The interleaved partial active shape model search (IPASM) algorithm – Preliminary results of a novel approach towards 3D ultrasound-based bone surface reconstruction

Christoph Hänisch[1*], Benjamin Hohlmann[1] and Klaus Radermacher[1]

[1] Chair of Medical Engineering, Helmholtz Institute for Biomedical Engineering,
RWTH Aachen University, Aachen, 52074 Germany
haenisch@hia.rwth-aachen.de

Abstract

In applications such as biomechanical simulations or implant planning, bone surfaces of the knee are most often reconstructed from computed tomography or magnetic resonance imaging data. Here, ultrasound (US) might serve as an alternative imaging modality. However, established methods cannot directly be transferred to US due to differences in imaging quality and underlying physics.

In this paper, we present a generalisation of the well-known active shape model search algorithm (ASM) that allows for segmenting various structures in US volume images that are too large to be captured with a single recording. The multi-view segmentation approach uses a-priori knowledge in the form of a statistical shape model (SSM) as is the case with the classical ASM. This allows to extrapolate missing information and to generate shapes that comply with the underlying distribution of some training data. The main differences are, however, that the SSM is not only adapted to a single image but to multiple images and that the adaption process is interleaved. As a result, within each iteration the surface information of all sub-volumes is propagated and used in all subsequent steps.

In-silico tests were conducted to investigate how this algorithm would perform in real tracked US data. US volume images were split in slightly overlapping sub-volumes, noise was added, and the alignment was distorted. We could show that the algorithm is capable of reconstructing shapes in the lower millimetre range and for some cases even with submillimetric accuracy. The algorithm is hardly affected by orientation errors below 5 degrees and displacement errors below 5 mm; above these limits, the average absolute SDE as well as its associated variance increases.

1 Introduction

In applications such as biomechanical simulations (D'Lima 2013, Asseln 2015) or implant planning (Bellemans 2005), bone surfaces of the knee are most often reconstructed from computed tomography (CT) or magnetic resonance imaging (MRI) data (Al Hares 2015). However, issues related to radiation (CT), additional costs (CT and MRI) and distortions (MRI) (Moro-oka 2007) have to be considered. For that reason, we investigate whether 3D ultrasound (US) may be used as an alternative imaging modality. US provides high resolution, dynamic (optional weight bearing and even soft tissue) imaging, is widely available and cost efficient. Its disadvantages, however, are the low signal-to-noise ratio, the relative small field of view and the acoustic shadowing. Thus, the image information from multiple images has to be combined and missing information has to be extrapolated; existing methods for CT and MRI image processing and reconstruction can often not directly be transferred to US due to differences in imaging quality and underlying physics.

To obtain a big volume image which contains the entire surface information, various sub-volumes might be recorded and registered. Since all volume images need to be at least roughly aligned, the US probe as well as the upper and lower leg must be tracked (e.g. non-invasively by optical markers, see figure 1.1). Unfortunately, the intensity values of various structures may vary severely when slightly changing the probe orientation (cf. Noble 2006). This makes intensity-based methods like mutual information based registration challenging. Bone surfaces might be segmented from the volume images (after enhancing these structures using phase congruency filters, for instance) and then be registered. But then, a lot of overlap is needed to robustly register the resulting patches. Such a procedure may be error prone and very time consuming due to the high amount of volume images to be recorded. Also, missing information must still be complemented. Statistical shape models (SSM) could alleviate this problem and could be used to segment each sub-volume. But in this case, a lot of models would have to be constructed and approximated model guided alignment of the probe by the operator could be cumbersome. In addition, the combination of various SSM which have been independently adapted to different parts of a single geometry is difficult and may result in artefacts.

We hypothesised that the use of a single SSM which is simultaneously adapted to all (roughly aligned) sub-volumes can provide an adequate quality of the 3D surface reconstruction of the relevant areas of the knee bone surface. This paper presents the novel interleaved partial active shape model (IPASM) search algorithm and its evaluation in an in-silico study which aims at investigating the effect of noise and localization errors on the surface reconstruction error of the distal femur.

2 Materials and methods

Construction of the SSM

The statistical shape models were trained from 414 bone surfaces of the distal femur which were semi-automatically segmented from CT data by clinical experts using subdivision surfaces; osteophytes were manually removed. For each SSM, an arbitrary reference shape was selected from the training data and all remaining shapes were aligned with this reference. Then, the reference was morphed into all other shapes which effectively transferred its topology and thus established the sought point correspondences. The morphing was done with the N-ICP-A algorithm introduced by (Amberg 2007). Eventually, the SSM was computed as described in (Heimann 2009) with the difference that a partial Procrustes analysis was performed (thus our models also encode size information). To reduce the bias from the reference shape, the procedure was repeated two times where each time the previously computed mean shape was taken as the new reference shape. The SSM is represented as a triangle mesh with a uniform vertex distance of about 1 mm and has about 22,000 vertex points.

The interleaved partial active shape model search (IPASM) algorithm

The proposed interleaved partial active shape model (IPASM) is a generalization of the active shape model (ASM) as described by (Cootes 1995). The main differences are that the SSM is not only adapted to a single image but to multiple images and that the adaption process is interleaved. As a result, within each iteration the surface information of all sub-volumes is propagated and used in all subsequent steps (cf. figure 1). The algorithm can be summarized as follows:

1. Initialization:

    a. Initialize the SSM such that it represents the mean shape.

    b. Roughly align all volume images with the mean shape (i.e. find all transformations from the image space into the shape space).

2. For each volume image do:

    a. For each vertex point of the SSM:

        i. Search along the vertex normal for a corresponding surface point (e.g. the brightest voxel) in the image data (figure 1.2).

    b. Use all valid correspondences and rigidly align the image with the current instantiation of the SSM (figure 1.3).

    c. Determine the residual errors dx between the vertex points and the corresponding points after the local alignment.

3. Combine all dx_i to a single difference vector d (e.g., a weighted sum which depends on the angle between the surface normal and the probe normal at each vertex point).

4. Compute the difference vector $db = \mathbf{P}d$ in the parameter space where $\mathbf{P}$ is the projection matrix whose rows are the modes of the SSM. Update the parameter vector $b \leftarrow b + db$ and limit b according to its distribution.

5. Reconstruct the shape $x = \bar{x} + \mathbf{P}^T b$ using the mean shape vector $\bar{x}$.

6. Go to item 2 until convergence or until any other stop criterion is met.

The IPASM search is performed on an image pyramid which lets the algorithm converge faster and makes it more robust. If the search length along the vertex normals is kept constant, the range is doubled at each pyramid level.

Evaluation

The aim of the experiments is to investigate the effect of tracking errors (position and orientation) and US image noise on the overall surface reconstruction performance.

   Artificial US volume images were generated from 12 femoral bone surfaces of the distal femur by mapping the signed distance functions of the meshes to grey-scale values. Noise and speckle were added to the volume images (cf. figure 2). In each case, the resulting volume images were split into 8 slightly overlapping pieces. The optimal alignment of each piece was distorted in terms of orientation and location. The resulting rotation errors ranged from zero degrees up to ten degrees and the location errors from zero millimetres up to ten millimetres. In each configuration, a SSM was built up from 413 training samples omitting the current test case (leave-one out cross-validation). The resulting surface reconstructions were rigidly aligned with their ground-truths and the signed surface distance errors (SDE) were computed. The 12 bone surfaces were chosen such that they span the entire range from very small to very big distal femora. In half of all cases, the first two modes were adapted to a

few randomly positioned points along the medial and lateral epicondyles with a method similar to the one proposed by (Fleute 1998). This is done with the assumption that this would remedy the possible aperture problem shown in figure 1.4; the first two modes were chosen since they primarily encode the overall size and the intercondylar width.
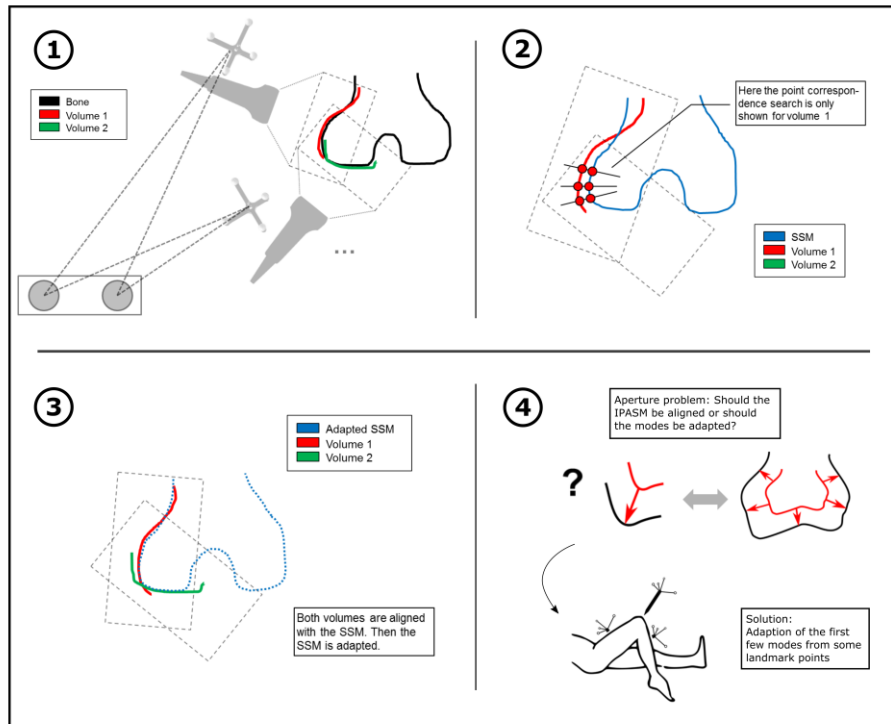

Figure 1: Visualization of the IPASM search algorithm. The principle setup with a tracked probe and tracked bones is shown in sub-figure 1. The correspondence search and the subsequent alignment of all shapes are shown in sub-figures 2 and 3. If only part of the bone is visible it may be undecidable whether a sub-volume needs to be aligned or whether the mode weights have to be adapted (sub-volume 4). Initializing the SSM in the IPASM search algorithm by prior adaption to a few landmark points solves this problem.
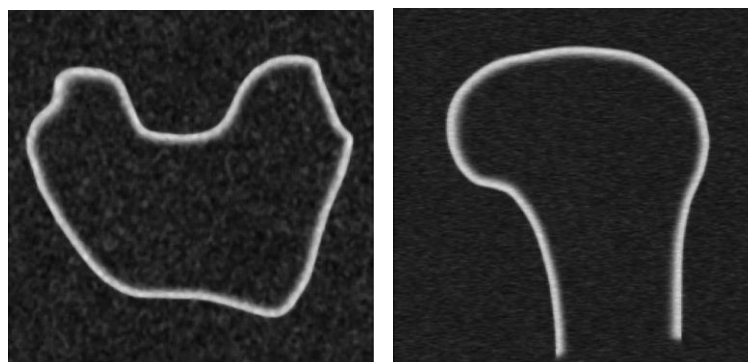

Figure 2: Slices of an artificial US volume image of the distal femur with very slight speckle noise before partitioning it (left: axial plane, right: sagittal plane).

3   Results

When adapting the first two modes of the SSM, the mean over all average absolute surface distance errors (SDE) for the noise-free cases is 0.61 mm ± 0.77 mm (the latter value denotes the standard deviation). The maximum absolute SDE may go up to 7 mm which appears to happen in more isolated areas. An example for a typical reconstruction can be seen in figure 3.

When the modes of the SSM were not adapted at all, the average absolute SDE increased to 0.99 ± 0.95 mm. The high increase in the error was mainly due to problems of the uninitialized IPASM to properly adapt to very small or to very big distal femora; two of the 12 bones had an average absolute SDE of over 2 mm (the remaining bones below 1.2 mm). For that reason, in the following only results of the initialized IPASM are shown.

With an independent rotation error of 5 degrees for each piece, the average absolute SDE remained nearly the same with a slightly increased standard deviation 0.7 ± 0.9 mm. Both increased to
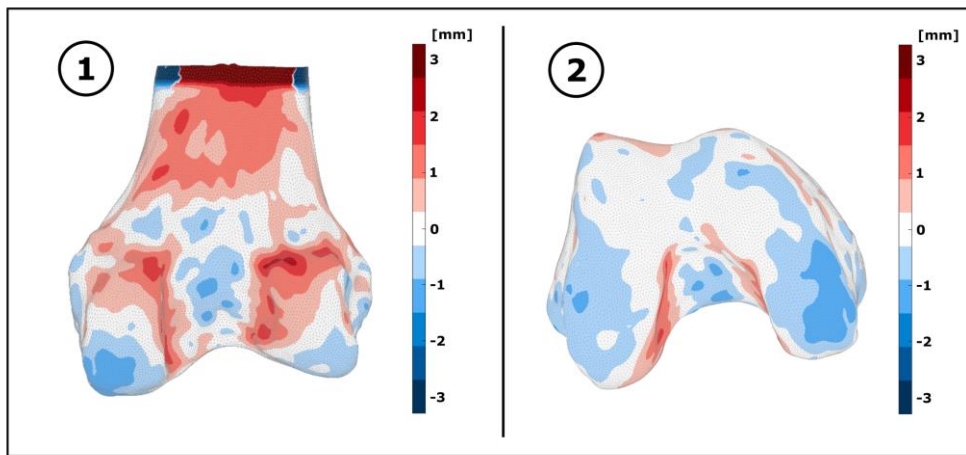

Figure 3: Reconstruction of a distal femoral bone surface with the IPASM. The signed distance error between the reconstruction and the synthetic ground truth is colour-coded.
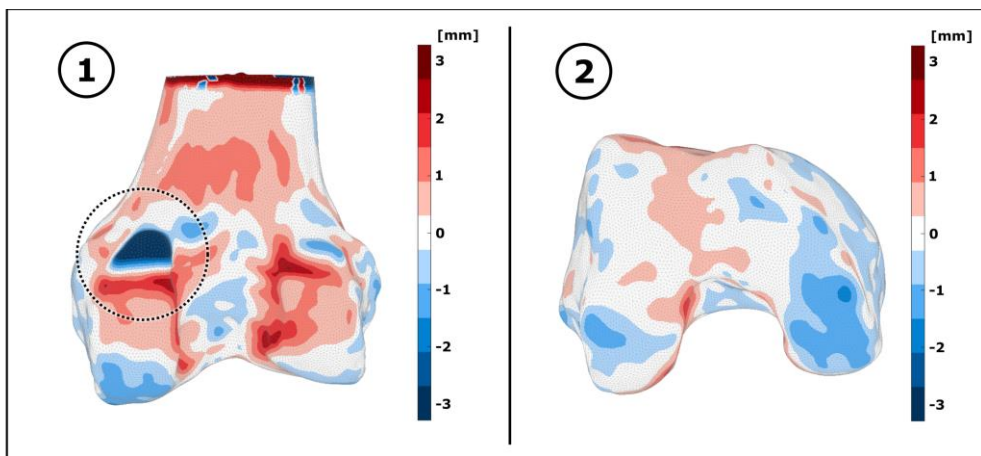

Figure 4: Erroneous reconstruction of a distal femoral bone surface with the IPASM due to wrong point correspondences in overly noisy images.

1.26 ± 1.5 mm, respectively, if the rotation error was increased to 10 degrees. A displacement error of up to 2 mm for each individual segment did not affect the SDE. The average absolute SDE increased to 0.95 ± 1.17 mm and 1.43 ± 1.63 mm for a displacement error of 5 mm and 10 mm, respectively.

In some cases the overall surface geometry was accurately reconstructed apart from a small area which exhibited a bump, effectively leading to a low average absolute SDE and a high maximum SDE (see figure 4). Further investigations suggested that this behaviour was likely due to the artificial high noise level in the volume images. In contrast to real US images, the noise often had brightness values much higher than the actual bone response.

4 Discussion

A new algorithm has been presented that allows for segmenting various structures in US volume images that are too large to be captured with a single recording. The multi-view segmentation approach uses a-priori knowledge in the form of statistical shape models. This allows to extrapolate missing information and to generate shapes that comply with the underlying distribution of some training data; accordingly, training data must be available. The interleaving during the adaption process allows for making full use of the entirely available surface information.

In-silico tests showed that the algorithm is capable of reconstructing shapes in the lower millimetre range and for some cases even with submillimetric accuracy. The algorithm is hardly affected by orientation errors below 5 degrees and displacement errors below 5 mm; above these limits, the average absolute SDE as well as its associated variance increases. In some cases, a very high level of noise may mislead the correspondence search. This, however, is very likely due to flaws in our current implementation which only looks for the brightest voxel along an intensity profile. This weakness could be mitigated, for instance, by filtering the input images using phase congruency filters (Belaid 2011, Hacihaliloglu 2012).

Our statistical shape model is based on a partial Procrustes analysis and thus also encodes size information. For that reason, the IPASM search only uses a rigid transformation (i.e. without scaling) as opposed to the ASM search (Cootes 1995). This allows decoupling the registration process for all sub-volumes (otherwise the scaling factor would have to be balanced between all parts) and thus provides a better compensation of prior registration errors. The drawback is that the algorithm cannot always decide whether it needs to adapt the modes or to align the image with the current model instance. This is often compensated by propagating the surface information from the other parts, but if a shape is too small or too big, this aperture problem needs to be addressed by initializing the first few modes (e.g. with some few points acquired transcutaneously with a pointer device as can be seen in figure 1.4).

(Saha 2011) already described a partial shape model that extrapolates surface information at sites of osteophytes, but this approach does not deal with multi-view segmentation. There is literature like (McDonagh 2013) that deals with the simultaneous registration of multi-view images, but since all these approaches do not make use of any a-priori knowledge (e.g. the shape), they necessarily need some overlapping data. In addition, the obtained surfaces may show deficiencies (no smooth surface) or might be non-physiologic. We think that these issues can be avoided with our approach.

Currently, all vertex points contribute equally to the resulting shape. However, certain parts of the shape may be more relevant than others (e.g. the articulating condyles). Here, weights might be introduced to address this issue.

This study is limited to in-silico data. In an upcoming study, in-vivo tests will be conducted.

References

- Al Hares G, et al., Combined magnetic resonance imaging approach for the assessment of in vivo knee joint kinematics under full weight-bearing conditions, Journal of Engineering in Medicine, 229 (6), 439 – 451, 2015

- Amberg B, Romdhani S, and Vetter T, Optimal Step Nonrigid ICP Algorithms for Surface Registration, Computer Vision and Pattern Recognition, pp. 1 – 8, 2007.

- Asseln M, et al. Prediction of in-vivo knee kinematics: Development of a patient-specific musculoskeletal model of the knee for clinical application, XXV Congress of the International Society of Biomechanics, pp. 471 – 472, 2015.

- Belaid A, et al., Phase-Based Level Set Segmentation of Ultrasound Images, IEEE Transactions on Information Technology in Biomedicine, pp. 138 – 147, 2011

- Bellemanns J, et al., Total Knee Arthroplasty – A Guide to Get Better Performance, Springer, 2005

- Cootes T, Active Shape Models—Their Training and Application, Computer Vision and Image Understanding, 61, pp. 38 – 59, 1995.

- D'Lima D, Patient-Specific Implants and Cutting Guides Better Approximate Natural Kinematics Than Standard Total Knee Arthroplasty, Bone & Joint Journal Orthopaedic Proceedings Supplement, 95, pp. 41 – 41, 2013.

- Fleute M, and Lavallee S, Building a complete surface model from sparse data using statistical shape models: Application to computer assisted knee surgery, Proceedings of MICCAI, 879 – 887, 1998

- Hacihaliloglu I, et al., Automatic Bone Localization and Fracture Detection from Volumetric Ultrasound Images using 3-D Local Phase Features, Ultrasound in Med. And Biology, 38 (1), pp. 128 – 144, 2012

- Heimann T, Statistical shape models for 3D medical image segmentation: A review, Medical image analysis, 13, pp. 543 – 563, 2009

- McDonagh S, Fisher R., Simultaneous registration of multi-view range images with adaptive kernel density estimation, Under review, pp. 31 – 62, 2013

- Moro-oka T, et al., Can Magnetic Resonance Imaging–Derived Bone Models Be Used for Accurate Motion Measurement with Single-Plane Three-Dimensional Shape Registration, Journal of orthopaedic research, 25 (7), pp. 867 – 872, 2007

- Noble A, and Boukerroui D, Ultrasound Image Segmentation – A Survey, Medical Imaging, 25 (8), 2006

- Saha P, et al., A new osteophyte segmentation algorithm using the partial shape model and its applications to rabbit femur anterior cruciate ligament transection via micro-ct imaging, IEEE Transactions on Biomedical Engineering, 58, pp. 2212 – 2227, 2011

Disclosures

Acknowledgements