GPU-Accelerated Linear and Deformable 3D Medical Image Registration

D. H. Adler[1,2], S. Chan[2,3], E. S. Penner[2,4], and J. R. Mitchell[2,5]

[1]Electrical and Computer Engineering, University of Calgary, Calgary, Alberta, Canada, [2]Seaman MR Centre, Foothills Med Ctr, Calgary Health Region, Calgary, Alberta, Canada, [3]Computer Science, Stanford University, Palo Alto, California, United States, [4]Computer Science, University of Calgary, Calgary, Alberta, Canada, [5]Clinical Neurological Sciences; and, Radiology, University of Calgary, Calgary, Alberta, Canada

INTRODUCTION

Image registration — the task of aligning or mapping multiple images to bring common features into spatial correspondence — is often required when images acquired at different times, from different subjects, or using different modalities need to be compared or integrated. Automatic image registration is an iterative procedure that maps a moving image onto a fixed image. Each iteration involves (up to) four steps, as shown in Fig. 1: (1) A transform is applied to the moving image; (2) the moving image is resampled into the space of the fixed image; (3) a similarity metric is calculated to measure correspondence in the images; and (4) the metric is passed to an optimization function that modifies the transform parameters to improve dataset alignment.

The purpose of our work is to design and validate methods to leverage recent advances in the programmability, data capacity, and computational power of commodity graphics processing units (GPUs) for the acceleration of automatic medical registration. We implemented affine registration and both control point- and demons-based [1] deformable registration. Our methods were evaluated on inter-modality registrations of clinical 3D magnetic resonance (MR), computed tomography (CT), and positron emission tomography (PET) images. Our new algorithms facilitate the implementation of a clinical registration toolbox for fast assessment of medical images.

METHODS

To date, GPU-based registration has focused on implementing at most two of the iterative steps described above on the GPU. We implement all computationally and memory intensive components of the iterative registration cycle on graphics hardware. Multi-resolution Gaussian pyramids of the fixed and moving images are stored as 3D floating-point textures in GPU video memory. Affine transforms are applied to the moving image by setting their texture transformation matrices. Deformable transforms are implemented via free-form deformation fields stored in 3D textures, and are applied via indirect texture lookup [2]. For control-point deformable registration, a 3D lattice of evenly spaced control points parameterizes the field. A vector at any arbitrary location within the field is computed by linear interpolation between neighbouring lattice points. For demons-based registration, a deformation vector is stored for each moving image voxel. To eliminate costly memory copies, we render directly to the deformation textures, rather than to the screen. The texturing hardware performs all linear interpolation to resample the moving image.

We implemented the following point similarity metrics using custom fragment shader programs: sum-of-squared-differences (SSD), normalized correlation coefficient (NCC), and normalized gradient field (NGF) [3]. We also accelerated the mutual information (MI) metric [4] by computing joint histograms on the GPU [5]. This is done by fetching fixed and moving texture intensity pairs in the vertex shader, then rendering vertices as point primitives to locations that represent histogram bins. Metrics evaluated by the shaders are summed over all slices using additive compositing functions to create a global volume metric. Optimization of the similarity metric is performed using the gradient descent simplex algorithm, with multi-threaded metric functions implemented for control point deformable registration.

We evaluated our methods on synthetically deformed data and clinical cases. Our synthetic test data were generated using MR images from the Montreal Neurological Institute Simulated Brain Database (available at http://www.bic.mni.mcgill.ca/brainweb/). We applied random rigid-body and non-linear deformations to this data, which were then recovered with registration. We quantitatively compared the known deformations to those recovered by our own equivalent GPU and CPU methods, as well as by widely used and publicly available tools, including FLIRT, ITK, and AIR. The clinical cases consisted of aligning PET and CT images to MR images of brain tumor patients from Vanderbilt University's Retrospective Image Registration Evaluation (RIRE) Project database (available at http://insight-journal.org/rire/). Gold standard "truth" transforms have been defined using fiducial markers (and remain sequestered) for all RIRE images. All experiments and timings were performed on an Apple Macintosh PowerMac G5 desktop computer, circa 2006, running OS X 10.4.8 and equipped with Quad 2.5 GHz G5 PowerPC processors, 3 GB of main memory, and an NVIDIA Quadro FX 4500 GPU with 512 MB of video memory.

RESULTS AND CONCLUSIONS

Our first evaluation was the quantification of time performance gains achieved using our GPU methods. Compared to standard CPU implementations of equivalent algorithms, our GPU methods achieved an average performance gain of between 10× to 20× for linear and deformable registration. The relative performance gain for computation of the MI metric itself using our novel implementation is approximately 25× for medical images. Rigid-body registration results for our GPU tool and ITK both return errors below 0.2 mm, with the relative acceleration between our tool and ITK being 13× and 20× for T1 to T1 and T1 to T2 registrations, respectively. Table 1 briefly summarizes our GPU framework's results for the rigid-body alignment of the PET and CT scans to the geometrically rectified MR images in the RIRE database. PET and MR images were aligned using the NCC metric; CT and MR images were aligned using the NGF metric. Voxel spacings for the corresponding PET, MR, and CT images are $2.59 \times 2.59 \times 8.00$ mm, $1.25 \times 1.25 \times 4.00$ mm, and $0.65 \times 0.65 \times 4.00$ mm, respectively. The errors for the RIRE alignments are computed as mean distances between voxels of segmented regions in the fixed and final moving images [6]. Gold standard results give the true distances between the segmentations.

We have developed a novel approach to 3D medical image registration using the GPU. We found no significant difference in the results computed by the registration algorithms when implemented on the GPU and on the CPU. Thus, we conclude that GPU methods yield results equivalent to CPU-based ones, yet require significantly less computation time. Our new methods perform 3D registration of multi-modal clinical data to sub-voxel accuracy within seconds on commodity computers, allowing for fast and efficient registration in clinical image assessment.

Modality	Mean Error (mm)	Time (s)	No. Trials
PET to MR T1	3.82 ± 0.07	3.81 ± 0.77	4
PET to MR T2	4.33 ± 0.06	3.41 ± 0.57	5
CT to MR PD	1.39 ± 0.01	10.19 ± 0.44	7
CT to MR T1	1.42 ± 0.01	11.77 ± 0.89	7
CT to MR T2	1.44 ± 0.03	11.85 ± 1.02	7

Table 1: Summary of results for GPU rigid-body alignments of PET and CT to geometrically rectified MR images from RIRE database.

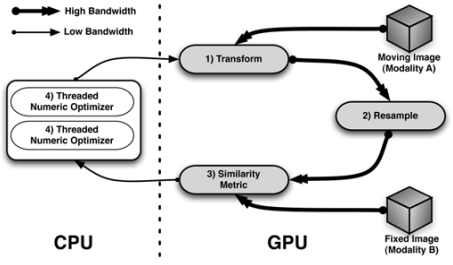

Figure 1: Schematic of the four core components of the iterative registration cycle. Shaded components are executed, and all data is stored on the GPU. Only a similarity metric and transform parameters are transferred between GPU and CPU.

References

[1] Thirion, J.-P. *Medical Image Analysis*, 2(3), pp. 243-260, 1998.
[2] Fernando, F. *GPU Gems*. Boston: Addison-Wesley, 2004.
[3] Haber, E. *et al. MICCAI*, vol. 2, pp. 726-733, 2006.
[4] Viola, P. *et al.. Int. Conf. on Computer Vision*, Boston, MA, pp. 16-23, 1995.
[5] Scheuermann, T. *et al. ACM I3D Proc.*, Seattle, WA, pp. 33-37, 2007.
[6] West, J. *et al. Journal of Computing Assisted Tomography*, 21(4), pp. 552-566, 1997.