Published in final edited form as:

Magn Reson Med. 2017 September ; 78(3): 897–908. doi:10.1002/mrm.26467.

Hybrid MRI-ultrasound acquisitions, and scannerless real-time imaging

Frank Preiswerk¹, Matthew Toews², Cheng-Chieh Cheng¹, Jr-yuan George Chiou¹, Chang-Sheng Mei³, Lena F. Schaefer¹, W Scott Hoge¹, Benjamin M Schwartz⁴, Lawrence P Panych¹, and Bruno Madore¹

¹Department of Radiology, Brigham and Women's Hospital, Harvard Medical School, Boston, MA, USA

² The Laboratory for Imagery, Vision and Artificial Intelligence, École de Technologie Supérieure, Montréal, QC, Canada

³Department of Physics, Soochow University, Taipei, Taiwan

⁴Google Inc, New York, NY, USA

Abstract

Purpose—To combine MRI, ultrasound and computer science methodologies toward generating MRI contrast at the high frame rates of ultrasound, inside and even outside the MRI bore.

Methods—A small transducer, held onto the abdomen with an adhesive bandage, collected ultrasound signals during MRI. Based on these ultrasound signals and their correlations with MRI, a machine-learning algorithm created synthetic MR images at frame rates up to 100 per second. In one particular implementation volunteers were taken out of the MRI bore with ultrasound sensor still in place, and MR images were generated on the basis of ultrasound signal and learned correlations alone in a ‘scannerless’ manner.

Results—Hybrid ultrasound-MRI data were acquired in eight separate imaging sessions. Locations of liver features, in synthetic images, were compared to those from acquired images: the mean error was 1.0 pixel (2.1mm), with best case 0.4 and worst case 4.1 pixels (in the presence of heavy coughing). For results from outside the bore, qualitative validation involved optically tracked ultrasound imaging with/without coughing.

Conclusion—The proposed setup can generate an accurate stream of high-speed MR images, up to 100 frames per second, inside or even outside the MR bore.

Keywords

Hybrid imaging; MR-ultrasound imaging; motion tracking; machine learning; image-guided therapy

Introduction

A main goal of the present work was to acquire ultrasound (US) and MR signals essentially at the same time, and to have a learning algorithm discover correlations between the two. Based on these correlations, US signals became a predictor or surrogate for MRI, in the sense that synthetic MR images could be generated from them.

While MRI has proven useful for the real-time guidance of clinical procedures (1, 2), the imaging process is typically too slow to properly capture breathing motion, especially in the presence of coughing or gasping. The so-called ‘in-bore’ application presented here involved creating synthetic MR images in-between acquired ones, to boost temporal resolution by up to two orders of magnitude. This high-rate stream of synthetic MR images for example could facilitate lesion tracking in the presence of breathing motion, for ablation purposes. Alternately, the so-called ‘out-of-bore’ application presented here involved moving volunteers out of the scanner room while pursuing synthetic MRI based on US signals and learned correlations alone. These out-of-bore synthetic images might help guide therapies that could not be performed within the confines of an MRI scanner, or toward registering images subsequently acquired from different modalities and scanners, for example.

The present work employed simple and relatively low-cost US hardware, as in (3–7). The single-element transducer was small enough to easily fit below or within the openings of a multi-element MR receiver coil, it did not need to be located or tilted in any specific way, and it was fixed to the subject’s abdomen using a simple adhesive bandage (Fig. 1). This contrasts with other existing setups that combine US and MR acquisitions (8–12), based on full-size imaging transducers, typically handheld or affixed to a holder. An especially notable example is the MR-compatible US scanner developed at the Fraunhofer Institute for Biomedical Engineering IBMT and used for breast imaging (e.g., link from (13)). A main realization at the basis of the present work is that although a small transducer is insufficient to generate 2D or 3D spatially-resolved images, it may not need to because MRI is spatially resolved and correlations exist between the two signal types. Compared to our prior work in (7), the present paper presents a new application (out-of-bore scanner-less imaging), a validation strategy, an algorithm for cough detection, an improved Bayesian learning algorithm and, of course, a completely different and larger set of human subjects.

US signals are employed here as a motion sensor. There are other motion sensors that have been used with considerable success in MRI, such as respiratory bellows and navigator echoes. Two bellows, one on the thorax and one on the abdomen, may be capable of properly capturing respiratory motion (14). However, two or more bellows would involve much fabric, velcro and cables, all liable to shift, de-adjust and/or contaminate sterile areas in an image-guided therapy context. In contrast, the US probe employed here contacts only a small area of a subject’s abdomen, out of the way of any interventionalist and any potential sterile area. Navigator echoes represent another alternative for monitoring motion, for example in 4D MRI (15, 16). The main advantages of employing US signals instead of navigator echoes are as follows: US signal acquisitions happen in parallel and simultaneously with the MRI scan and do not reduce the time available for MR image data acquisition, unlike most navigated schemes. Furthermore, US signals are available outside as

well as inside the MR bore, leading to the intriguing possibility of synthesizing real-time MR images from patients who are not even in the scanner.

The proposed method was tested in twenty-two real-time acquisitions, performed over eight separate imaging sessions, inside and/or outside the bore. Subjects were instructed to occasionally cough or gasp to further challenge the algorithm. Validation inside the bore involved comparing synthetic US-based MR-like images to acquired ones, while outside the bore qualitative validation involved time-matched and optically-tracked US imaging (USI) data.

Methods

Organ-Configuration Motion (OCM) sensor

A simple MR-compatible US transducer, referred to as an Organ-Configuration Motion (OCM) sensor, was used here to characterize motion (Fig. 1). Its field was not focused; ideally, it would penetrate and reflect multiple times within the abdomen for the received signal to act as a unique signature of the arrangement of internal organs at any given moment. In contrast, conventional USI may involve transducers with hundreds of elements, and images reconstructed with a ‘delay-and-sum’ (D&S) beamforming algorithm (17, 18) or some related alternative (19–22). A D&S reconstruction rejects much of the raw data it operates on, as it strives to only preserve signals from US waves that traveled at 1540 m/s and reflected only once. Even though the hardware used here is simply a single-element transducer, we aim to exploit its signals as fully as possible without rejecting any motion-related information.

In-bore and out-of-bore applications

The proposed work involves acquiring a time series of MRI images, I_t , along with a simultaneous time series of ultrasound OCM sensor signals, U_t . A continuously-learning algorithm finds correlations between the two, so that the OCM signals become a surrogate for MRI images. Synthetic OCM-based images were exploited here in two different ways, referred to as in-bore and out-of-bore applications. The in-bore application involved generating images at the rate of the OCM acquisition, to visualize respiratory organ motion with greatly improved temporal resolutions. Possible applications might include tracking moving lesions during tumor ablation. The out-of-bore application involved stopping the I_t stream while pursuing with the U_t stream. Using learned correlations and the ongoing U_t stream, synthetic MRI results were generated from volunteers even after they had been removed from the MRI suite. Possible applications might include registering multi-modality datasets using the synthetic MRI results as a common thread between all successive scans and acquisitions. The learning algorithm is described below, first for a single imaging plane and then for the more general multi-plane case.

Learning algorithm – Single plane

Let T represent the collection of time points when MR images were actually acquired and $D = \{I_T, U_T\}$ the collection of all N_T available MR images and associated OCM data. Each image I_T was obtained from a set of individual MR signals or k -space lines, acquired over a

number of TR intervals, and similarly, each U_T may also consist of a set of several individual OCM signals (see timing diagram, Fig 2). Our method seeks to estimate a new MR image, I_b , from the most current set of OCM data available, U_b , based on past experience, D . More specifically the expectation of I_t is computed as:

$$E[I_t|U_t, D] = \int I_t p(I_t|U_t, D) dI_t = \frac{\int I_t p(I_t, U_t|D) dI_t}{p(U_t|D)}. \quad (1)$$

The second equality results from applying Bayes' rule, where $p(I_t, U_t|D)$ is the joint density of the MR image I_t and observed OCM data U_t , conditioned on previously seen data D . We propose below an instance-based method for computing Eq. 1. The joint density in the numerator is estimated using a Kernel Density Estimation (KDE) (23) of the form:

$$p(a, b) \approx \frac{1}{N} \sum_{i=1}^N k_a(a - a_i) k_b(b - b_i), \quad (2)$$

where $k(\cdot)$ represents the KDE kernel. A critical parameter in KDE is the selection of a suitable kernel bandwidth, to determine the spread of the kernel weights such that

$$\lim_{\sigma \rightarrow \infty} k(a, b; \sigma) = 1, \quad (3)$$

$$\lim_{\sigma \rightarrow 0} k(a, b; \sigma) = \delta(a - b), \quad (4)$$

where $\delta(\cdot)$ is a Dirac delta function. Using the training set D as defined above, a Gaussian model is chosen here for the KDE kernel $k(\cdot)$:

$$k(U_t, U_T) = \nu \cdot \exp \left(-\frac{1}{2} (U_t - U_T)^T \Sigma^{-1} (U_t - U_T) \right) := N(U_t; U_T, \Sigma), \quad (5)$$

where $\nu = \left(\sqrt{(2\pi)^n \|\Sigma\|} \right)^{-1}$, Σ is a covariance matrix and the operator T represents a transpose. For one of the two kernels, k_a in Eq. 2, the bandwidth is infinitesimally small resulting in the Dirac delta function $\delta(I_t - I_T)$ centered at I_t while for k_b we use a Gaussian kernel with covariance matrix Σ . The numerator in Eq. 1 becomes:

$$\int I_t p(I_t, U_t|D) dI_t \approx \frac{1}{N_T} \int I_t \sum_i \delta(I_t - I_i) N(U_t; U_i, \Sigma) dI_t = \frac{1}{N_T} \sum_i I_i N(U_t; U_i, \Sigma). \quad (6)$$

where i loops over entries in T . The denominator in Eq. 1 is obtained from:

$$p(U_t|D) \approx \frac{1}{N_T} \sum_i N(U_t; U_i, \Sigma). \quad (7)$$

Combining Eq. 6 and Eq. 7, the computational form of the expectation of I_t becomes:

$$E[I_t|U_t, D] \approx \frac{\sum_i I_i N(U_t; U_i, \Sigma)}{\sum_i N(U_t; U_i, \Sigma)}, \quad (8)$$

where Eq. 8 is of a form consistent with a Nadaraya-Watson Kernel Regression (24).

Individual entries in I_T were obtained from k-space data acquired over a period of $TR \times N_{TR} = TR \times N_y / (N_{echoes} \times R)$, with N_y the image matrix size, N_{echoes} the number of k-space lines per TR period and R the acceleration factor. Similarly, U_T consisted of N_{OCM} individual concatenated OCM traces acquired over a period of $TR \times N_{OCM}$, at a rate of one OCM trace per TR. One such OCM trace was time-matched with the acquisition of k-space center, at $k_y = 0$, and the others stretched back in time up to a moment $(t - TR \times (N_{OCM} - 1))$ in the past. Accordingly, Eq. 8 is causal as it involves only data acquired at or prior to the current time t , and the computed frame I_t emulates an MR image whose k-space center would have been acquired at t . Because different motion types lead to different OCM signal evolutions, and that U_T captures a time window of width $TR \times N_{OCM}$, Eq. 8 can intrinsically differentiate between inspiration and expiration periods. OCM signals were minimally processed: Envelope detection was performed as $|H(U_t)|$, where $H(\cdot)$ represents the Hilbert transform (25) and $|\cdot|$ the magnitude operator. As is usual for ultrasound signals, a logarithm operator was applied to help handle the wide variations in dynamic range.

Learning algorithm – Extension to multiple planes

When acquiring intersecting planes, as done here, one could apply Eq. 8 to each plane independently. However, the motion of different planes is highly correlated especially if/where they intersect. With I_t and J_t the MR images at a first and second plane, Eq. 1 is adapted to compute both planes jointly:

$$E\{[I_t, J_t]|U_t, D\} = \int \int [I_t, J_t] p(I_t, J_t|U_t, D) dI_t dJ_t, \quad (9)$$

where $[I_t, J_t]$ is a concatenation of the two images.

Let $[I_t]_L$ and $[J_t]_L$ represent all pixels at the intersection between the two planes, and $[I_t]_{\bar{L}}$ and $[J_t]_{\bar{L}}$ all other locations. The joint distribution $p(I_t, J_t|U_t, D)$ is separated in terms of dependent regions (where the planes intersect) and independent regions (where they do not):

$$p(I_t, J_t|U_t, D) = p([I_t]_{\bar{L}}|U_t, D) p([J_t]_{\bar{L}}|U_t, D) p([I_t]_L, [J_t]_L|U_t, D). \quad (10)$$

We represent $p([I]_L, [J]_L | U_t, D)$ with a Gaussian model $\mathcal{N}([I]_L; [J]_L, \Sigma_L)$, with Σ_L for covariance and $[I]_L$ (or $[J]_L$) for mean. Applying KDE as above, the two-plane equivalent of Eq. 8 becomes:

$$E\{[I_t, J_t] | U_t, D\} \approx \frac{\sum_i \sum_j [I_i, J_j] N(U_t; U_i, \Sigma) N(U_t; U_j, \Sigma) N([I_t]_L; [J_t]_L, \Sigma_L)}{\sum_i \sum_j N(U_t; U_i, \Sigma) N(U_t; U_j, \Sigma) N([I_t]_L; [J_t]_L, \Sigma_L)}. \quad (11)$$

where U_i and U_j are OCM signals associated with images I_i and J_j respectively. Σ_L was evaluated directly from the OCM data by calculating the standard deviation at all points along the OCM trace, in the first five seconds of OCM acquisition, and assuming these points to be independent in terms of noise. Generalization of Eq. 11 to an arbitrary number of planes parallel to I and J involves repeated application of the product rule in Eq. 10.

Reconstruction of out-of-bore results

Out-of-bore results, just like in-bore results, were reconstructed using Eq. 11; one difference however, the past experience D extended all the way to the current time point t for the in-bore case but stopped sometime in the past for the out-of-bore case. Because D was fixed in time and continuous learning stopped, out-of-bore results were susceptible to unexpected changes in OCM signals, such as those associated with any displacement of the OCM sensor. To alleviate this problem, a non-rigid registration step was added, to make OCM signals inside and outside the scanner agree better prior to reconstruction (26). More specifically, a transformation T was introduced that minimized the cost function:

$$C = -C_{\text{similarity}}(\bar{U}_{\text{in}}, T(\bar{U}_{\text{out}})) + \lambda C_{\text{smooth}}(T), \quad (12)$$

where $\bar{U}_{\text{in}}$ and $\bar{U}_{\text{out}}$ are averages of 100 OCM exhalation traces, manually selected within a couple of seconds from inside the scanner, and from a short acquisition outside the scanner. The sum of squared differences was used for similarity,

$$C_{\text{similarity}}(A_t(x), B_t(x)) = \sum_{x=1}^{N_x} (A_t(x) - B_t(x))^2 \text{ and the smoothness of the deformation}$$

was constrained according to $C_{\text{smooth}}(T) = \frac{1}{N_x} \int_{x=1}^{N_x} \left(\frac{d^2 T}{dx^2} \right)^2 dx$, where N_x is the number of samples per OCM trace. A one-dimensional 15 B-spline grid point based deformation model was used at 6 resolution levels with $\lambda = 1.0$ to efficiently solve Eq. 12 using the Simplex algorithm (27).

Cough/gasp detector

Coughing and gasping may cause rapid motion that Eq. 11 cannot readily handle. A derivative-based statistical algorithm capable of detecting such instances of rapid motion was developed; to be helpful, the algorithm must robustly reject problematic time frames and yet accommodate physiological variations. The scalar quantity v_t defined below, based on OCM signals U_b , proved sensitive to motion-induced changes (see Fig. 3):

$$v_t := \left| \frac{\partial U_t}{\partial t} \right|_1 \approx \left| \frac{U_t - U_{t-1}}{\text{TR}} \right|_1, \quad (13)$$

where $|\cdot|_1$ is an l1-norm operator. While v_t tends to show clear peaks in the presence of rapid motion, it is also sensitive to noise and physiological signals, e.g., heartbeats. A more sophisticated classifier ℓ was constructed as follows: A Gaussian model $\mathcal{N}(v_b; \mu, \sigma)$ was fitted to the first $s_{init} = 10$ s of v_b to estimate μ and σ , and a threshold was set to $\tau = \mu \pm 3 \cdot \sigma$. The core of the algorithm was based on the following set of rules:

$$\ell(t) = \begin{cases} 1, & \text{if } v_{t'} > \tau \text{ for all } t' \in \{t - \alpha, \dots, t\} \\ 1, & \text{if } \ell(t-1) = 1, v_{t''} > \tau \text{ for any } t'' \in \{t - \beta, \dots, t\} \\ 0, & \text{otherwise} \end{cases} \quad (14)$$

where $\ell = 1$ means that rapid motion was detected. The first rule in Eq. 14 required the threshold τ to be exceeded for α consecutive timesteps before rapid-motion might be detected, with $\alpha \cdot \text{TR} = 30\text{ms}$ here. This setting was chosen to avoid spurious cough detections from noise alone, and could presumably be reduced in duration if/when SNR were higher. The second rule introduced a minimum alert period after each instance of rapid-motion whereby subsequent timesteps get systematically labeled as well, until the signal v_t returns below the threshold τ for at least β timesteps. The rule helped avoid any unrealistically rapid flip-flopping between labeled and non-labeled timesteps, with $\beta \cdot \text{TR} = 600$ ms here. To account for baseline changes in v_b , μ was replaced by a time-varying average μ_t performed over the last $s_\mu = 10$ s of accepted data, resulting in an adaptive value for the threshold, τ_t .

Inspiration for the algorithm came, to some degree, from research on cognitive information processing and decision science (28): In Kahneman's dual model of the human mind, "System 1" (29) is a mode of operation of the brain capable of quickly and automatically processing large amounts of data, generally associated with intuition. The brain's System 1 mode of operation is known to make regular, predictable errors if unchecked by a "System 2" mode, a more analytical mind triggered upon detecting anomalies and generally associated with reasoning. The processing from Eq. 11 and Eq. 14 were meant to be analogous to System 1 and System 2, respectively.

Computational considerations

Finding the best match for each OCM trace can prove computationally expensive given their high dimensionality and high acquisition rate. Simplifications/optimizations were implemented here for Eq. 11 to compute more rapidly: OCM traces were subsampled by calculating the standard deviation at each location x_i along the trace, over the initial five seconds of acquisition, and retaining only the $\hat{N}_x = 200$ indices with highest variation. Although the Gaussian kernel in Equation 11 was always evaluated for the entire database up to any given time, the resulting scores were sorted in order to limit the summation in

Equation 11 to only $K = 6$ terms, which seemed sufficient for our datasets. Higher values led to increased SNR and increased overall reconstruction stability but also increased blurring.

The drop-off of $\mathcal{N}(U_i; U_T, \Sigma)$ for samples further away in data space proved quite pronounced, partly because of the relatively high number of sample points in each OCM trace. This problem, often referred to as the “curse of dimensionality”, led to numerical instability when evaluating the kernel. Additionally, the number of close matches had a dependency on TR. Both issues were handled by appropriately scaling the covariance matrix Σ by an empirically-determined factor taking TR and dimensionality into account:

$$\hat{\Sigma} = \Sigma \times \text{TR} \times (\hat{N}_x \times N_{TR})^2, \quad (15)$$

Our non-optimized Matlab implementation required 90 ms per frame (both planes) once the data-base had filled with roughly two-minute worth of data (2.8 GHz Intel Core i7, 16 GB Ram), i.e., from a real-time perspective, the algorithm currently runs at about 11 fps. Optimized code and better hardware (e.g., GPUs) would be needed to reach the full speed of the OCM data stream, up to 100 fps here.

Hybrid OCM-MRI and OCM-USI imaging setup

An experimental hybrid setup allowed both MRI and OCM signals to be gathered simultaneously, within the bore of a 3T MR system. The method was implemented on a GE Signa HDxt (40 mT/m, 150 T/m/s) system as well as a Siemens Verio system (45 mT/m, 200 T/m/s), and the latter was located in a multi-modality interventional suite (AMIGO suite, Brigham and Women’s Hospital, <http://www.ncigt.org/amigo>). The setup was based on a very simple single-element MR-compatible OCM sensor (Imasonics, 8-mm diameter, 5 MHz, impedance matching layer of 1.5 MRayl, Fig. 1a) small enough for a commercial MR-imaging flexible coil array to wrap over with no detectable loss in MR image quality (GE 8-channel cardiac array, or Siemens body matrix). The OCM sensor was inserted into a specially carved rubber disc (3.5 cm diameter, 1.4 cm thickness, Fig. 1b), positioned on the abdomen of the subject and held in place using an adhesive bandage (Walgreens, bordered gauze 10.2×10.2 cm, Fig. 1c). For safety, fiberoptic temperature probes (Neoptix ReFlex, Qualitrol Company LLC of Fairport, NY) were placed on the OCM sensor, to detect potential heating of the device. Small canals dug into the rubber disk allowed two fiberoptic probes to reach the casing of the OCM sensor while a third probe monitored the temperature of its co-axial cable (Fig. 1b). Ultrasound gel coated both the front and back of the OCM sensor, on the front for proper acoustic coupling with the skin and on the back for proper thermal coupling with the fiberoptic probes. An Olympus 5072PR pulser receiver (Olympus Scientific Solutions Americas Inc, Waltham, MA) fired the OCM sensor through a penetration panel and suitably long coaxial cables. The MR pulse sequence provided the trigger for the pulser receiver, and firings occurred outside of any MR readout window, to avoid artifacts in the MR images. Triggered transducer firings occurred once per TR; although higher OCM acquisition rates would be possible, a rate of one OCM acquisition per TR proved sufficient for present purposes. A PCI digitizer card NI 5122, 150 MHz, 200 MS/s, 12-bit, 512 MB (National Instruments Corporation, Austin, TX) mounted into an off-

the-shelf PC sampled the OCM signals. All electronic devices were located either in the computer room (GE scanner) or in the console room (Siemens scanner).

In addition, “out-of-bore” experiments were performed as follows. Hybrid MRI-OCM data were acquired first, in the Siemens 3T system, followed by USI outside the MRI scan room (BK scanner, Analogic Corporation, Peabody, MA). A Polaris Vicra optical-tracking system (NDI Medical, Cleveland, OH) monitored the movements of the abdominal BK probe. USI and optical-tracking results were sampled, calibrated, time stamped and recorded using the ‘Public software Library for UltraSound’ (PLUS) software package (30). Optically tracked USI results provided a reference standard for the synthetic out-of-scanner real-time MRI results. In the present implementation, reconstruction was performed off-line using code written in a combination of C, C++ and Matlab languages. Great care was taken to make the algorithm causal and compatible with on-line real-time reconstruction.

Human scans

Twenty-two separate time series of MR images were acquired, in 8 consecutive imaging sessions (labeled A through H) involving 7 different human volunteers (with one volunteer imaged twice, on different days), using a steady-state gradient-echo sequence. Informed consent was obtained according to an IRB-approved protocol. Two time series from each imaging session were selected for further processing, typically the last two series of each session. Early in each session, volunteers were instructed to breathe in a slow and regular manner, but as the session progressed they were requested to include coughs and/or gasps. For this reason the last two series of each session were selected here, to ensure that nearly all of the more challenging datasets, those with coughs and/or gasps, would be included in the study. Only for imaging sessions C and H were series other than the last two selected instead: In session C, five time series were acquired but the OCM signal became weak for series #3–5, possibly due to the OCM sensor becoming displaced, and for this reason series #1–2 were selected (as opposed to #4–5). In session H, four series were acquired but only series #2 featured interesting coughs/gasps and series #2–3 were selected (as opposed to #3–4).

Respiration-induced organ motion occurs primarily in the superior/inferior (S/I) direction and to a lesser extent in right/left (R/L) and anterior/posterior (A/P) directions (31–33). For this reason all acquired slices included the S/I direction, as the acquisition process alternated between a pair of intersecting sagittal and coronal slices. Imaging parameters are listed in Table 1 for sessions A through H: N_{echoes} represents the number of k_y lines acquired per TR period, N_t the number of time frames per plane, BW the readout bandwidth, and R the acceleration factor. An R of 3.0 or 3.2 corresponded to 2-fold parallel imaging combined with 2/3 or 5/8 partial-Fourier imaging, respectively. The acquisition time per plane, t , equals $N_y \times TR / N_{echoes} / R$, where $N_y = 192$ was the matrix size in the phase-encoding direction. Temporal resolution was $2 \times t$ (because two planes were acquired), and total scan time was $N_t \times 2 \times t$. Parameters common to all scans included: slice thickness = 5 mm, flip angle = 30°, matrix size = 192 × 192, FOV = 38 × 38 cm², $N_{OCM} = N_{TR}$, and OCM scan rate = 1/TR, which ranged from 56 to 156 Hz. A single frame from session A is shown in Fig. 4, as an example.

Scanning would have stopped if a fiberoptic probe had detected heating beyond the body temperature (threshold at 38°C) on either the OCM sensor or its cable. Sessions A-C were performed on the GE system and sessions D-H on the Siemens system in the AMIGO suite. For sessions D-H, after a few minutes of hybrid OCM-MRI, the volunteer was taken out of the scan room with OCM sensor still in place, and further imaged using an optically-tracked USI probe (2D B-mode, 5 MHz, 12.5 cm depth, 65 dB dynamic range, acquired at 23 fps and frame-grabbed at 15 fps).

In-bore application – Qualitative and quantitative validation

Qualitative validation of synthetic MRI results took the form of M-mode displays, to convey visually the finely detailed nature of the time axis as captured by the high-rate synthesized results. In contrast, quantitative validation involved only the N_t time points T when MRI images were actually acquired (Table 1), providing a reference standard for comparison purposes.

MR images do not, in general, clearly capture a precise point in time; instead, they consist of data acquired over an extended time window (e.g., see Fig. 2). When in need of a clear timestamp, the usual choice would be the instant when k-space center and the main bulk of the signal got sampled. In other words, an MR image acquired at time t is not yet available at time t , because much of its k-space information remains to be sampled. All of our processing, including validation, was causal, which ensured that the not-yet-available MR image at time t did not contribute to the computation of the synthetic MR image reconstructed at time t , an important point when comparing these images for validation purposes.

Landmarks such as blood vessels were tracked in both acquired and synthesized images, using template matching supervised by a human reader. The reader clicked on all landmarks in all frames, to ensure the correct vessels were picked, and template matching was performed afterward to fine tune the location of each landmark. The estimated error was computed as the absolute Euclidean difference in landmark location, between acquired and synthesized images, while excluding frames labeled as cough/gasp by the algorithm. The mean, median and 90th percentile error value for all landmarks and subjects were calculated. A time average was also performed, once sufficient learning had occurred, over a time interval when performance had plateaued; the chosen interval extended here from timestep #48 to #91.

Out-of-bore application – Qualitative validation

Qualitative validation was performed, in Volunteers E-H (Table 1), by comparing synthetic MR images to optically-tracked USI data. At least 4 short sequences (< 1 min) of simultaneous OCM and optically-tracked USI were performed for each volunteer. Although the stream of acquired MRI data had stopped, the creation of synthetic MRI continued based solely on OCM signals and previously learned correlations. Time-matched sagittal synthetic MRI and acquired USI slices were obtained, but these slices did not match spatially thus complicating the task of comparing them. Only for the left lobe were we able to acquire sagittal US liver images throughout the breathing cycle. Because of the upward curvature of

the ribcage toward the solar plexus, the left lobe was seen even at expiration. The chosen MRI sagittal plane, through the main bulk of the organ, was typically more to the right than the USI plane through the left lobe. Furthermore, MRI and USI gave drastically different tissue contrasts; after all, MRI and USI being so notoriously difficult to register was part of the rationale for the present hardware-based solution to begin with. For all these reasons, it was not possible, from our data, to track the same landmarks in both modalities, for quantitative validation purposes. On the other hand, qualitative comparisons of the timing and amplitude of displacements was possible:

Between 1 and 3 clearly-visible structures, e.g. blood vessels and/or the lung-liver boundary, were manually selected in a master frame and automatically tracked in time with template matching. Optical tracking of the USI transducer allowed the in-plane motion of the hand-held probe to be accounted for. Results from all landmarks were averaged for each subject.

Cough and rapid motion detection – Quantitative validation

The human reader also marked the acquired MR images for the presence of motion artifacts. Presumably, motion artifacts should be most visible at times when the algorithm from Eq. 13 detected rapid motion. The presence of motion as detected by the algorithm was validated against the presence of motion artifacts as detected by the human reader. False positives arose when rapid motion failed to cause visible artifacts (e.g., if it happened far from k-space center), and false negatives when artifacts appeared even in the absence of rapid motion (e.g., due to slow motion at a most inopportune time).

A second human reader labeled time segments in the OCM data when rapid motion was readily visible. The purpose of the automated rapid-motion detector from Eq. 14 was to sense coughs/gasps and other instances of rapid motion based on OCM data, and results were validated here against the same task performed by human eyes instead.

Results

In-bore application – Qualitative and quantitative validation

M-mode results from all eight scanning sessions are shown in Fig. 5. Orange lines in the annotated images, on the left side of Fig. 5, indicate the 1D set of locations selected for M-mode display. As can be observed from Fig. 5, Volunteer B took to enthusiastic extremes our suggestion to occasionally cough or gasp. A movie showing an example of within-the-bore results is also available (Supporting Video S1). Furthermore, open-source code and sample data were made publicly available at <https://github.com/fpreiswerk/OCMDemo>.

Green circles in Fig. 5 indicate the tracked landmarks. Figure 6 shows the Euclidean error between synthetic and acquired frames, for all landmarks and all subjects combined, as a function of time. As time progressed, the algorithm showed rapid initial learning (see pink line in Fig. 6) and a tendency to plateau after a few breathing cycles. Over all acquisitions, a mean, median and 90th percentile error of 1.0 pixel (2.1 mm), 0.6 pixel (1.15 mm) and 2.2 pixel (4.4 mm) was measured over the plateau (Table 2). Perhaps not surprisingly, datasets featuring calmer and/or more regular breathing patterns typically led to smaller errors:

volunteers C, F, G and H had some of the lowest tracking errors while volunteer B led to the highest error.

Outside the MRI bore, qualitative validation

Figure 7 shows MRI data acquired inside the bore with subject H, and synthetic MRI data generated outside the bore based on OCM and learned correlations alone, along with time-matched USI data. The leftmost and rightmost green dashed lines in Fig. 7 highlight full-inspiration and full-expiration phases, respectively. Portions of the time axes, marked with blue annotations, are zoomed to better show how the very coarse temporal sampling with MRI (1.7 fps) can be made much smoother (100 fps) in the synthesized results, even from outside the bore.

In Fig. 8a, poor agreement was obtained between synthetic MRI and USI as subject E breathed deeply inside the scanner and more normally outside, affecting the algorithm. Much better agreement was obtained in Fig. 8b, with subject F. In Fig. 8c, coughs (gray-shaded areas) affected performance but the algorithm could recover to some degree in-between coughs. As in Fig. 8b, results in Fig. 8d from subject H showed good qualitative agreement between synthetic MRI and USI.

Cough and rapid motion detection – Quantitative validation

In Fig. 5, intervals highlighted in red and yellow represent instances of coughs/gasps, as detected by our algorithm (Eq. 14) and by a human reader looking at OCM data, respectively. The resulting sensitivity and specificity values obtained for the algorithm from Eq. 14 were 0.90 / 0.98 inside the scanner, and 0.87 / 0.99 outside the scanner. In other words, Eq. 14 proved an effective and automated alternative to human eyes in the task of detecting periods of rapid motion based on OCM data. When testing the ability of the motion-detection algorithm to predict MR image degradations, as seen by a human reader, lower values were obtained for sensitivity and specificity, 0.42 and 0.97, respectively.

Discussion

Detecting and compensating for organ motion can be an important component of MR-guided therapeutic procedures such as biopsies or ablations. Rapid and irregular motion, such as that caused by coughing or gasping, generally proves particularly challenging to address. A strategy was proposed here based on a hybrid US-MRI system, which used a Bayesian reconstruction to handle the flow of incoming hybrid data and a rapid-motion detection algorithm to detect periods of unusual motion activity. While there was nothing special about the ultrasound signal or the MR images obtained here, interesting behaviors emerged when detecting correlations between the two. After a training period extending for a few breathing cycles the US signal became, in essence, a surrogate for MR images (4, 8). While previously unseen behavior, such as a deeper breath, could confuse the algorithm momentarily, continuous-learning ensures that in time the algorithm would adapt and properly deal with repeated instances. The method could be used to generate extra time frames in-between actually-acquired ones and thereby greatly boost temporal resolution in the process. Alternately MR images could be generated even after the volunteer has been

physically removed from the MR scanner employing only the incoming US signals and previously-learned correlations. This intriguing ‘scannerless real-time MRI’ behavior, in principle at least, might potentially allow image-guided procedures to be performed outside the confines of the imaging bore.

As a limitation, several potentially important tradeoffs were left mostly unexplored here. At the acquisition stage, the operating frequency, size, acoustic power and precise placement of the OCM sensor were not optimized. Similarly, some relevant MRI parameters such as the tradeoff between temporal and spatial resolution were not optimized in any specific manner. More specifically, smaller pixels might have led to smaller location errors: the error as averaged over all sessions and landmarks was 1.0 pixel here, suggesting that smaller pixels might have given smaller errors. On the other hand, faster MRI could have helped to better capture motion. In real-time imaging, there is a direct relationship between frame rate, MRI scan time and temporal resolution. In contrast, in our method temporal resolution and frame rate were both improved, but not scan time per frame. Because the strongest MRI signal comes from k-space center, MRI images tend to represent a narrower time window than their acquisition time would suggest, allowing temporal resolution to be improved here by increasing frame rate alone. However, a more compact MRI acquisition with shorter scan time should in principle also help improve temporal resolution, no matter how high a frame rate the proposed method might achieve.

Further limitations included the small number of imaging sessions (eight), and the fact that all volunteers were healthy. Out-of-bore results with optically-tracked USI were obtained from only five sessions, mostly as a proof of concept. Most envisioned applications would likely require a real-time implementation of the method, whereby results are generated and displayed in near real-time, as opposed to saved and reconstructed later. Such an implementation remains as future work.

Some of the main possible alternatives to our proposed OCM sensor for monitoring organ motion include navigator echoes and respiratory bellows. While very effective at motion detection, navigator echoes often compete with the image acquisition process, decreasing the efficiency of the overall scan. Furthermore, navigator echoes are, of course, unavailable outside the MRI scanner; this latter point may prove most relevant, as the out-of-bore application presented here would not be possible using navigator echoes. While respiratory bellows detect organ motion only indirectly, through changes in torso perimeter, two or more of them placed at different locations might allow the proper characterization of motion from abdominal and/or thoracic breathing. However, two or more bellows would likely be cumbersome and get in the way of an interventionalist. In contrast, the present OCM hardware contacts only a small area of a subject’s torso, as covered with an adhesive bandage. Furthermore, the OCM hardware generates much richer signals.

Possible applications were envisioned for the proposed method, such as monitoring tissue ablations for the in-bore implementation, and for the out-of-bore implementation, needle / applicators placement and/or hardware-based registration. Effects on the algorithm from the presence of a needle / applicator, or how long pre-learned correlations may remain valid, are examples of questions left unanswered in this study. The OCM sensor has been successfully

scanned in a PET-CT scanner and caused no sizeable artifact, suggesting that the OCM hardware might also be used as a common denominator to help register images between MR and other modalities.

Conclusion

A hybrid system based on a single-element US transducer and an MR scanner was proposed, to help boost temporal resolution in real-time MRI. A Bayesian algorithm was developed to handle the flow of hybrid data and criteria were developed to detect periods of unusual activity (e.g., coughing). The system enables MRI contrast to be achieved, even for subjects outside the MR bore, and thus may have interesting applications for the field of image-guided therapy.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Financial support from grants NIH R01CA149342, P41EB015898, R21EB019500 and SNSF P2BSP2 155234 is duly acknowledged. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

References

1. Silverman SG, Collick BD, Figueira MR, Khorasani R, Adams DF, Newman RW, Topulos GP, Jolesz FA. Interactive MR-guided biopsy in an open-configuration MR imaging system. *Radiology*. 1995; 197(1):175–181. [PubMed: 7568819]
2. Panych, LP., Tokuda, J. Real-time and interactive MRI. In: Jolesz, F., editor. *Intraoperative Imaging and Image-Guided Therapy*. New York, USA: Springer; 2014. p. 193-210.
3. Schwartz BM, McDannold NJ. Ultrasound echoes as biometric navigators. *Magn Reson Med*. 2013; 69(4):1023–1033. [PubMed: 22648783]
4. Toews, M., Mei, C-S., Chu, R., Hoge, WS., Schwartz, BM., Wang, G., Panych, LP., Madore, B. Boosting MR Temporal Resolution using Rapid Ultrasound Measurements, for Motion-Tracking Purposes. *Proceedings of the International Society of Magnetic Resonance in Medicine*; Salt Lake City, UT, USA. 2013. p. 0478
5. Toews, M., Mei, C-S., Chu, R., Hoge, WS., Panych, LP., Madore, B. Detecting rapid organ motion using a hybrid MR-ultrasound setup and Bayesian data processing. *Proceedings of the International Society of Magnetic Resonance in Medicine*; Milan, Italy. 2014. p. 0681
6. Preiswerk, F., Hoge, WS., Toews, M., Chiou, J-yG, Chauvin, L., Panych, LP., Madore, B. Speeding-up MR acquisitions using ultrasound signals, and scanner-less real-time MR imaging. *Proceedings of the International Society of Magnetic Resonance in Medicine*; Toronto, Canada. 2015. p. 0863
7. Preiswerk, F., Toews, M., Hoge, WS., Chiou, J-y, Panych, L., Wells, W., III, Madore, B. Hybrid Ultrasound and MRI Acquisitions for High-Speed Imaging of Respiratory Organ Motion. In: Navab, N.Hornegger, J.Wells, W., Frangi, A., editors. *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. Springer International Publishing; 2015. p. 315-322.
8. Arvanitis CD, Livingstone MS, McDannold N. Combined ultrasound and MR imaging to guide focused ultrasound therapies in the brain. *Phys Med Biol*. 2013; 58(14):4749–4761. [PubMed: 23788054]
9. Kording, F., Schoennagel, B., Much, C., Ueberle, F., Kooijman, H., Yamamura, J., Adam, G., Wedegaertner, U. MR compatible Doppler-ultrasound device to trigger the heart frequency in

Cardiac MRI: comparison to ECG. Proceedings of the International Society of Magnetic Resonance in Medicine; Salt Lake City, UT, USA. 2013. p. 1421

10. Petrusca L, Cattin P, De Luca V, Preiswerk F, Celicanin Z, Auboiroux V, Viallon M, Arnold P, Santini F, Terraz S, Scheffler K, Becker CD, Salomir R. Hybrid ultrasound/magnetic resonance simultaneous acquisition and image fusion for motion monitoring in the upper abdomen. *Invest Radiol*. 2013; 48(5):333–340. [PubMed: 23399812]
11. Günther M, Feinberg DA. Ultrasound-guided MRI: preliminary results using a motion phantom. *Magn Reson Med*. 2004; 52(1):27–32. [PubMed: 15236363]
12. Feinberg DA, Giese D, Bongers DA, Ramanna S, Zaitsev M, Markl M, Günther M. Hybrid ultrasound MRI for improved cardiac imaging and real-time respiration control. *Magn Reson Med*. 2010; 63(2):290–296. PMID:PMC2813925. [PubMed: 20025068]
13. Breast cancer – combining imaging techniques for quicker and gentler biopsies. [Accessed February 29, 2016] Fraunhofer-Gesellschaft website. <http://www.fraunhofer.de/en/press/research-news/2013/oktober/breast-cancer-combining-imaging-techniques-for-quicker-and-gentler-biopsies>. Published October 29, 2013
14. Odille F, Vuissoz PA, Marie PY, Felblinger J. Generalized reconstruction by inversion of coupled systems (GRICS) applied to free-breathing MRI. *Magn Reson Med*. 2008; 60(1):146–157. [PubMed: 18581355]
15. von Siebenthal M, Székely G, Gamper U, Boesiger P, Lomax A, Cattin P. 4D MR imaging of respiratory organ motion and its variability. *Phys Med Biol*. 2007; 52(6):1547–1564. [PubMed: 17327648]
16. Celicanin Z, Bieri O, Preiswerk F, Cattin P, Scheffler K, Santini F. Simultaneous acquisition of image and navigator slices using CAIPIRINHA for 4D MRI. *Magn Reson Med*. 2015; 73(2):669–676. [PubMed: 24604250]
17. Thomenius, KE. Evolution of ultrasound beamformers. *IEEE Ultrasonics Symp*; San Antonio, TX, USA. 1996. p. 1615–1622.
18. Shattuck DP, Weinshenker MD, Smith SW, von Ramm OT. Explososcan: a parallel processing technique for high speed ultrasound imaging with linear phased arrays. *J Acoust Soc Am*. 1984; 75(4):1273–1282. [PubMed: 6725779]
19. Capon J. High resolution frequency-wavenumber spectrum analysis. *Proc IEEE*. 1969; 57(8):1408–1418.
20. Synnevag JF, Austeng A, Holm S. Adaptive beamforming applied to medical ultrasound imaging. *IEEE Trans Ultrason Ferroelectr Freq Control*. 2007; 54(8):1606–1613. [PubMed: 17703664]
21. Madore B, Meral F. Reconstruction algorithm for improved ultrasound image quality. *IEEE Trans Ultrason Ferroelectr Freq Control*. 2012; 59(2):217–230. NIHMS360420. [PubMed: 24626030]
22. Ellis MA, Viola F, Walker WF. Super-resolution image reconstruction using diffuse source models. *Ultrasound Med Biol*. 2010; 36(6):967–977. [PubMed: 20447760]
23. Parzen E. On estimation of a probability density function and mode. *Ann Math Statist*. 1962; 33(3):1065–1076.
24. Watson GS. Smooth regression analysis. *Sankhya: The Indian Journal of Statistics, Series A*. 1964; 26:359–372.
25. Marple SL. Computing the discrete-time "analytic" signal via FFT. *IEEE Trans Sig Proc*. 1999; 47(9):2600–2603.
26. Rueckert D, Sonoda LI, Hayes C, Hill DL, Leach MO, Hawkes DJ. Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Trans Med Imaging*. 1999; 18(8):712–721. [PubMed: 10534053]
27. Lagarias JC, Reeds JA, Wright MH, Wright PE. Convergence properties of the Nelder-Mead simplex method in low dimensions. *Siam J Optim*. 1998; 9(1):112–147.
28. Kahneman, D. Attention and effort. Englewood Cliffs, N.J.: Prentice-Hall; 1973. p. 246
29. Kahneman, D. Thinking, fast and slow. New York, N.Y.: Farrar, Straus and Giroux - Macmillan; 2011. p. 483
30. Lasso A, Heffter T, Rankin A, Pinter C, Ungi T, Fichtinger G. PLUS: open-source toolkit for ultrasound-guided intervention systems. *IEEE Trans Biomed Eng*. 2014; 61(10):2527–2537. PMID:PMC4437531. [PubMed: 24833412]

31. Wade OL. Movements of the thoracic cage and diaphragm in respiration. *J Physiol.* 1953; 124:193–212.
32. Suramo I, Paivansalo M, Myllyla V. Cranio-caudal movements of the liver, pancreas and kidneys in respiration. *Acta Radiol Diagn (Stockh).* 1984; 25(2):129–131. [PubMed: 6731017]
33. Wang Y, Riederer SJ, Ehman RL. Respiratory motion of the heart: kinematics and the implications for the spatial resolution in coronary imaging. *Magn Reson Med.* 1995; 33(5):713–719. [PubMed: 7596276]

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Fig. 1.

a) A single-element 8-mm MR-compatible transducer was employed. b) A soft-plastic holder was fabricated to hold the transducer and the fiberoptic temperature sensors in place. c) The transducer, temperature sensors and holder were kept in place on the volunteer's abdomen using an adhesive bandage. A blue sheath of foam material surrounded the transducer cable over much of its length, to insulate it thermally from the volunteer.

Fig. 2.

A few of the main variables involved in the proposed method are depicted here, such as the number of TR periods required to acquire an MR image, N_{TR} , the number of OCM traces employed to characterize an image, N_{OCM} , and the collection of past knowledge, D , available to guide the reconstruction of a current synthetic image, I_t .

Fig. 3. OCM signals, sagittal and coronal images from subject B are shown. A cough occurred during the displayed interval, causing artifacts in the MR images. Note how clearly the OCM signals captured the instance of rapid motion. In Eqs 13 and 14, the partial derivative with respect to time of the OCM signal was employed to help detect gasps and/or coughs.

Fig. 4.

Time series of sagittal and coronal slices were acquired in a time-interleaved fashion. Because sagittal and coronal images are acquired one after the other and cannot capture the exact same respiratory state, a discontinuity appears where the planes intersect. By greatly improving temporal resolution for both planes, the proposed method essentially resolves this issue. Please see the M-mode displays, under the images, where the 1D locations where the planes intersect are associated with the blue color (sagittal image) and the yellow color (coronal image). For our results on the right, due to much improved temporal resolution, the two planes move in synchrony and nearly everywhere blue and yellow combine into green. In contrast, for the acquired data of the left, motion was jagged and discontinuities appeared between the planes, creating many regions where the color deviated from green. Readout gradients for coronal and sagittal slices had different polarities; for this reason, a relative shift of 2 pixels along the superior-inferior direction was applied when generating the M-mode displays, to compensate for susceptibility-induced shifts.

Fig. 5. Hybrid OCM-MRI results are shown in M-mode format for each of the eight imaging sessions, for the first of two acquisitions. As usual in M-mode displays, the horizontal axis is time and the vertical axis represents a 1D image. Orange lines in the images on the left mark the 1D locations selected for M-mode display. Green circles in these same annotated images indicate the landmarks employed for validation purposes, in Fig. 6 and Table 2. Intervals highlighted in red in the M-mode displays represent cough/gasp events as detected by our algorithm (Eq. 14). In contrast, intervals highlighted in yellow represent cough/gasp events

as detected by a human reader looking at OCM signals. Clearly, Volunteer B proved especially enthusiastic in following the request to challenge the algorithm with coughing events.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Fig. 6.

Landmarks in the liver, indicated with green circles in Fig. 5, were tracked both in synthetic OCM-MRI and in acquired MRI images. The Euclidean error between synthetic and acquired images is plotted here as a function of time: the algorithm learned from incoming timeframes, rapidly at first (see pink dashed line) and plateaued after a few breathing cycles (blue dashed line). The mean, median and 90th percentile error, over all landmarks and all subjects, plateaued here at 1.0, 0.6 and 2.2 pixel(s), respectively.

Fig. 7. MRI data were acquired (left column), and then the subject was taken out of the MRI scan room. Based on OCM signals and learned correlations only, streams of synthetic MRI images could still be generated (middle column). Green lines indicate examples of full inspiration and expiration, while blue markings show the time segment zoomed in the lower row; note the coarser nature of the temporal resolution from acquired MRI (1.7 fps here) compared to synthetic MRI (100 fps). Images in the right-most column show the 1D locations selected for M-mode display.

Fig. 8.

The motion of structures as seen in synthetic MRI images (red and orange curves), generated from outside the MRI scan room, was compared to that obtained from time-matched images from a commercial ultrasound scanner (green curve). Dots represent measured locations and solid lines represent fitted splines. Except for the first part of (a), where Subject E challenged the algorithm with a different breathing pattern, the results showed good qualitative agreement between synthetic MRI and ultrasound imaging results. In (c), the algorithm detected all three coughs (shaded regions) and the algorithm did, at least in part, recover in-between coughs.

Table 1

MR imaging parameters.

Session	TR (ms)	N_{echoes}	R	t (s)	N_t	BW (kHz)
A, B.1, C	18	2	3.0	0.576	92	± 250
B.2	18	2	3.0	0.576	192	± 250
D.1	10	1	3.2	0.6	200	± 37.4
D.2	10	1	3.2	0.6	165	± 37.4
E-H	10	1	3.2	0.6	100	± 37.4

Table 2

Mean absolute errors over timesteps 48–91 (plateau in Fig. 6) expressed in pixels and in mm, for all eight human scanning sessions A–H, two scans per session.

MAE	A1	A2	B1	B2 [†]	C1	C2	D1	D2	E1	E2	F1	F2	G1	G2	H1	B2 [‡]
	px	1.4	0.8	2.0	4.1	0.6	1.1	1.3	1.9	1.2	1.3	0.6	0.5	0.6	0.7	0.4
	mm	2.8	1.6	4.0	8.1	1.1	2.3	2.6	3.8	2.3	2.7	1.0	0.9	1.2	1.4	0.9

[†]Worst case, with highest mean absolute error

[‡]Best case, with lowest mean absolute error