Bringing Compressed Sensing to Clinical Reality: Prototypic Setup for Evaluation in Routine Applications

Kai Tobias Block[1], Robert Grimm[2], Li Feng[1], Ricardo Otazo[1], Hersh Chandarana[1], Mary Bruno[1], Christian Geppert[3], and Daniel K. Sodickson[1]
[1]*Center for Biomedical Imaging, NYU Langone Medical Center, New York, NY, United States,* [2]*Pattern Recognition Lab, University of Erlangen-Nuremberg, Erlangen, Germany,* [3]*Siemens Medical Solutions, New York, NY, United States*

Introduction: Compressed sensing (CS) is a rather novel concept for reconstructing images from incomplete data[1], which has potential to overcome some of the major limitations of today's MR imaging techniques. However, because it requires solving a non-linear numerical problem with an iterative optimization technique, the computations are very expensive and reconstruction times exceed those of established MR techniques by magnitudes. While CS has raised considerable enthusiasm in the scientific community and initial results from research applications have been published, CS is not available for real-world clinical applications yet. Hence, at this point it is not clear if CS is robust enough for routine patient exams and if it delivers consistent diagnostic image quality in daily use. Manufacturers are currently hesitant to implement CS approaches commercially because the high computational demand requires redesign of the image-reconstruction architectures. Consequently, it is necessary to validate the clinical benefit of CS on a scientific level before manufacturers will make such substantial investment into the technology.
Recently, we developed a CS-based approach for motion-robust dynamic T1-weighted imaging entitled GRASP[2], which provided promising initial results for patients that are unable to suspend respiration. In this work, we describe how we integrated the method into our standard clinical workflow, allowing us to evaluate its clinical performance in a large number of patient exams and to obtain direct feedback from the radiologists reading the conventional exams of these patients.

Methods: The GRASP method is based on acquisitions with a 3D radial stack-of-stars GRE sequence that uses the golden-angle ordering scheme[3]. This sequence was implemented using the IDEA development environment (Siemens AG, Healthcare Sector, Germany), which enables us to use the sequence across all our scanners including 1.5 Tesla and 3 Tesla systems. When performing dynamic scans with the sequence, data is acquired continuously over a time period of up to 6 min while the contrast agent is injected 20 sec after the start of the scan with a power injector. Because the user interface of the sequence is identical to that of product sequences, all of our technicians are able to run these scans without further support from research staff. Time-averaged gridding images are reconstructed on the MRI system so that the technicians can verify the quality of the exam before the patient leaves the table.

Data Transport: Because dynamic GRASP reconstructions would block the MR system for long time and would affect patient throughput, the computations are performed on an external server. To transport the data to our dedicated GRASP server, a raw-data storage software has been developed, which runs as background service on all our clinical MR systems. Once per night, the software automatically parses all scans performed during the day for protocols that have been configured for offline reconstruction and sends the raw data to the server through a network share. If data for a specific patient is needed more urgently, the transfer can also be initiated right after the exam with a single mouse click, so that the data is transferred while the patient is taken off the table. Noteworthy, because the software runs on the console computer and employs a standard mechanism for data export, operating the software does not require any configuration change that would affect the MR system's certification for clinical use (as opposed to many other solutions that run on the reconstruction computer). In this way, no manual steps are required to handle raw data, which proved to be highly important for systems with busy patient scheduling. Furthermore, it helps us monitoring for which patients dynamic GRASP scans were performed.

Image Reconstruction: When raw data arrives on the reconstruction server, the GRASP reconstruction is started using predefined reconstruction parameters for different clinical applications. To achieve sufficiently high reconstruction speed, the GRASP algorithm was implemented as standalone application on the Linux operating system using the C++ programming language. As initial processing step, a SVD-based channel compression is performed to reduce the amount of raw data[4], which can take up to 6 GB per scan with the protocols we use. Afterwards, the reconstruction problem is separated into individual slices by performing a FFT along the partition direction of the stack-of-stars geometry. Slices are then processed in parallel where the interpolation coefficients for the gridding operation are pre-calculated and shared across all threads. Parallelization was achieved using the OpenMP interface[5] where memory allocations were supported by the libnuma library[6] to ensure that working memory is bound to the computing nodes on multi-core systems with non-uniform memory access architecture (NUMA). To solve the reconstruction problem, a C implementation of the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) algorithm[7] was used. After the iterations have been finished, the reconstructed images are saved as DICOM files using the libraries provided by the DCMTK package[8]. Image orientation and relevant patient information are extracted from the header of the raw-data file and written into the corresponding DICOM tags. When all images are saved, the DICOM files are sent to our PACS archive using the storescu tool from DCMTK. In this way, our radiologists can read and evaluate the images together with the conventional exams using standard PACS workstation software. Currently, we are running the reconstructions on a server equipped with 4 AMD Opteron 6272 16-core CPUs and 128 GB memory. Depending on the protocol settings for different clinical applications, the reconstruction of a full 3D dynamic data set with high temporal resolution usually takes between 5 and 45 min on this system.

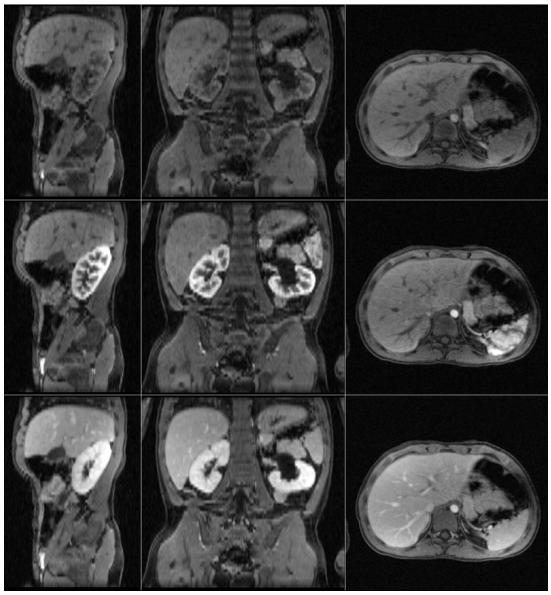

Results: Dynamic imaging using the GRASP method is currently evaluated at our institution for a variety of clinical applications. However, because GRASP is not yet an established and approved imaging technique, usage of GRASP is confined to applications where no interference with standard clinical imaging occurs. Thus, if the clinically required protocols involve dynamic imaging, we cannot use GRASP for these patients whereas if only pre- and post-contrast images are required, we can acquire GRASP scans during the injection of the contrast agent. This helps exploring the possibilities of GRASP and identifying potential limitations in an upmost realistic scenario. Dynamic studies are currently acquired for pediatric body imaging, breast imaging, neck imaging, imaging of the orbits, spine imaging, lung imaging, and enterography. At time of writing, a total of 343 dynamic GRASP patient studies were conducted.
As one example, the figure shows selected MPR slices from an axial dynamic body scan of a 3-year old patient after partial nephrectomy. Imaging had to be done in free breathing as the patient was sedated. Due to the high motion robustness of the GRASP approach, diagnostic image volumes are obtained with a temporal footprint of 7.2 sec for different dynamic phases, including the pre-contrast, arterial, and late phase as shown.

Conclusion: Compressed sensing has high potential to impact the way MR imaging is performed in the future, but its feasibility and reliability for routine clinical applications remain to be validated. With the prototypic setup that we described in this work, we are now able to address this question in large patient studies for the case of CS-supported dynamic imaging. Multiple clinical studies are currently conducted and will be published in a series of manuscripts with our clinical collaborators. In the next step, the setup will be made available to researchers at external sites for multi-center studies.

References: [1] Lustig et al, MRM 58, 2007 [2] Chandarana et al, Invest Radiol, in press, 2013 [3] Chandarana et al, Invest Radiol 46, 2011 [4] Buehrer et al, MRM 57, 2007
[5] http://www.openmp.org [6] http://oss.sgi.com/projects/libnuma [7] Nocedal et al, Math of Comp 35, 1980 [8] http://dicom.offis.de