standard deviations. This reveals that in general the morphology of these bones is quite similar in the population and the differences are very subtle. The tibia and fibula are the two long shinbones and extend from the ankle joint up the leg all the way to the knee joint, but in this case since only their surface at the ankle joint is of interest their segmentations are truncated at a set distance up from the ankle.

ASMs can be useful not only in understanding the variations between people in the shape of their bones, but more interestingly for surgical applications and particularly automatic segmentation. From a CT scan of a patient's ankle, the gradient version of the image can be taken to highlight just the edges, where the gradient is high. The mean models of each bone are then initialized to the image by rigid registration to the gradient edges. Then, in an iterative process to update the shape vector b, the ASM for each bone is deformed and shifted in the positions of the vertices so that they best fit the edges of the image.

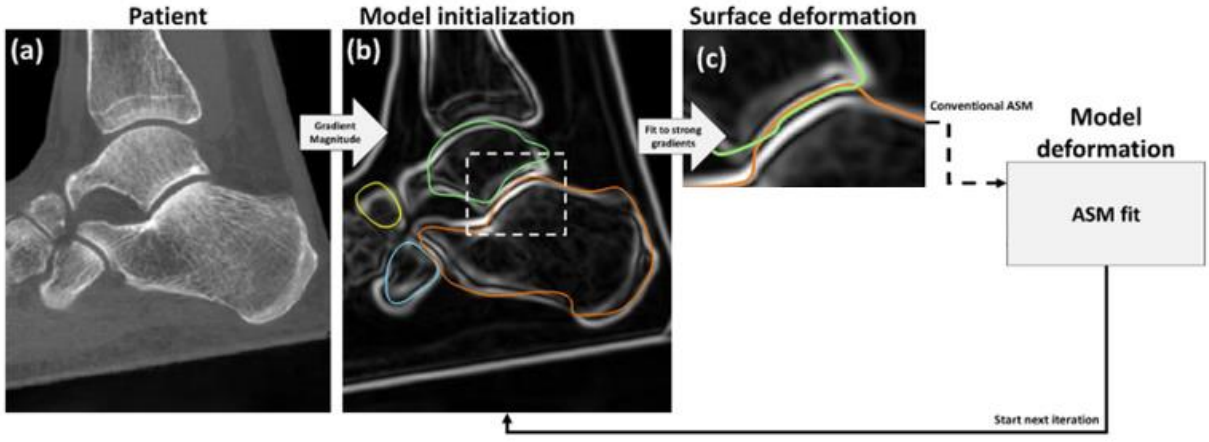

Calculating the amount of error in any automatic segmentation first requires a manual ground truth segmentation. Both the test and truth segmentations are represented as a mesh of vertices in space. Then, the surface error is calculated as the sum of squared distances between paired points on the two meshes, and the root mean square distance can be taken across all points.

$$\text{surf\_err}_i = \sqrt{(x_i - x_{nn,i})^2 + (y_i - y_{nn,i})^2 + (z_i - z_{nn,i})^2}$$

$$\text{RMSD} = \sqrt{[\sum_1^N (surf\_err_i)^2] / N}$$

Methodology

Paper 1 followed this process in using ASMs for automatically segment the 3 bones of interest the calcaneus, talus, and cuboid, and thus map the edges and articular regions between them. The

clinical exploration came in the form of a cadaver study to detect surgical implants by approximating them as cylinders in the image.

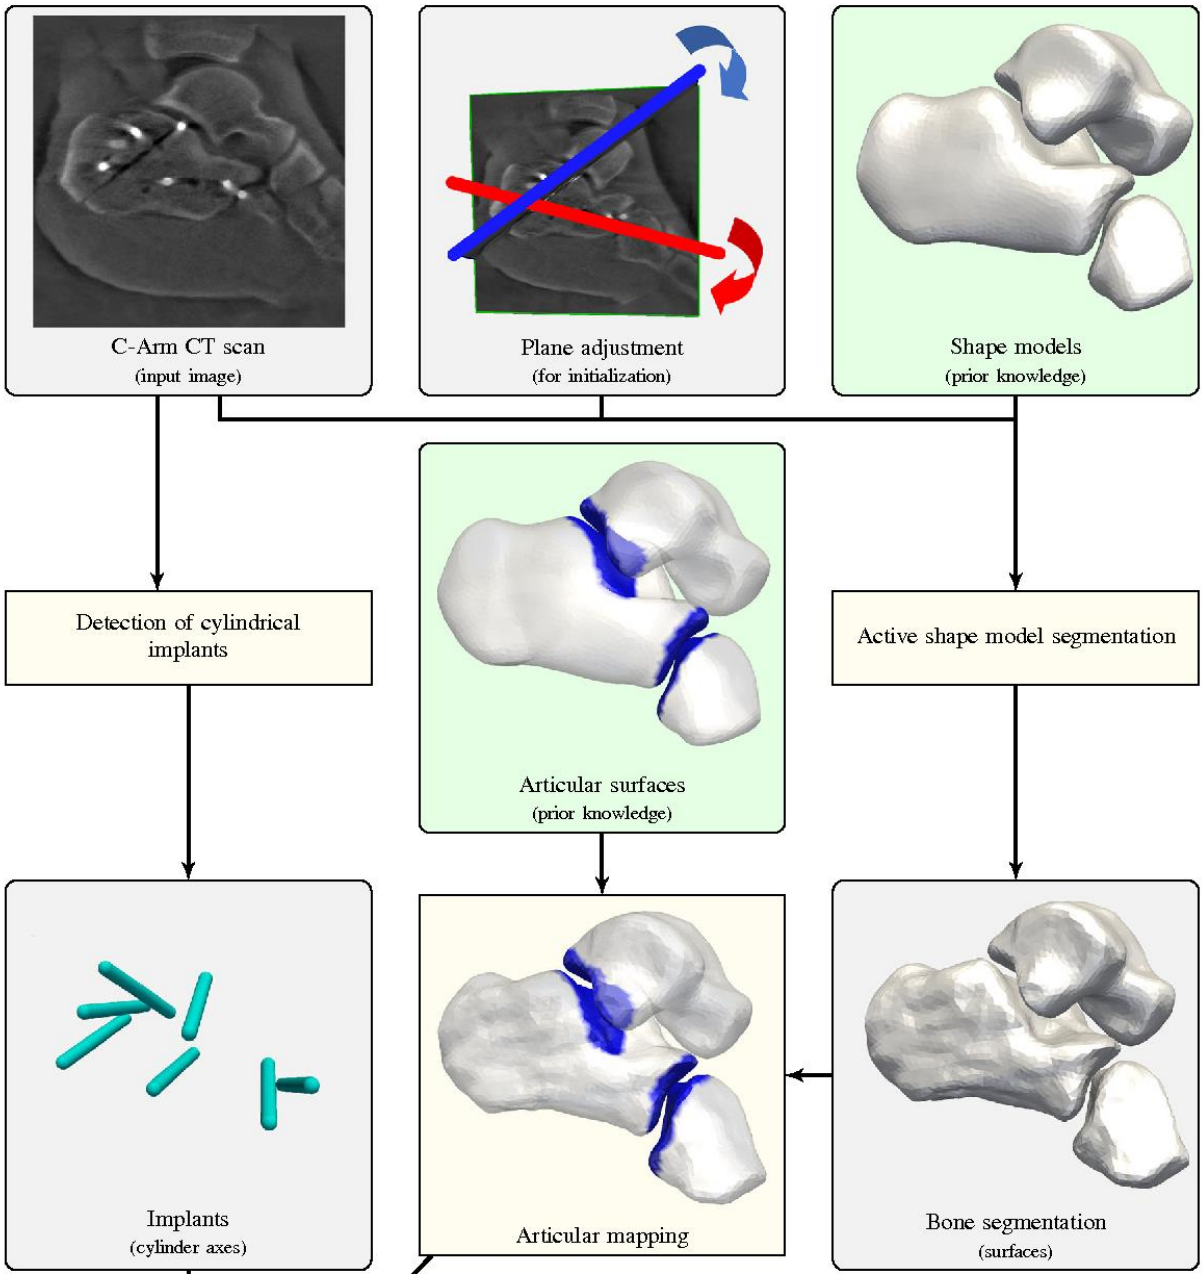

Implants were inserted into cadaver ankles, and then scans taken. The locations of implant articular intersection were manually annotated in the images as ground truth. Then, these locations were automatically detected based on the surface segmentations and articular edge mappings by intersection detection, and the errors were calculated.

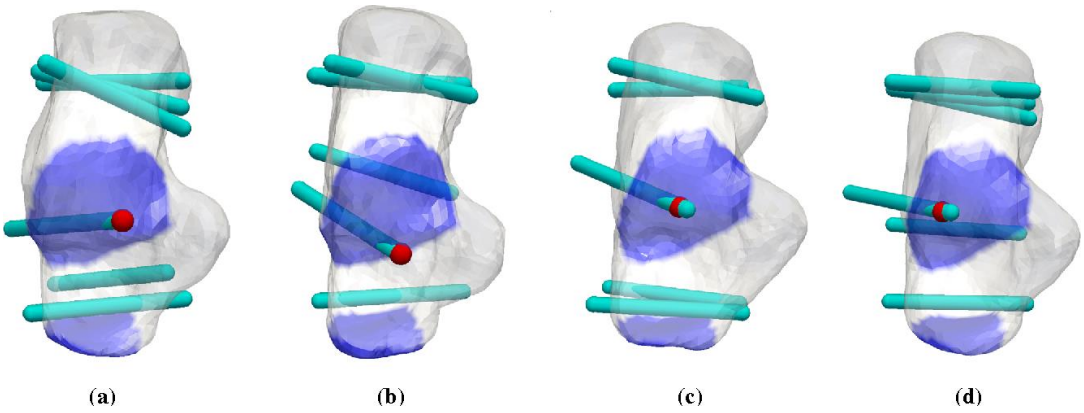

(a) (b) (c) (d)

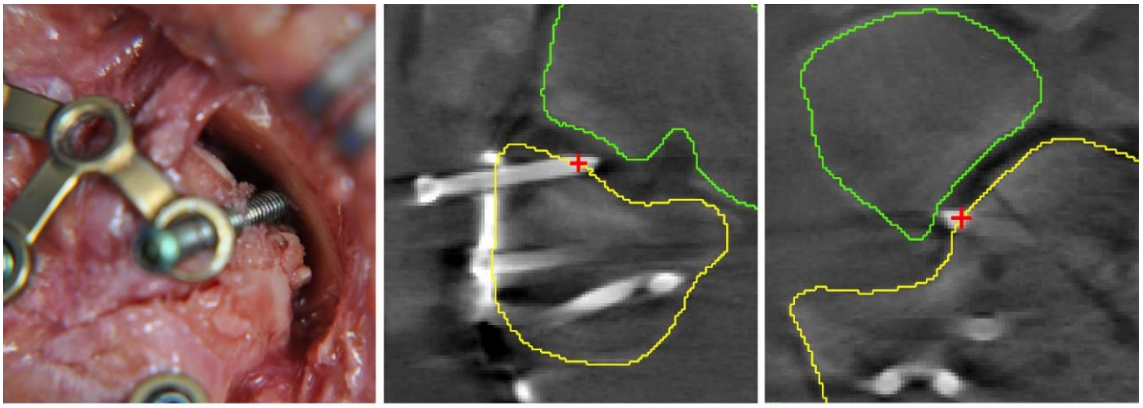

(a) (b) (c)

Results and Significance

The RMS surface segmentation errors were 0.81, 0.45 and 0.59 mm for the calcaneus, talus, and cuboid respectively. It was found that the algorithm was able to compensate suboptimal initializations of up to 16 mm translation and 21 degrees of rotation off. In the cadaver experiment, the automatically detected positions of the implants had a 0.80 mm average error. Thus, reliable segmentation and ability to accurate localize implants in the clinical application were shown.

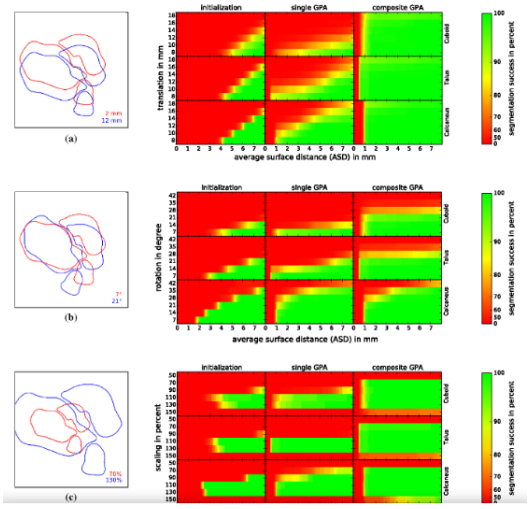

Method bone	ASD bone (mm)	ASD LA (mm)	ASD CC (mm)
Multiple ASMs with optimal surface integration			
"Active shape models" section			
Calcaneus	0.81 ± 0.13	0.76 ± 0.29	0.66 ± 0.24
Talus	0.45 ± 0.18	0.44 ± 0.09	
Cuboid	0.59 ± 0.27		0.38 ± 0.10
Model-independent regions			
"Model-independent regions" section			
Calcaneus	0.79 ± 0.14	0.59 ± 0.29	0.46 ± 0.16
Talus	0.44 ± 0.17	0.34 ± 0.07	
Cuboid	0.57 ± 0.26		0.30 ± 0.09
Deformable surface extension			
Heimann et al. [15]			
Calcaneus	0.65 ± 0.12	0.55 ± 0.21	0.44 ± 0.24
Talus	0.36 ± 0.11	0.37 ± 0.15	
Cuboid	0.57 ± 0.27		0.31 ± 0.12

The errors at articular surface regions are separately shown for the whole bone surface (bone), the lower ankle joint (LA), and the calcaneocuboid joint (CC)

Assessment

The paper validated the theoretical ASM framework for automatic ankle segmentation and this surgical application by showing sub-mm accuracy of both surface segmentation and implant localization. In particular, the cadaver study demonstrated the potential clinical use for this application. However, this study only focused on 3 bones, whereas there are 7 total in the ankle joint. In addition, although there was compensation of initialization error to some extent, the segmentation would fail in the case of excessively inaccurate initialization. In addition, although the cadaver study showed sub-mm accuracy when using manual visual annotations as a ground truth, it didn't necessarily compare to a surgical standard of care for implant localization.

Paper 2

Purpose and Significance

The aim of this paper was to improve traditional ASMs especially in narrow joint spaces, and prevent overlaps that are common due to poor initialization in the case of unreliable gradients due to low contrast or noise. It included all 7 bones of the ankle, and took into account the distances between them in enforcing proximity constraints. This new framework was termed coupled active shave models (cASMs).

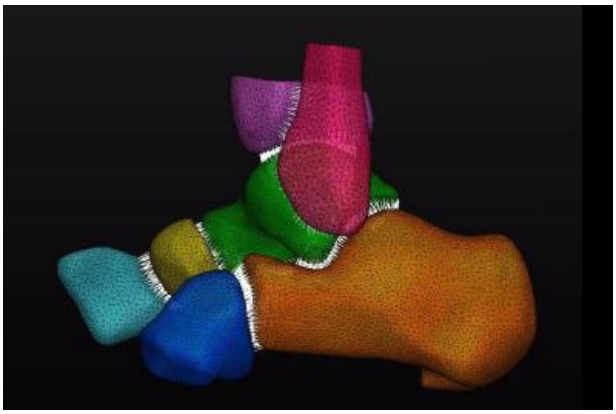

Methodology

The algorithm is the same as traditional ASMs, with the added feature of overlap detection due to triggering of proximity conditions within each iteration. Traditional ASMs fit each bone individually, but cASMs jointly fit multiple bones together and can thus take into account distances between them. Thus, it can enforce proximity constraints by detecting if there is an overlap between bones during the surface deformation iterations, and then correcting them before proceeding to the next iteration.

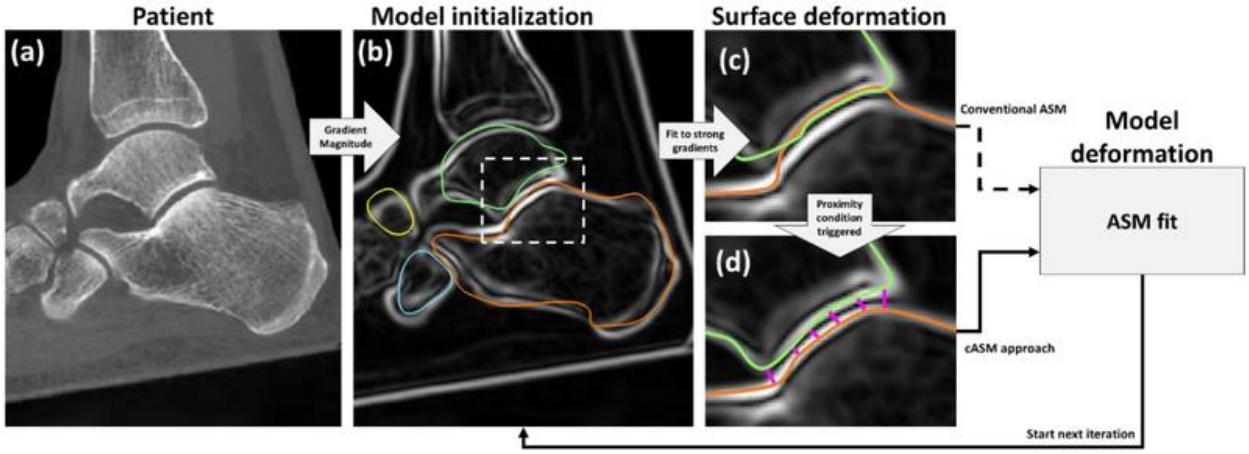

Results and Significance

The cASM framework was tested on a set of 21 ankle segmentations through a leave one out scheme, where each subject was used as a test example once and the remaining 20 used to train the model each time. Due to the strict proximity constraints that are enforced, cASM segmentations never allowed overlap, and detected and corrected them in 50-80% cases for each bone.

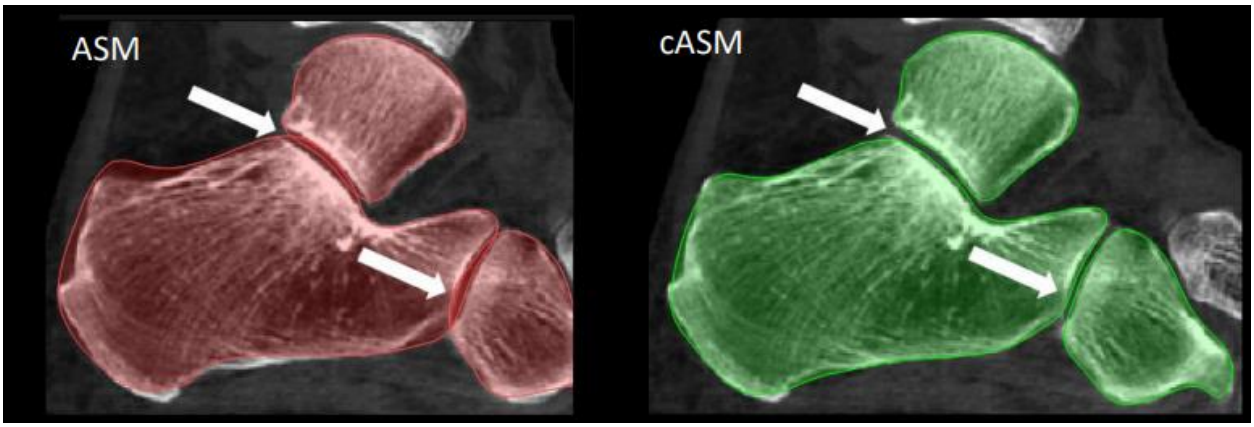

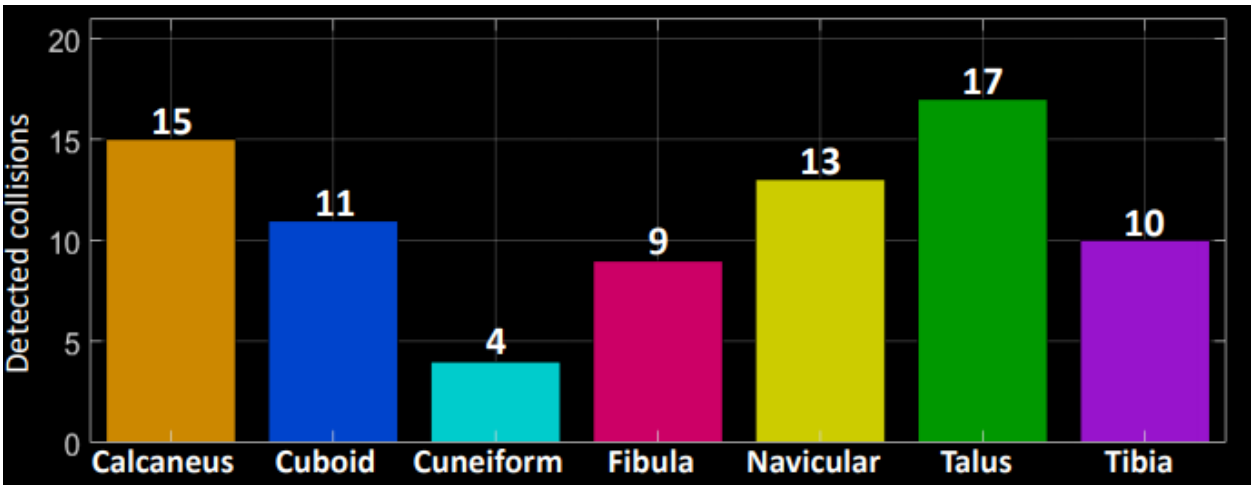

The segmentation accuracy overall improved as well, with an RMSD of 2.7 mm compared to 3.6 mm for traditional cASMs. The cASMs also converged faster on average (~40 vs ~60 iterations), and there was a 25-55% error reduction in the positions of 8 landmark points.

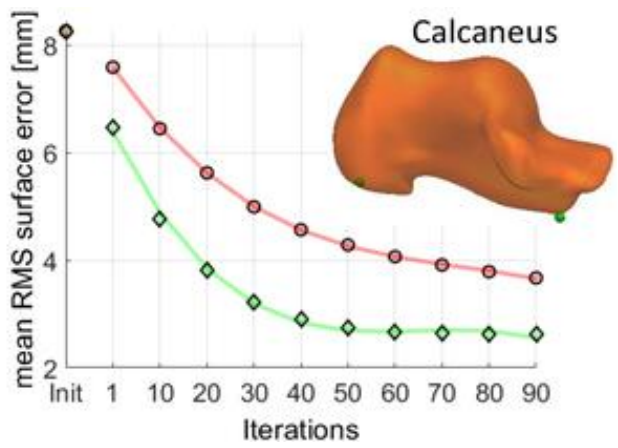

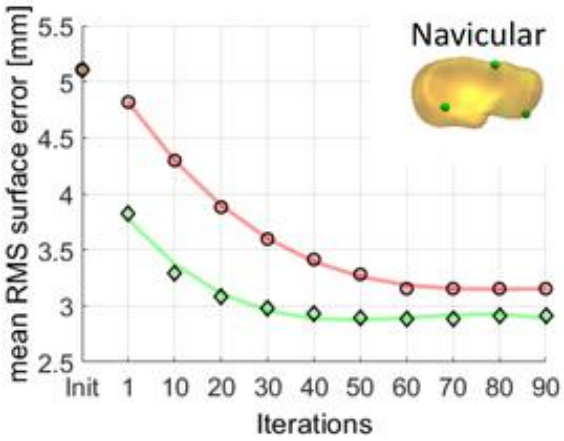

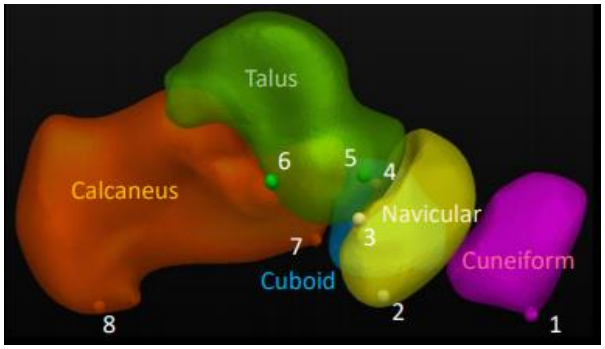

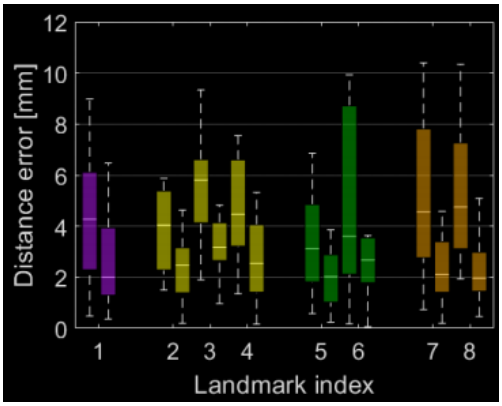

Assessment

Paper 2 demonstrated improved accuracy of segmentations overall with cASMs in comparison to ASMs, especially in narrow spaces between bones and cases of weak gradients. It also showed detection of correction of collisions, faster convergence, and improved landmark accuracy. There are thus many potential clinical applications for this improved accuracy automatic segmentation. However, the sample size of this study was relatively small (21) and could be increased further in later work. In addition, this paper did not show how the framework would perform in the case of artifacts like metal implants or surgical tools in the field of view which may be common during surgical applications and may add noise to the image. Although it may improve performance with poor contrast over regular ASMs, if the image quality overall is too poor and the contrast exceedingly low to the point that that edges can't reliably be detected at all, then it will not work. For these reasons, although surgical applications have been envisioned, they still have to be proven through later studies such as the ankle fracture project and clinical trials by surgeons, both in terms of functionality and ease of use.