Hardware-Assisted 2D/3D Intensity-Based Registration for Assessing Patellar Tracking

T.S.Y. Tang¹, N.J. MacIntyre², H.S. Gill³, R.A. Fellows², N.A. Hill²,
D.R. Wilson⁴, and R.E. Ellis¹²

¹ School of Computing, Queen's University, Kingston, Canada
{ttang, ellis}@cs.queensu.ca

² Department of Mechanical Engineering, Queen's University, Kingston, Canada

³ Nuffield Orthopaedic Center, University of Oxford, Oxford, UK

⁴ Department of Orthopaedics, University of British Columbia, Vancouver, Canada

1 Introduction

The problem in 2D/3D intensity-based registration of computed tomography (CT) images is to find a pose such that a digitally reconstructed radiograph (DRR) of the 3D image matches a given 2D image. Generating DRRs, a computationally intensive process, can be accelerated by precomputation [1], or by using custom hardware that accelerates volume rendering. Recent algorithmic advances for 3D texture mapping using off-the-shelf video cards have further improved volume rendering. We propose to generate DRRs generated in real time, without any precomputation, for 2D/3D registration.

This work is a continuation of the point-based fluoroscopic technique for assessing patellar tracking [4], in which a cadaver patellofemoral joint was tracked by a single-plane fluoroscope using a point-based registration method [3] with accuracy determined by comparison to Roentgen Stereophotogrammetric Analysis (RSA). We have developed a hardware-assisted intensity-based registration method and applied it to the same set of data.

2 Materials and Methods

Details of the experimental set-up can be found in [4]. The original CT scan of the knee was roughly divided into three volumes corresponding to each of the femur, patella and tibia, so that only one of the bone would show up in a DRR. To generate a DRR, the CT slices were first transferred to the texture memory of the video card; then, a particular pose could be rendered by setting the camera geometry according to the information from fluoroscope calibration. An NVidia GeForce FX 5800 video card on a 2.4 GHz PC was used. For registration purposes, soft-tissue was not rendered, although this can be done easily by adjusting the transfer function according to the Hounsfield units of the CT voxels.

A *single* fluoroscopic image was used in the intensity-based registration. Each of the femur, patella, and tibia was registered separately. For the similarity measure, which estimated how similar the original fluoroscopic image was to the DRR, we used the weighted gradient correlation (a modified version of gradient

correlation [2]). A rough region of interest was specified in the fluoroscopic image, and when a DRR was generated, pixels in the DRR that were outside the ROI had a negative weight added. To find a registration pose that maximized the similarity measure we used the downhill simplex algorithm in a multiscale manner: given an initial guess, it was perturbed m times and registration was done at a low resolution (128^2), the n best answers were used as initial guesses at a medium resolution (512^2), and its result was used for registering in full resolution (980^2). We used $m = 5$ and $n = 2$.

3 Results and Discussion

With a $512^2 \times 160$ CT, a DRR of size 128^2 , 512^2 , and 980^2 took about 0.01, 0.11, and 0.38 seconds, respectively, to generate. A registration with a single volume took 3 to 5 minutes to complete. Results were compared with the ground truth obtained from RSA. The table shows the mean absolute errors and standard deviations for the variables describing patellar orientation and translation.

	Point-Based		Intensity-Based	
	Error	SD	Error	SD
Patella Motion				
Orientation (degrees)				
Flexion	1.71	0.64	1.81	0.63
Internal Rotation	1.75	0.64	1.80	0.80
Lateral Tilt	0.84	0.55	0.79	0.57
Translation (mm)				
Lateral	0.99	0.81	1.29	0.77
Proximal	0.51	0.35	0.97	0.64
Anterior	2.09	1.35	2.49	1.65

Hardware generation of DRRs did not need precomputation and was about 25 times faster than software computation, yet preserved the image quality. It worked well with intensity-based registration, and it made the implementation straightforward. Also, in less than 5% of all instances, an incorrect registration was found with a similarity value greater than that of the correct registration. This situation could perhaps be resolved by using another similarity measure [2].

Errors in our intensity-based registration were comparable to those reported in literature. The main advantage of our method is that only a single fluoroscopic image is used, while most others use at least two. It allows the assessment of patellar tracking to be done using commonly available graphics hardware.

References

1. D. A. LaRose *et al.* Transgraph: interactive intensity based 2D/3D registration of X-ray and CT data. In *SPIE Image Processing*, 2000.
2. G. P. Penney *et al.* A comparison of similarity measures for use in 2-D-3-D medical image registration. *IEEE Transactions on Medical Imaging*, 17(4):586-595, 1998.
3. T. S. Y. Tang *et al.* Fiducial registration from a single X-Ray image: a new technique for fluoroscopic guidance and radiotherapy. In *MICCAI 2000*, pages 502-511.
4. T. S. Y. Tang and N. J. MacIntyre *et al.* Accuracy of a fluoroscopy technique for assessing patellar tracking. In *MICCAI 2003*, pages I:319-326.