1 PUBLISHABLE SUMMARY

Summary description of the project objectives: The main purpose of the project was to apply a novel graphics processing units (GPUs) hardware acceleration framework from the fellow's previous work in real-time Computer Graphics and interactive Visualization to current computational problems in medical image analysis, especially problems arising from imaging moving fetuses in-utero with Magnetic Resonance Imaging (MRI). Making a diagnosis during fetal life permits expectant parents to make informed decisions on choices available in pregnancy. For selected lesions, fetal intervention may also be considered with the aim of improving the natural history of the fetal abnormality. Most importantly, knowing the diagnosis of an abnormality has a direct impact on how the new-born is managed, allowing access to specialist units and appropriate treatments from birth, rather than having a severely ill new-born present without a clear diagnosis or treatment plan. A significant number of infants require long-term care, but targeting follow-on services to children who need them is difficult, particularly as early prediction of neurodevelopmental impairment is currently inaccurate. However, understanding the fetal development is challenging, but crucial for the early detection of congenital abnormalities as well as abnormal development following pre-term birth. Fetal and neonatal MRI has the potential to identify and characterize more subtle aspects of abnormal development. So far, the evaluation of fetal MRI has been a tedious, labour-intensive, and time-consuming endeavour. The reasons for this are motion induced artefacts. Fetuses are usually scanned without sedation and slow image modalities like MRI make inconsistencies between individual slices of a target volume extremely likely. During the past years the fellow developed a fully automatic and hardware-accelerated framework, which allows fast high-resolution reconstruction of motion corrupted data to support clinical examination in three and more dimension. This framework fulfils the initially for the project planned objectives as they are summarized in the following.

Objective 1: Applying recent scheduling methods from the fellow's previous work for GPU-accelerated methods to recent GPU-accelerated medical image processing algorithms to provide a generic and fast scheduled medical image processing framework.

Objective 2: Providing methods for an automatic evaluation of errors and uncertainties of modern medical image processing algorithms concurrently to their execution.

Objective 3: Better utilization of the GPU hardware and reduction of computational bottlenecks to allow novel visualization methods for algorithmic uncertainties and an online parameter correction of the underlying image processing parameters with prioritized user-input.

Objective 4: Enabling steering of fast medical scanning sequences with uncertainty information from interactive post-processing algorithms and reduction of wrong decisions during diagnosis by showing this uncertainty information additionally to the structural data in real-time (especially for prenatal MR imaging sequences).

Work performed since the beginning of the project and main results achieved: The fellow applied flexible hardware acceleration to a novel method for fetal motion correction. Thereby new possibilities arose, which allowed researching the reconstruction and automatic segmentation of the brain and any organ of the fetus and the use of uncertainty defined point-spread-functions with continuous support. The speed of this approach allowed also developing a full body quasi-non-rigid reconstruction of full uterus scans and visualization of the resulting data uncertainties. The newly developed methods are currently integrated into the clinical research practice together with newly developed visualization methods aiming for optimal scan strategies. In the course of the project it turned out that reconstruction speed is the most important issue when developing fetal MRI research applications. Therefore, the fellow focused on performance optimization research and established a novel high-performance fetal organ scanning and reconstruction pipeline. This pipeline consists of the following four steps: (1) fully automatic organ localization in motion corrupted stacks, (2) fast motion correction and super-resolution reconstruction from stacks of 2D slices, (3) visualization of the reconstructed organ and potential data uncertainties, and (4) feedback to the scanner to acquire missing data. The results about the individual steps have been published in the leading forums for medical image analysis.

Organ localization in (1) has been investigated first and achieved through the use of rotation and motion

robust feature descriptors. Automatic detection of the fetal brain in Magnetic Resonance (MR) Images is especially difficult due to arbitrary orientation of the fetus and possible movements during the scan. In this paper, we propose a method to facilitate fully automatic brain voxel classification by means of rotation invariant volume descriptors. To test step (1) of the pipeline, we calculated features for a set of 50 prenatal fast spin echo T2 volumes of the uterus and learn the appearance of the fetal brain in the feature space. We evaluate our novel classification method and show that we can localize the fetal brain with an accuracy of 100% and classify fetal brain voxels with an accuracy above 97%. Furthermore, we show how the classification process can be used for a direct segmentation of the brain by simple refinement methods within the raw MR scan data leading to a final segmentation with a Dice score above 0.90.

Furthermore, a semi-automatic method for analysis of the fetal thorax in genuine three-dimensional volumes has been developed using these techniques. After one initial click we localize the spine and accurately determine the volume of the fetal lung from high resolution volumetric images reconstructed from motion corrupted prenatal Magnetic Resonance Imaging (MRI). We compare the current state-of-the-art method of segmenting the lung in a slice-by-slice manner with the most recent multi-scan reconstruction methods. We use fast rotation invariant spherical harmonics image descriptors with Classification Forest ensemble learning methods to extract the spinal cord and show an efficient way to generate a segmentation prior for the fetal lung from this information for two different MRI field strengths. The spinal cord can be segmented with a DICE coefficient of 0.89 and the automatic lung segmentation has been evaluated with a DICE coefficient of 0.87. We evaluated this method on 29 fetuses with a gestational age (GA) between 20 and 38 weeks and show that our computed segmentations and the manual ground truth correlate well with the recorded values in literature.

We also investigated methods for the automatic evaluation of fetal motion from specialised MRI sequences. Being able to automate the location of individual fetal body parts has the potential to dramatically reduce the work required to analyse time resolved fetal Magnetic Resonance Imaging (cine-MRI) scans, for example, for use in the automatic evaluation of the fetal development. Currently, manual preprocessing of every scan is required to locate body parts before analysis can be performed, leading to a significant time overhead. With the volume of scans becoming available set to increase as cine-MRI scans become more prevalent in clinical practice, this stage of manual preprocessing is a bottleneck, limiting the data available for further analysis. Any tools which can automate this process will therefore save many hours of research time and increase the rate of new discoveries in what is a key area in understanding early human development. Here we present a series of techniques which can be applied to fetal cine-MRI scans in order to first locate and then differentiate between individual body parts. A novel approach to maternal movement suppression and segmentation using Fourier transforms is put forward as a preprocessing step, allowing for easy extraction of short movements of individual fetal body parts via the clustering of optical ow vector fields. These body part movements are compared to a labelled database and probabilistically classified before being spatially and temporally combined to give a final estimate for the location of each body part.

Pipeline steps (2) and (3) have been investigated using novel flexible methods for the evaluation of super-resolution reconstruction on the GPU. Capturing an enclosing volume of moving subjects and moving organs using fast individual image slice acquisition turned out to provide a promising way to be robust to motion artefacts in 2D. Motion between slice acquisitions results in spatial inconsistencies, which can be resolved by slice-to-volume reconstruction (SVR) methods proving high quality 3D image data. However, existing algorithms have been typically very slow, specialised to a specific application, and use inaccurate assumptions thus impeding the potential clinical use of these methods.

At these steps of the pipeline we presented a fast generalized multi-GPU accelerated framework for slice-to-volume reconstruction. It is based on optimised 2D/3D registration and super-resolution with automatic outlier rejection and additional optional intensity bias correction. We introduced a novel and fully automatic selection procedure for the image stack with the least motion, which will serve as an initial registration target. We fully evaluate the presented method using artificially motion corrupted phantom data and clinical data including tracked freehand ultrasound of the liver and fetal Magnetic Resonance Imaging. We achieved speed-ups of more than -30× compared to a single CPU system and more than 10× compared to currently available state-of-the-art multi-core CPU methods.

High reconstruction accuracy is ensured by computing a fully continuous point-spread function for every input data point, which has not been possible before due to computational bottlenecks. Our framework and its implementation is scalable to the available computational infrastructure and tests show a speed-up of 1.70× for each additional GPU. This paves the way for online application of image based reconstruction methods as part of clinical examinations. An example for motion corrupted data and its reconstruction is shown in Figure 1.

The reconstruction framework allows to compute the quality of the reconstruction at any point in the volume and advanced visualization methods to communicate this additional information are investigated in the course of a by the fellow supervised student project which will go on beyond the end of this project.

To achieve the last step (4) of the developed pipeline, several novel methods to account for fetal movements during fetal Magnetic Resonance Imaging (fetal MRI) have been explored. We showed how slice-to-volume reconstruction methods can be used to account for motion adaptively during the scan. Three candidate methods have been tested for their feasibility and integrated into a computer simulation of fetal MRI. The first alters the main orientation of the stacks used for reconstruction, the second stops if too much motion occurs during slice acquisition and the third steers the orientation of each slice individually. Reconstruction informed adaptive scanning can provide a peak signal-to-noise ratio (PSNR) improvement of up to 2 dB after only two stacks of scanned slices and is more efficient with respect to the uncertainty of the final reconstruction.

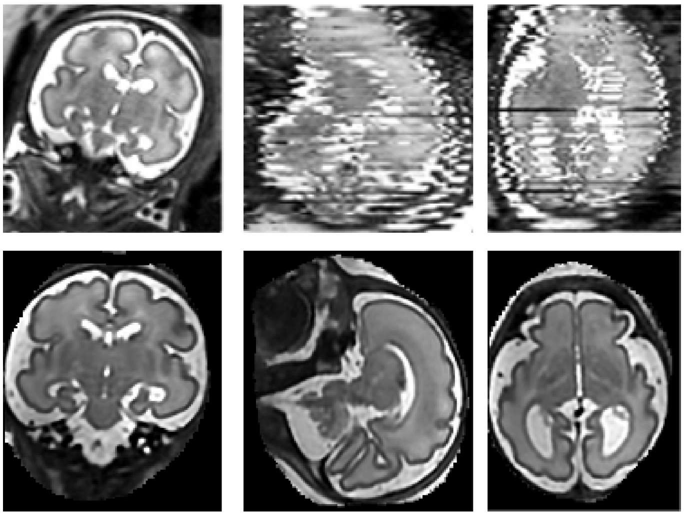

Figure 1: Top row: An example of three orthogonal views through a stack of motion corrupted MRI slices of the fetal brain. Note the significant motion artefacts between the slices and the intensity bias. The left image shows an acquired slice and the other two images orthogonal planes through a stack of these slices. Bottom row: The resulting reconstruction at 0.75mm isotropic voxel size after applying my methods that the fellow developed during the project.

The source code for the developed approach, which supports parallel computing on multiple GPUs is publicly available on github (https://github.com/bkainz/fetalReconstruction.git). It has been published in the most important journal for medical image analysis (IEEE Transactions on Medical Imaging, doi:10.1109/TMI.2015.2415453) and has been used for advanced analysis of fetal data (for example published in the most relevant conference for medical image analysis: MICCAI, doi:10.1007/978-3-319-10470-6_36 or one of the most recognised journals for brain imaging: NeuroImage, doi:10.1016/j.neuroimage.2014.07.023). Furthermore, the fellow maintained a blog as project website, which can be reached via http://mc-faust.blogspot.co.uk/.

The expected final results and their potential impact and use: The fellow will continue to integrate the developed methods into the clinical research practice at King's College London and at various sites all over the world. Their immediate impact is to provide so far unseen and previously impossible structural data of growing fetuses. This project builds the foundation for various other projects researching improvements of standard ultrasound examination methods and supports the research of the developing human connectom.