Hybrid MRI-Ultrasound Acquisitions, and Scannerless Real-Time Imaging

Frank Preiswerk,[1]* Matthew Toews,[2] Cheng-Chieh Cheng,[1] Jr-yuan George Chiou,[1] Chang-Sheng Mei,[3] Lena F. Schaefer,[1] W. Scott Hoge,[1] Benjamin M. Schwartz,[4] Lawrence P. Panych,[1] and Bruno Madore[1]

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 the 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 with those from acquired images: The mean error was 1.0 pixel (2.1 mm), 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. **Magn Reson Med 78:897–908, 2017. © 2016 International Society for Magnetic Resonance in Medicine.**

Key words: 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,

[1]Department of Radiology, Brigham and Women's Hospital, Harvard Medical School, Boston, Massachusetts, USA.
[2]The Laboratory for Imagery, Vision and Artificial Intelligence, École de Technologie Supérieure, Montréal, QC, Canada.
[3]Department of Physics, Soochow University, Taipei, Taiwan.
[4]Google Inc, New York, New York, USA.

*Correspondence to: Frank Preiswerk, PhD, Department of Radiology, BWH, Thorn Building, Room 328, Boston, MA 02115, USA.

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.

Although 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 (13). A main realization at the basis of the present work is that although a small transducer is insufficient to generate two-dimensional (2D) or three-dimensional (3D) spatially resolved images, it may not need to because MRI is spatially resolved and correlations exist between the two signal types. Compared with 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
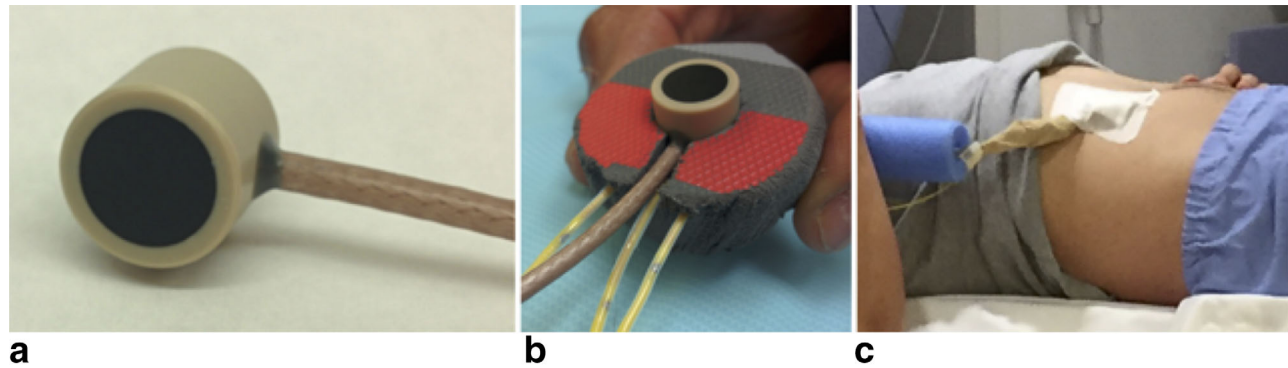
FIG. 1. A single-element 8-mm MR-compatible transducer was employed (a). A soft-plastic holder was fabricated to hold the transducer and the fiberoptic temperature sensors in place (b). The transducer, temperature sensors, and holder were kept in place on the volunteer's abdomen using an adhesive bandage (c). A blue sheath of foam material surrounded the transducer cable over much of its length, to insulate it thermally from the volunteer.

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, such as in four dimensional MRI (4D MRI) (15,16). The main advantages of employing US signals instead of navigator echoes are as follows: (i) 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; and (ii) US signals are available outside and 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 22 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 MR images to acquired ones, whereas outside the bore, qualitative validation involved time-matched and optically tracked US imaging (USI) data.

METHODS

Organ-Configuration Motion 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 multimodality data sets using the synthetic MRI results as a common thread between all successive scans and acquisitions. The learning algorithm is described subsequently, first for a single imaging plane and then for the more general multiplane 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 I_T image was obtained from a set of individual MR signals or k-space lines, acquired over a number of repetition time (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_t, from the most current set of OCM data available, U_t, based on past experience, D. More specifically the expectation of I_t is computed as

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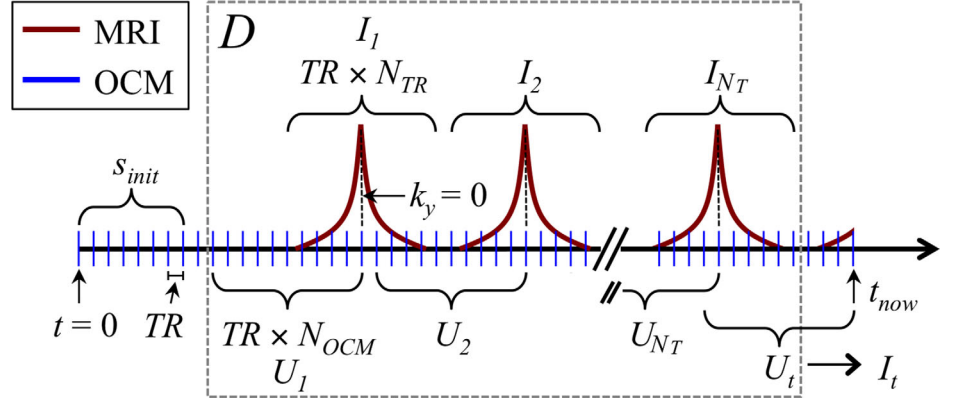
$$E[I_t|U_t, D] = \int I_t \; p(I_t|U_t, D) dI_t = \frac{\int I_t \; p(I_t, U_t|D) dI_t}{p(U_t|D)}. \quad [1]$$

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 an instance-based method for computing Equation [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

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

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

where $\delta(\cdot)$ is a Dirac delta function. Using the training set D as defined previously, 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 Equation [2], the bandwidth is infinitesimally small, resulting in the Dirac delta function $\delta(I_t - I_T)$ centered at I_t, whereas for k_b we use a Gaussian kernel with covariance matrix Σ. The numerator in Equation [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 Equation [1] is obtained from

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

Combining Equations [6] and [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 Equation [8] is of a form consistent with 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 being 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 $\left(t - TR \times (N_{OCM} - 1)\right)$ in the past. Accordingly, Equation [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 U_T captures a time window of width $TR \times N_{OCM}$, Equation [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 Equation [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 being the MR images at a first and second plane, Equation [1] is adapted to compute both planes jointly, as follows:

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

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 $[J_t]_{\bar{L}}$ represent 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):

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

We represent $p([I_t]_L, [J_t]_L|U_t, D)$ with a Gaussian model $N([I_t]_L; [J_t]_L, \Sigma_L)$, with Σ_L for the covariance and $[I_t]_L$ (or $[J_t]_L$) for the mean. Applying KDE as previously, the two-plane equivalent of Equation [8] becomes

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

where U_i and U_j are the 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 5 s of OCM acquisition, and assuming these points to be independent in terms of noise. Generalization of Equation [11] to an arbitrary number of planes parallel to I or J involves repeated application of the product rule in Equation [10].

Reconstruction of Out-of-Bore Results

Out-of-bore results, just like in-bore results, were reconstructed using Equation [11]; one difference, however, is that 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 nonrigid 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 a cost function as follows:

$$C = -C_{similarity}\left(\bar{U}_{in}, \ T(\bar{U}_{out})\right) + \lambda C_{smooth}(T), \qquad [12]$$

where $\bar{U}_{in}$, $\bar{U}_{out}$ 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, $C_{similarity}\left(A_t(x), B_t(x)\right) = \sum_{x=1}^{N_x} (A_t(x) - B_t(x))^2$ and the smoothness of the deformation was constrained according to $C_{smooth}(T) = \frac{1}{N_x}\int_{x=1}^{N_x}\left(\frac{d^2 T}{dx^2}\right)^2 dx$, 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 six resolution levels with $\lambda = 1.0$ to efficiently solve Equation [12] using the simplex algorithm (27).

Cough/Gasp Detector

Coughing and gasping may cause rapid motion that Equation [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, based on OCM signals U_t, proved sensitive to motion-induced changes (see Fig. 3), as follows:

$$v_t := \left|\frac{\partial U_t}{\partial t}\right|_1 \approx \left|\frac{U_t - U_{t-1}}{TR}\right|_1, \qquad [13]$$

where $|\cdot|_1$ is an l1-norm operator. Although v_t tends to show clear peaks in the presence of rapid motion, it is also sensitive to noise and physiological signals (eg, heartbeats). A more sophisticated classifier ℓ was constructed as follows: A Gaussian model $N(v_t; \mu, \sigma)$ was fitted to the first $s_{init} = 10$ s of v_t, 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:

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

where $\ell = 1$ indicates that rapid motion was detected. The first rule in Equation [14] required the threshold τ to be exceeded for α consecutive time steps before rapid motion might be detected, with $\alpha \cdot TR = 30$ ms. This setting was chosen to avoid spurious cough detections from noise alone, and could presumably be reduced in duration if/when the signal-to-noise ratio (SNR) was higher. The second rule introduced a minimum alert period after each instance of rapid motion, in which subsequent time steps get systematically labeled until the signal v_t returns below the threshold τ for at least β time steps. The rule helped to avoid any unrealistically rapid flip-flopping between labeled and nonlabeled time steps, with $\beta \cdot TR = 600$ ms. To account for baseline changes in v_t, 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
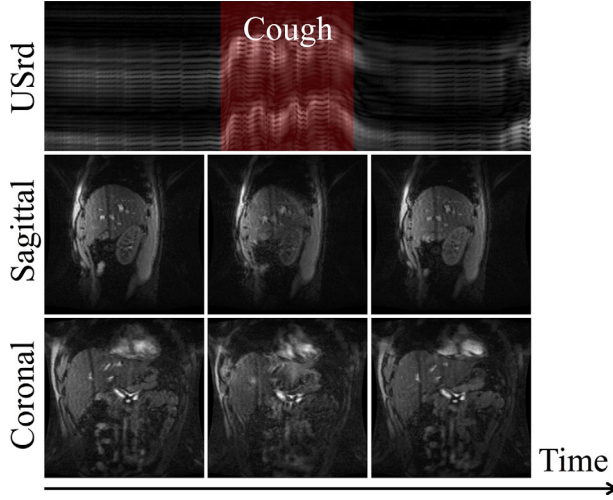
FIG. 3. OCM signals as well as sagittal and coronal images from subject B. A cough occurred during the displayed interval, causing artifacts in the MR images. Note how clearly the OCM signals captured the instance of rapid motion. In Equations [13] and [14], the partial derivative with respect to time of the OCM signal was employed to help detect gasps and/or coughs.

the brain that is 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 Equations [11] and [14] were intended to be analogous to Systems 1 and 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 Equation [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 5 s 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 appeared to be sufficient for our data sets. Higher values led to increased SNR and increased overall reconstruction stability, but also increased blurring.

The drop-off of $N(U_t; U_T, \Sigma)$ 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 nonoptimized MATLAB (The MathWorks, Natick, Massachusetts, USA) implementation required 90 ms per frame (both planes) once the database had filled with approximately 2 min worth of data (2.8 GHz Intel Core i7, 16 GB Ram) (ie, from a real-time perspective, the algorithm currently runs at approximately 11 fps). Optimized code and better hardware (eg, graphical processing units) would be needed to reach the full speed of the OCM data stream, up to 100 fps.

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 3 Tesla (T) MR system. The method was implemented on a GE Signa HDxt (40 mT/m, 150 T/m/s) system (Milwaukee, Wisconsin, USA) as well as a Siemens Verio system (45 mT/m, 200 T/m/s) (Erlangen, Germany), and the latter was located in a multimodality 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, Voray-sur-l'Ognon, France, 8-mm diameter, 5 MHz, impedance matching layer of 1.5 MRayl, Fig. 1a) that is 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, Deerfield, Illinois, USA bordered gauze 10.2 × 10.2 cm, Fig. 1c). For safety, fiber-optic temperature probes (Neoptix ReFlex, Qualitrol Company, Fairport, New York, USA) were placed on the OCM sensor to detect potential heating of the device. Small canals dug into the rubber disk allowed two fiber-optic probes to reach the casing of the OCM sensor, while a third probe monitored the temperature of its coaxial 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 fiber-optic probes. An Olympus 5072PR pulser receiver (Olympus Scientific Solutions Americas, Waltham, Massachusetts, USA) 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, Austin, Texas) 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, Peabody, Massachusetts, USA). A Polaris Vicra optical-tracking system (NDI Medical, Cleveland, Ohio, USA) 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 Ultra-Sound (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 eight consecutive imaging sessions (labeled A through H) involving seven 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 to ensure that nearly all of the more challenging data sets, 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 the 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 two-fold parallel imaging combined with 2/3 or 5/8 partial-Fourier imaging, respectively. The acquisition time per plane, Δt, equals $N_y \times TR / N_{echoes} / R$, in which $N_y = 192$ was the matrix size in the phase-encoding direction. Temporal resolution was $2 \times \Delta t$ (because two planes were acquired), and the total scan time was $N_t \times 2 \times \Delta t$. Parameters common to all scans included slice thickness = 5 mm, flip angle = 30°, matrix size = 192 × 192, field of view (FOV) = 38 × 38 cm^2, $N_{OCM} = N_{TR}$, and OCM

scan rate = 1/TR, which ranged from 56 to 156 Hz. A single frame from session A is shown in Figure 4 as an example.

Scanning would have stopped if a fiber-optic 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 the 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 time step #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 four 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

Table 1
MR Imaging Parameters

Session	TR (ms)	N_{echoes}	R	Δt (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

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. 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 of these reasons, it was not possible, from our data, to track the same landmarks in both modalities for quantitative validation purposes. However, qualitative comparisons of the timing and amplitude of displacements were possible:

Between one and three clearly visible structures, eg, 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.
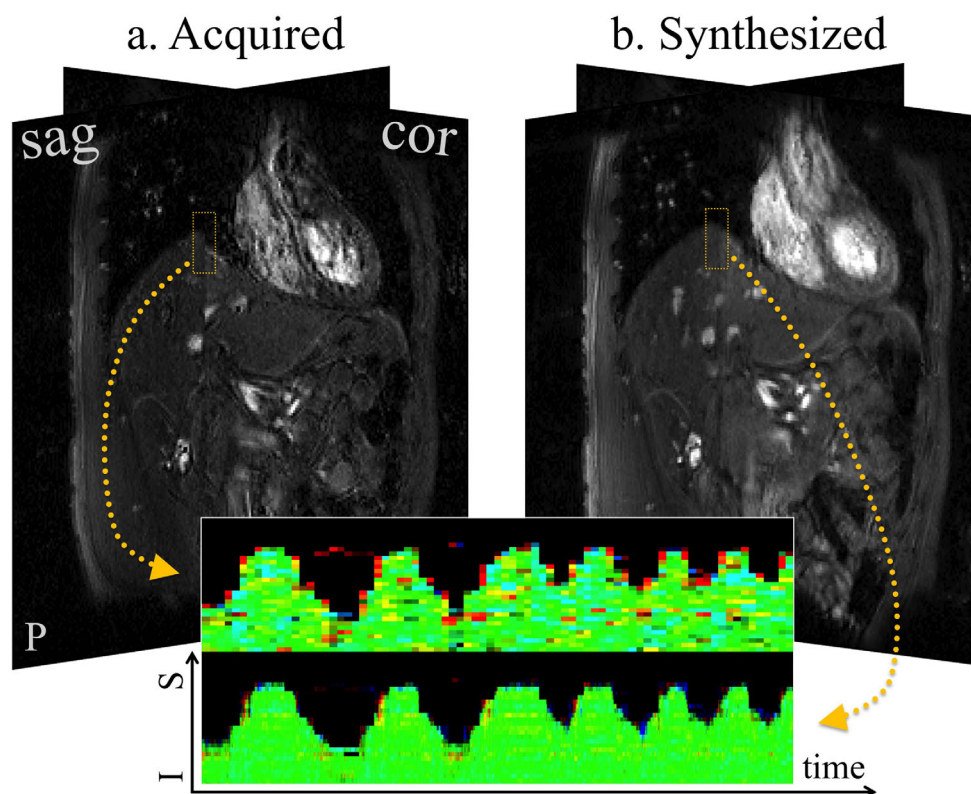

FIG. 4. Time series of sagittal and coronal slices were acquired in a time-interleaved fashion. Because sagittal and coronal images are acquired one after the other and cannot capture the exact same respiratory state, a discontinuity appears where the planes intersect. By greatly improving temporal resolution 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 on 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.
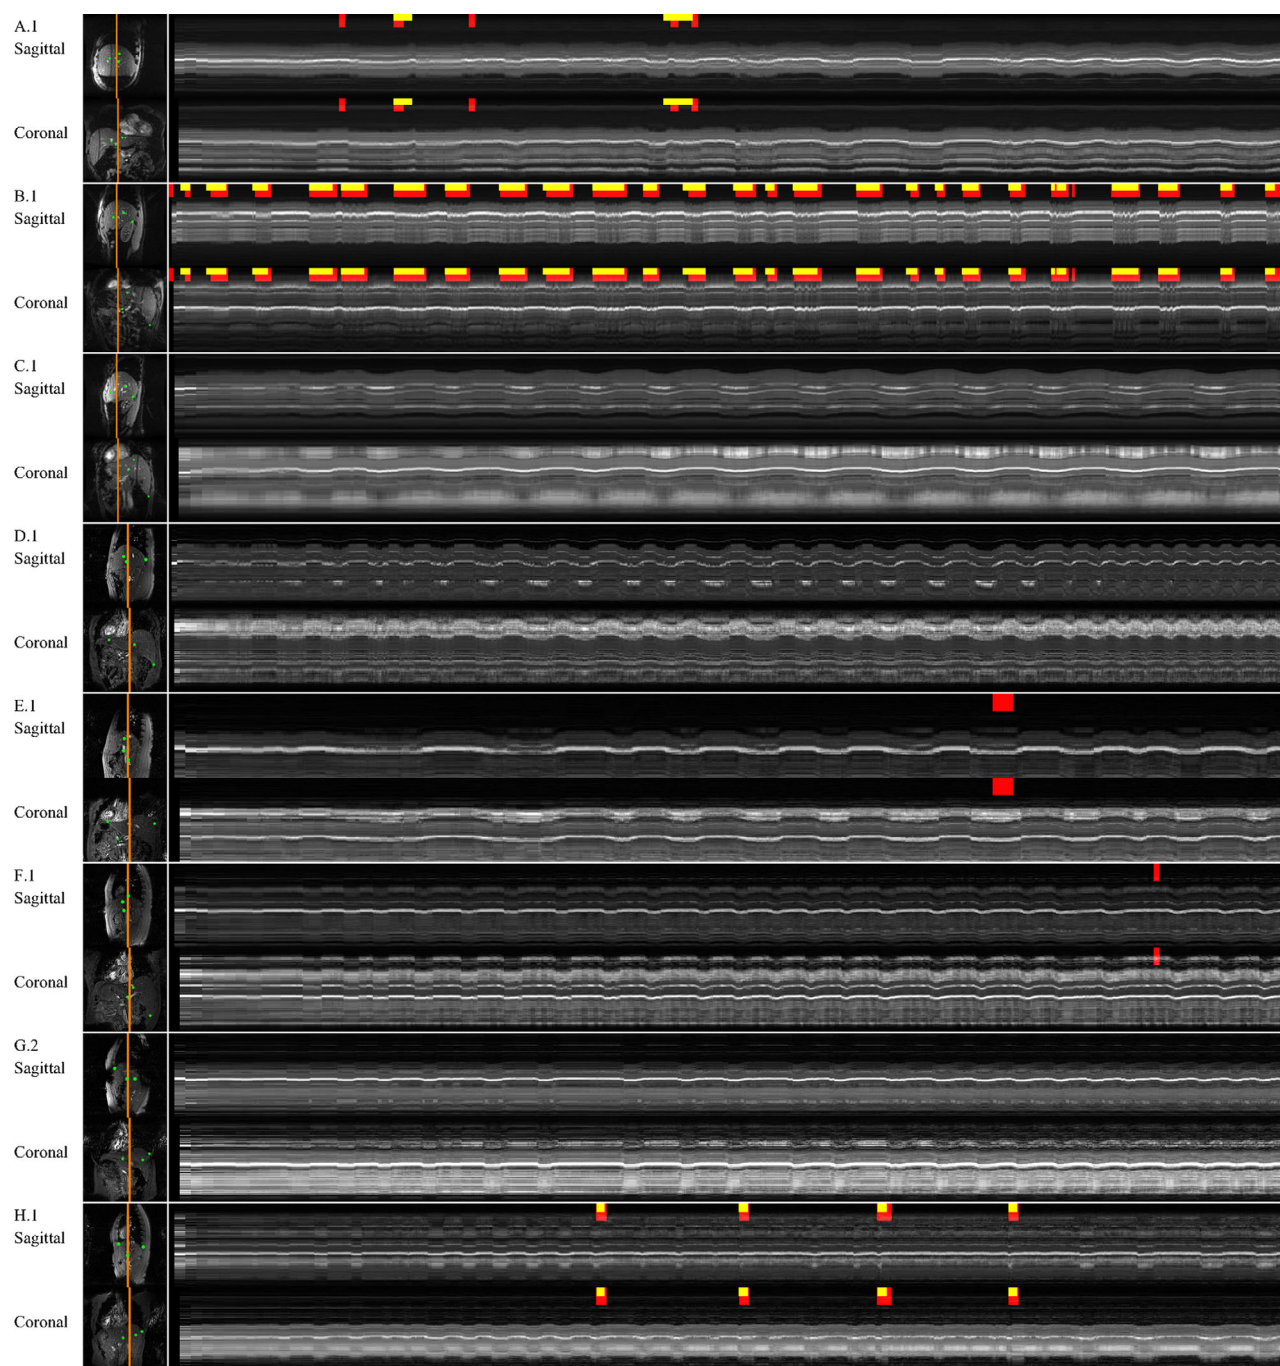
FIG. 5. Hybrid OCM-MRI results are shown in M-mode format for each of the eight imaging sessions, for the first of two acquisitions. As usual in M-mode displays, the horizontal axis is time and the vertical axis represents a 1D image. Orange lines in the images on the left mark the 1D locations selected for M-mode display. Green circles in these same annotated images indicate the landmarks employed for validation purposes, in Figure 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.

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 Equation [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 (eg, if it happened far from k-space center), and false negatives when artifacts appeared even in the absence of rapid motion (eg, due to slow motion at a most inopportune time).

Table 2

Mean Absolute Errors over Time Steps 48–91 (Plateau in Fig. 6), Expressed in Pixels and in Millimeters, for All Eight Human Scanning Sessions A–H, Two Scans per Session

		A1	A2	B1	B2†	C1	C2	D1	D2	E1	E2	F1	F2	G1	G2	H1	H2‡
MAE	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	0.5	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	1.1	1.0	0.9	1.2	1.4	**0.9**

†Worst case, with highest mean absolute error.

‡Best case, with lowest mean absolute error.

A second human reader labeled the time segments in the OCM data when rapid motion was readily visible. The purpose of the automated rapid-motion detector from Equation [14] was to sense coughs/gasps and other instances of rapid motion based on OCM data, and results were validated 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 Figure 5. Orange lines in the annotated images, on the left side of Figure 5, indicate the one-dimensional (1D) set of locations selected for M-mode display. As can be observed from Figure 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.

The green circles in Figure 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
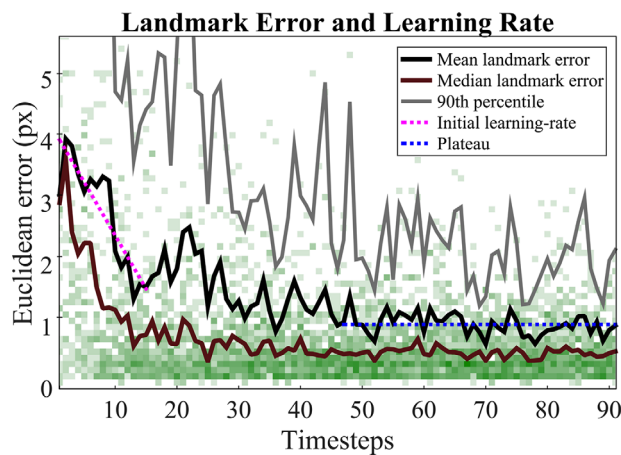

FIG. 6. Landmarks in the liver, indicated with green circles in Figure 5, were tracked both in synthetic OCM-MRI and in acquired MRI images. The Euclidean error between synthetic and acquired images is plotted as a function of time: The algorithm learned from incoming timeframes, rapidly at first (see 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 at 1.0, 0.6, and 2.2 pixels, respectively.

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 pixels (1.15 mm), and 2.2 pixels (4.4 mm) was measured over the plateau (Table 2). Perhaps not surprisingly, the data sets 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, whereas 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 Figure 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 Figure 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 Figure 8b, with subject F. In Figure 8c, coughs (gray-shaded areas) affected performance, but the algorithm could recover to some degree in between coughs. As in Figure 8b, results in Figure 8d from subject H showed good, qualitative agreement between synthetic MRI and USI.

Cough and Rapid Motion Detection—Quantitative Validation

In Figure 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 the OCM data, respectively. The resulting sensitivity and specificity values obtained for the algorithm from Equation [14] were 0.90/0.98 inside the scanner, and 0.87/0.99 outside the scanner. In other words, Equation [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.

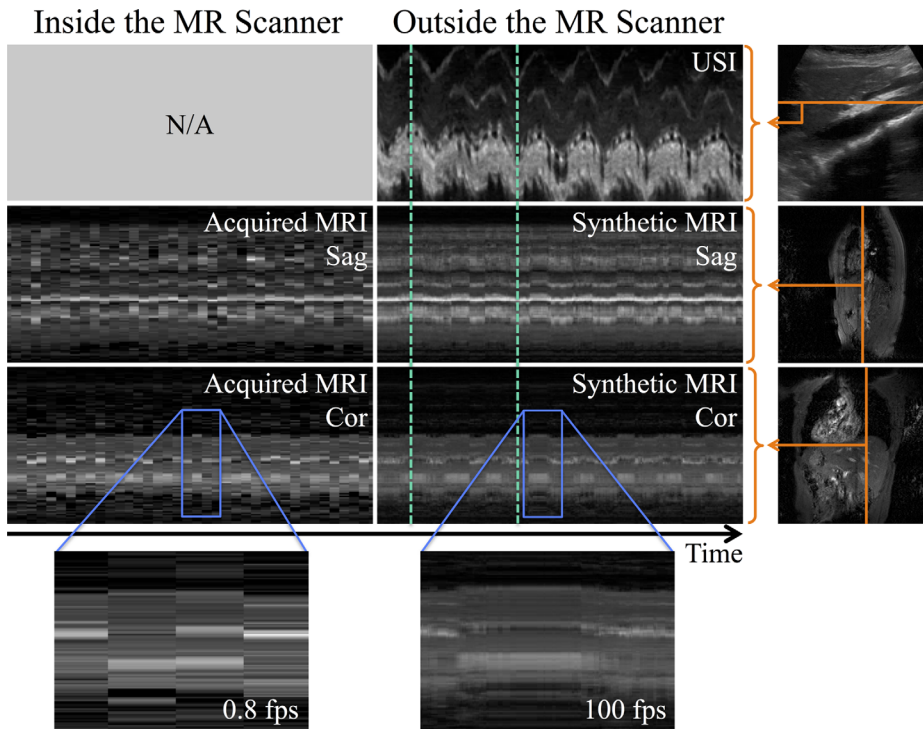Inside the MR Scanner   Outside the MR Scanner

FIG. 7. MRI data were acquired (left column), and then the subject was taken out of the MRI scan room. Based on the 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, whereas 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) compared with synthetic MRI (100 fps). Images in the rightmost column show the 1D locations selected for M-mode display.

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 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. Although there was nothing special about the ultrasound signal or the MR images obtained, interesting behaviors emerged when detecting correlations between the two. After a training period extending for a few breathing
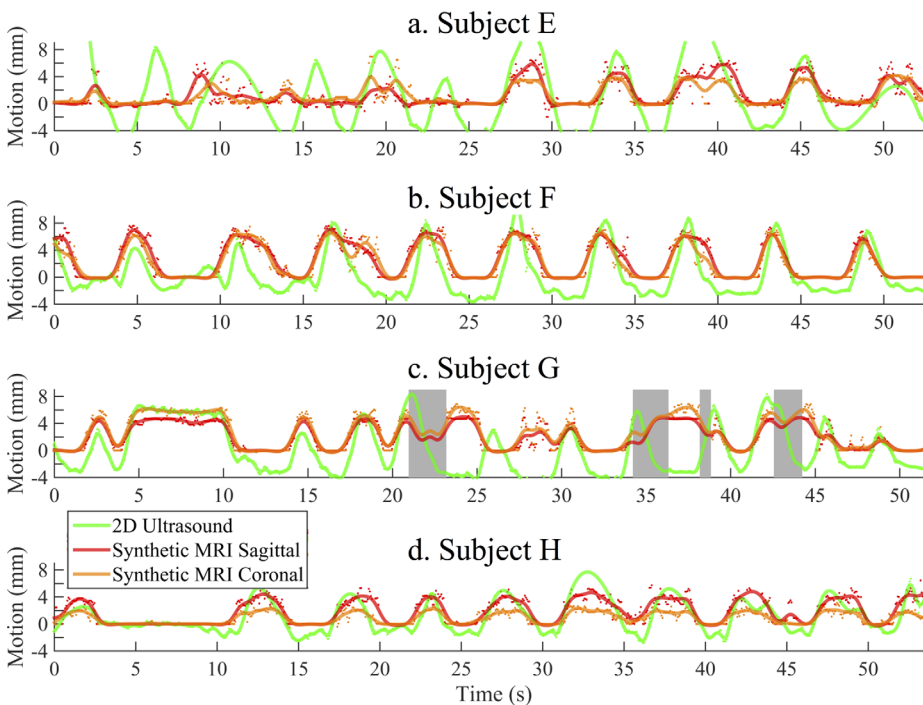

FIG. 8. The motion of structures as seen in synthetic MRI images (red and orange curves), generated from outside the MRI scan room, was compared with that obtained from time-matched images from a commercial ultrasound scanner (green curve). Dots represent measured locations and solid lines represent 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.

cycles, the US signal became, in essence, a surrogate for MR images (4,8). Although 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. 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 in this study, suggesting that smaller pixels might have given smaller errors. Furthermore, faster MRI could have helped to better capture motion. In real-time imaging, there is a direct relationship among frame rate, MRI scan time, and temporal resolution. In contrast, in our method the temporal resolution and frame rate were both improved, but not the 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 by increasing the frame rate alone. However, a more compact MRI acquisition with shorter scan time should in principle also help to 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. Although 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. Although 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 prelearned 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.

CONCLUSIONS

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 (eg, 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.

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;197:175–181.
2. Panych LP, Tokuda J. Real-time and interactive MRI. In: Jolesz F, ed. Intraoperative Imaging and Image-Guided Therapy. New York: Springer; 2014. p 193–210.
3. Schwartz BM, McDannold NJ. Ultrasound echoes as biometric navigators. Magn Reson Med 2013;69:1023–1033.
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. In Proceedings of the 21st Annual Meeting of ISMRM, Salt Lake City, Utah, USA, 2013. p. 478.
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. In Proceedings of the 22nd Annual Meeting of ISMRM, Milan, Italy, 2014. p. 681.
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. In Proceedings of the 23rd Annual Meeting of ISMRM, Toronto, Canada, 2015. p. 863.
7. Preiswerk F, Toews M, Hoge WS, Chiou J-y, Panych L, Wells W, III, Madore B. Hybrid utrasound and MRI acquisitions for high-speed imaging of respiratory organ motion. In: Navab N, Hornegger J, Wells W, Frangi A, eds. 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;58:4749–4761.
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. In Proceedings of the 21st Annual Meeting of ISMRM, Salt Lake City, Utah, USA, 2013. p. 1421.

10. Petrusca L, Cattin P, De Luca V, e al. Hybrid ultrasound/magnetic resonance simultaneous acquisition and image fusion for motion monitoring in the upper abdomen. Invest Radiol 2013;48:333–340.

11. Günther M, Feinberg DA. Ultrasound-guided MRI: preliminary results using a motion phantom. Magn Reson Med 2004;52:27–32.

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: 290–296.

13. Breast cancer—combining imaging techniques for quicker and gentler biopsies. 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. Accessed February 29, 2016.

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;60:146–157.

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;52:1547–1564.

16. Celicanin Z, Bieri O, Preiswerk F, Cattin P, Scheffler K, Santini F. Simultaneous acquisition of image and navigator slices using CAIPI-RINHA for 4D MRI. Magn Reson Med 2015;73:669–676.

17. Thomenius KE. Evolution of ultrasound beamformers. IEEE Ultrasonics Symp, San Antonio, Texas, 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;75:1273–1282.

19. Capon J. High resolution frequence-wavenumber spectrum analysis. Proc IEEE 1969;57:1408–1418.

20. Synnevag JF, Austeng A, Holm S. Adaptive beamforming applied to medical ultrasound imaging. IEEE Trans Ultrason Ferroelectr Freq Control 2007;54:1606–1613.

21. Madore B, Meral F. Reconstruction algorithm for improved ultrasound image quality. IEEE Trans Ultrason Ferroelectr Freq Control 2012;59:217–230.

22. Ellis MA, Viola F, Walker WF. Super-resolution image reconstruction using diffuse source models. Ultrasound Med Biol 2010;36:967–977.

23. Parzen E. On estimation of a probability density function and mode. Ann Math Statist 1962;33:1065–1076.

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;47:2600-2603.

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;18:712–721.

27. Lagarias JC, Reeds JA, Wright MH, Wright PE. Convergence properties of the Nelder-Mead simplex method in low dimensions. Siam J Optim 1998;9:112–147.

28. Kahneman D. Attention and effort. Englewood Cliffs, NJ: Prentice-Hall; 1973. p 246.

29. Kahneman D. Thinking, fast and slow. New York: 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;61:2527–2537.

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;25:129–131.

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;33:713–719.

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article

Video S1. A movie showing the results from inside the scanner (with subject C) has been made available as an example.

Link S1. Open-source code and sample data has been made available at https://github.com/fpreiswerk/OCMDemo.