Hybrid referenceless and multibaseline subtraction MR thermometry for monitoring thermal therapies in moving organs

William A. Grissom^{a)}

Department of Radiology, Stanford University, Stanford, California 94304 and Department of Electrical Engineering, Stanford University, Stanford, California 94304

Viola Rieke

Department of Radiology, Stanford University, Stanford, California 94304

Andrew B. Holbrook

Department of Bioengineering, Stanford University, Stanford, California 94304

Yoav Medan

InSightec Ltd., Haifa 32762, Israel

Michael Lustig

Department of Electrical Engineering and Computer Sciences, UC Berkeley, Berkeley, California 94704

Juan Santos

HeartVista, Inc., Los Altos, California 94304

Michael V. McConnell

Department of Cardiovascular Medicine, Stanford University, Stanford, California 94304

Kim Butts Pauly

Department of Radiology, Stanford University, Stanford, California 94304

(Received 12 February 2010; revised 5 July 2010; accepted for publication 13 July 2010; published 30 August 2010)

Purpose: Magnetic resonance thermometry using the proton resonance frequency (PRF) shift is a promising technique for guiding thermal ablation. For temperature monitoring in moving organs, such as the liver and the heart, problems with motion must be addressed. Multi-baseline subtraction techniques have been proposed, which use a library of baseline images covering the respiratory and cardiac cycle. However, main field shifts due to lung and diaphragm motion can cause large inaccuracies in multi-baseline subtraction. Referenceless thermometry methods based on polynomial phase regression are immune to motion and susceptibility shifts. While referenceless methods can accurately estimate temperature within the organ, in general, the background phase at organ/tissue interfaces requires large polynomial orders to fit, leading to increased danger that the heated region itself will be fitted by the polynomial and thermal dose will be underestimated. In this paper, a hybrid method for PRF thermometry in moving organs is presented that combines the strengths of referenceless and multi-baseline thermometry.

Methods: The hybrid image model assumes that three sources contribute to image phase during thermal treatment: Background anatomical phase, spatially smooth phase deviations, and focal, heat-induced phase shifts. The new model and temperature estimation algorithm were tested in the heart and liver of normal volunteers, in a moving phantom HIFU heating experiment, and in numerical simulations of thermal ablation. The results were compared to multi-baseline and referenceless methods alone.

Results: The hybrid method allows for *in vivo* temperature estimation in the liver and the heart with lower temperature uncertainty compared to multi-baseline and referenceless methods. The moving phantom HIFU experiment showed that the method accurately estimates temperature during motion in the presence of smooth main field shifts. Numerical simulations illustrated the method's sensitivity to algorithm parameters and hot spot features.

Conclusions: This new hybrid method for MR thermometry in moving organs combines the strengths of both multi-baseline subtraction and referenceless thermometry and overcomes their fundamental weaknesses. © 2010 American Association of Physicists in Medicine.

[DOI: [10.1118/1.3475943](https://doi.org/10.1118/1.3475943)]

Key words: interventional MRI, thermometry, thermal therapy

I. INTRODUCTION

MRI's ability to construct maps of *in vivo* temperature distribution allows MR image guidance of thermal therapy for benign and malignant diseases.¹ Although several MR parameters exhibit a temperature dependence, the proton resonance frequency (PRF) shift with temperature^{2,3} has been the method of choice for monitoring thermal ablation. Advantages of the PRF shift method are its linearity over the temperature range of interest, a constant of proportionality that is largely independent of tissue type (with the exception of adipose tissue),⁴ and temperature map acquisition with high spatial and temporal resolution.

In conventional baseline subtraction thermometry, phase images acquired prior to heating (i.e., baseline images) are subtracted from phase images acquired during heating. However, baseline subtraction's high sensitivity to motion is the most prevalent problem for temperature monitoring with the PRF method. Tissue motion causes misregistration of the images and changes in the magnetic susceptibility field,⁵ leading to large artifacts in the temperature maps. In recent years, various approaches have been developed to monitor thermal ablation in moving organs, such as the liver and the heart. Some of the proposed methods try to solve the problem by avoiding motion, e.g., using general anesthesia combined with breath holds⁶ or single lung ventilation⁷ for liver ablation, which make conventional baseline subtraction applicable. However, not all patients may qualify for breath holds or general anesthesia, and the methods are not suitable for cardiac ablation. Conventional respiratory gating has been used under general anesthesia⁸ and free breathing,⁹ but gating can fail when the respiratory cycle is irregular, leading to motion artifacts and errors in the temperature maps.¹⁰

Methods for temperature estimation in the presence of motion can be divided into two categories: methods based on a multibaseline strategy and methods based on a referenceless strategy. Multibaseline methods collect background phase information at various stages of the respiratory and/or cardiac cycle prior to heating, so that baseline data exist for all positions of the organ. The baseline images are stored in a library. The baseline subtraction is then performed by matching the image acquired during heating with the corresponding library baseline data in order to mitigate motion-induced misregistration. Selecting the corresponding baseline image can be performed by determining the organ's position with a navigator echo,¹¹ or based on a similarity criterion, such as nonsimilarity coefficients¹² or intercorrelation coefficients.^{13,14} The multibaseline methods require cyclic organ motion in order to acquire a full set of all possible baseline images and require additional setup time to assemble the baseline library. They are generally much more robust to motion than conventional baseline subtraction, but remain sensitive to susceptibility changes during therapy.

Referenceless methods estimate heating from a treatment image itself, without a baseline image. Operating under the assumption that the image phase directly surrounding the heating region would otherwise smoothly extend into the heating region, referenceless or self-referenced methods fit a

set of smooth, low-order polynomial functions to the surrounding phase¹⁵ or to a unit magnitude complex image with the same phase¹⁶ using a weighted least-squares fit. The polynomial is then extrapolated to locations within the heating region, and the result is used as a baseline for subtraction. Referenceless thermometry has been validated in a retrospective study in 33 patients undergoing uterine fibroid ablation¹⁷ and has been used for real-time monitoring of focused ultrasound ablation of the liver during free breathing.¹⁸ Recently, we introduced a new approach to the referenceless method that uses reweighted ℓ_1 regression,¹⁹ a form of robust regression, to obtain background phase coefficients. That method does not require prior knowledge of the heating region in order to obtain an accurate temperature map and does not require tracking the heated area during motion.²⁰ This approach makes the referenceless methods much less dependent on user interaction and appropriate region-of-interest choice as compared to the original (ℓ_2 -based) referenceless method.¹⁵ Although referenceless methods are immune to motion, they are sensitive to rapid anatomical phase variations, since these cannot be accurately expressed as a weighted sum of smooth functions. Such variations commonly exist at organ interfaces. Proximity to the lungs is especially problematic due to the large susceptibility difference between the organ and the air-filled lungs.

To address the drawbacks of multibaseline and referenceless methods, we introduce a new hybrid image model for PRF-shift thermometry and a penalized-likelihood algorithm for fitting MRI images to that model, which yield temperature change maps in the presence of motion. Compared to multibaseline methods, the new method is robust to susceptibility changes during therapy. Compared to referenceless methods, the hybrid method is robust to rapidly varying anatomical phase, for example, at organ interfaces. The new method is therefore more robust than previous thermometry techniques when applied to moving organs.

II. ALGORITHM

II.A. Treatment image model

PRF shift thermometry typically uses GRE images as input data (although spin-echo images have also been used²), and temperature is estimated from image phase. In the presence of therapy-induced temperature change, a hot spot appears in the phase, since proton resonance frequency decreases with increasing temperature. However, image phase can comprise several other sources in addition to the hot spot. The goal of temperature estimation is to separate phase shifts resulting from thermal therapy from other sources. To accomplish this, we introduce a thermal treatment image model. First, image phase contains anatomical features that are created by local field inhomogeneities and phase variations imposed by the transmit and receive coils. At a given point in the target organ's cycle of motion, these sources can generally be regarded as static over the course of thermal therapy, and are effectively captured by baseline images acquired prior to treatment. Other sources of phase variation are smooth main field shifts resulting from, e.g., respiration

and bowel filling, and can be modeled as a linear combination of smooth basis functions. Combining these sources yields the following model for image voxel j during treatment:

$$y_j = \left(\sum_{b=1}^{N_b} x_{b,j} w_b \right) e^{i(\{A\mathbf{c}\}_j + \theta_j)} + \varepsilon_j, \quad (1)$$

where the $\{x_{b,j}\}_{b=1}^{N_b}$ are complex images from the baseline library acquired prior to treatment (comprising $N_b \geq 1$ images), the w_b are the baseline image weights, A is a matrix of smooth (e.g., low-order polynomial) basis functions, $\mathbf{c}$ is a polynomial coefficient vector, θ is a temperature-induced phase shift, and ε is complex Gaussian noise.

The temperature-induced phase shift θ is related to a map of temperature change (in units of $\Delta^\circ\text{C}$) via multiplication with a constant c that is given by

$$c = -\frac{1}{\gamma \alpha \text{TE} B_0}, \quad (2)$$

where $\alpha = -0.01 \text{ ppm}/^\circ\text{C}$ is the PRF change coefficient for aqueous tissue, γ is the proton gyromagnetic ratio, B_0 is the main magnetic field strength, and TE is the echo time of the GRE sequence.

II.B. Iterative regularized temperature estimation

The image model in Eq. (1) can be fitted to the treatment image $\mathbf{y}$ by minimizing the following penalized-likelihood cost function with respect to the multibaseline weights $\mathbf{w}$, the smooth phase shift coefficients $\mathbf{c}$, and the temperature-induced phase shift θ ,

$$\Psi(\mathbf{w}, \mathbf{c}, \theta) = \frac{1}{2} \sum_{j=1}^{N_s} \left| y_j - \left(\sum_{b=1}^{N_b} x_{b,j} w_b \right) e^{i(\{A\mathbf{c}\}_j + \theta_j)} \right|^2 + \lambda \|\theta\|_0, \quad (3)$$

where the first term on the right-hand side is proportional to the negative log likelihood under the additive Gaussian noise model, $\|\theta\|_0$ equals the number of nonzero elements in θ , and λ is a regularization parameter that controls the sparsity of θ . Note that the baseline images are implicitly assumed to be noise-free, which is a practical requirement to prevent the estimated baseline phase from absorbing the temperature-induced phase shift. The temperature-induced phase shift θ is required to be non-negative in the case of heating and non-positive in the case of cooling. The elements of $\mathbf{w}$ are required to be non-negative and sum to one. An ℓ_0 penalty on θ is used to encourage sparse estimates because the hot spot is expected to be localized to a small region of the image. The ℓ_0 penalty is also effective in separating background polynomial phase from sparse heat-induced phase shifts because vectors with small ℓ_0 norms will not coincide with vectors in the range of the polynomial matrix A . Without this penalty, a nonsparse vector θ that contains background polynomial phase would yield the same cost function value as the true, sparse θ , and the method could converge to this erroneous result. This penalized-likelihood approach can be re-

garded as an extension of the regularized referenceless thermometry method in Ref. 21.

To minimize Eq. (3), we cycle through $\mathbf{w}$, $\mathbf{c}$, and θ , minimizing Eq. (3) with respect to one parameter while holding the other two constant. Detailed update equations for $\mathbf{w}$, $\mathbf{c}$, and θ are given in Appendix A. The baseline weight vector $\mathbf{w}$ is obtained by solving a quadratic program [Eq. (A1)]. Optimization transfer²² is used to ensure that the maximum likelihood term in Eq. (3) is decreased monotonically with respect to θ and $\mathbf{c}$. The ℓ_0 penalty is implemented using an iteratively reweighted least-squares strategy.²³ Since the regularized cost function is nonconvex in the temperature map θ and the polynomial coefficients $\mathbf{c}$, the proposed gradient-based algorithm will converge to a local minimum. Therefore, an initial estimate must be provided that is reasonably close to the global minimum. The initial $\mathbf{w}$ is obtained by solving a quadratic program using magnitude images, yielding a baseline image estimate. Then, the estimated baseline image's phase is subtracted from the treatment image and finite difference-based reweighted ℓ_1 referenceless thermometry¹⁹ is used to obtain initial $\mathbf{c}$ and θ . A detailed description of this initialization procedure is given in Appendix B. Prior to estimation, in order to reduce variations in the optimal choice of λ , we normalize the treatment image $\mathbf{y}$ and the baseline library images $\{x_{b,j}\}_{b=1}^{N_b}$ by the mean across baseline library images of the median image magnitudes.²²

The hybrid model and numerical methods can be easily extended to multicoil image reception, by jointly estimating baseline coefficients $\mathbf{w}$ across coils, independent smooth phase shift coefficients $\mathbf{c}$ for each coil, and a common heat phase shift θ . The cost function in this scenario is given by

$$\Psi(\mathbf{w}, \{\mathbf{c}\}_{k=1}^{N_c}, \theta) = \frac{1}{2} \sum_{j=1}^{N_s} \sum_{k=1}^{N_c} \left| y_{k,j} - \left(\sum_{b=1}^{N_b} x_{b,k,j} w_b \right) e^{i(\{A\mathbf{c}_k\}_j + \theta_j)} \right|^2 + \lambda \|\theta\|_0, \quad (4)$$

where N_c is the number of receive coils, $y_{k,j}$ is the signal in coil k at voxel j , $x_{b,k,j}$ is the signal in coil k 's baseline image b at voxel j , and $\mathbf{c}_k$ is the polynomial coefficient vector for coil k .

III. METHODS

The penalized-likelihood hybrid algorithm was implemented on a dual core 2.4 GHz laptop computer with 4 GB RAM (Apple Computer, Inc., Cupertino, California, USA) in MATLAB 7.4 (Mathworks, Inc., Natick, Massachusetts, USA) and was applied in the following experiments and simulations. To illustrate the algorithm's convergence properties and execution times, for one image in each data set the penalized-likelihood cost function value at each iteration was recorded, as well as the total time spent updating each model parameter.

III.A. Model validation experiments

The multibaseline, referenceless, and hybrid thermometry methods differ in how they model pretreatment image phase. To compare the accuracy of the their pretreatment phase models, the three methods were applied to cardiac and liver images acquired in healthy volunteers without heating. Short-axis free-breathing cardiac images (dimension 66×51 , $N_s=3366$) were acquired in a volunteer in real time,²⁴ using a spiral acquisition with four interleaves (TR=23 ms, TE=5 ms) on a 3T scanner (GE 3T Signa HDx, GE Healthcare, Waukesha, WI, USA). Images were reconstructed using a sliding window reconstruction in which, with the acquisition of each new interleaf, the fourth oldest one was discarded and the image was reconstructed with the new interleaf and the preceding three interleaves.²⁴ A sixth-order background polynomial basis set was used in the referenceless and hybrid methods, and the baseline libraries comprised 3.45 s of free-breathing, encompassing approximately four heartbeats ($N_b=150$ images). The hybrid algorithm was compared with the referenceless thermometry method described in Ref. 21 and with multibaseline subtraction performed using the hybrid algorithm with zeroth-order polynomial phase regression for field drift compensation.^{11,13} The methods were applied to eight images that were acquired 2.3 s after the library and covered approximately one-half of a cardiac cycle. Using masking, the referenceless method was applied separately in the septum and left ventricular (LV) free wall [the two temperature overlay regions in Fig. 1(b)], while the other two methods were applied over circular regions of interest (ROIs) containing the entire left ventricle [outlined on the gradient echo (GRE) images in Fig. 1(b)]. The sparsity regularization parameter λ was set very high ($\lambda=10^{10}$) in all methods to force the estimated temperature θ to zero everywhere, effectively deleting it from the image model. All methods were run for 25 iterations. For quantitative comparison of the image models, the standard deviation of the residual phase after model fitting (i.e., the phase that remains after subtraction of the estimated library image phase and the estimated polynomial phase from the treatment image phase) was calculated separately in the septum and LV free wall. The phase standard deviations were converted to temperature standard deviations using Eq. (2).

The three methods were also applied to sagittal liver images (dimension 337×112 , $N_s=37744$) acquired in a healthy volunteer during free breathing. Images were acquired on a 3T scanner (GE 3T Signa Excite, GE Healthcare, Waukesha, WI, USA) with a real-time reduced-field-of-view (rFOV) readout-segmented echo planar imaging (RS-EPI) sequence (TE=16.2 ms, TR=122 ms, three interleaves, 365 ms/image)¹⁸ using an eight-channel cardiac array. Out of the eight-channel array, $N_c=4$ of the coils that covered the FOV of the selected slice were used for temperature reconstruction. The baseline library for the multibaseline and hybrid methods comprised $N_b=10$ images acquired over one breathing cycle; a sixth-order polynomial basis set was used in the hybrid and referenceless methods. The methods were applied to eight images that were acquired directly after the baseline

library acquisition and covered approximately three breathing cycles. The referenceless method was run separately in two rectangular regions; one well inside the liver and one at the liver/rib interface [the temperature overlay regions in Fig. 2(c)], whereas the other two methods estimated temperature in the whole image. All algorithms used 25 iterations and $\lambda=10^{10}$. For quantitative comparison of the image models, the temperature standard deviations were calculated separately in the inner liver and liver/rib interface regions, in the same manner as the cardiac experiment. Prior to calculating temperature standard deviations, the residual phase from the individual receive coils was combined using an image-magnitude weighted mean.

The hybrid model was also compared to multibaseline and referenceless models as a function of baseline library size and polynomial order. To compare the hybrid and multibaseline models in the heart, 100 images from the heart data set were processed with library sizes ranging from 25 to 250 images. Temperature errors were calculated within the entire myocardium and were averaged across the processed images. To compare the hybrid and referenceless models in the heart, the same 100 images were processed for polynomial orders between 1 and 10. Temperature errors were calculated separately within the septum and LV free wall and were averaged across the processed images. To compare the hybrid and multibaseline models in the liver, 10 images from the liver data set were processed with library sizes ranging from 5 to 40 images; temperature errors were calculated over the entire liver, and averaged across the 10 processed images. To compare the hybrid and referenceless models in the liver, the same ten images were processed for polynomial orders between 1 and 10; temperature errors in the liver interior and liver/rib interface were averaged across the ten processed images. All algorithms used 25 iterations and $\lambda=10^{10}$.

III.B. Phantom heating experiment

A high intensity focused ultrasound (HIFU) phantom heating experiment was performed to validate the hybrid method and compare its performance to the multibaseline and referenceless methods during thermal treatment. Heating was applied to the center of a polyacrylamide gel phantom using an InSightec ExAblate 2000 HIFU system (InSightec Ltd., Haifa, Israel) (1000 element transducer, 50 s sonication duration, 550 kHz, 44 W acoustic power). Imaging was performed on a 3T scanner (GE 3T Signa HDx, GE Healthcare, Waukesha, WI, USA) with the same rFOV, RS-EPI sequence that was used in the volunteer liver experiment. Reception was performed using a custom coil built for liver imaging. During sonication, the phantom/transducer assembly was manually moved in and out of the scanner along the z-axis in a cyclical manner [as depicted in Figs. 4(a) and 4(b)]. Approximately 30 s into the experiment, a bottle phantom placed near the gel phantom but out of the imaging FOV was tipped over in order to induce a smooth main field change in the gel phantom. Images were reconstructed on a 184×96 grid ($N_s=17664$). Prior to sonication, a baseline library for the multibaseline and hybrid methods was acquired that

FIG. 1. Comparison of pretreatment image phase model accuracy in free-breathing, real-time short-axis cardiac images of a healthy volunteer. (a) Magnitude (top row) and phase images (bottom row) of the baseline library over approximately one heartbeat. Phase in the septum is relatively smooth, while a large, steep phase ramp is present in the LV free wall near the lung. (b) Comparison of pretreatment phase modeling errors in the ventricular myocardium. The hybrid model is more accurate throughout the myocardium than the other two methods. Standard deviations were calculated separately within the septum and LV free wall.

comprised $N_b=100$ images covering two cycles of motion. The hybrid and referenceless methods used a sixth-order polynomial basis set. All methods were run for 25 iterations, with $\lambda=10^{-2}$.

III.C. Heating simulations

Heating simulations were performed to investigate the proposed algorithm's sensitivity to the sparsity regularization parameter λ and sensitivity to hot spot location. A rectangu-

lar object was generated on a 64×64 grid ($N_s=4096$), with anatomical phase synthesized from randomly generated low frequency Fourier components, as shown in Fig. 5. The object was then shifted up and down within a range of motion similar to liver displacement during free breathing to form a multibaseline library with $N_b=50$ images [Fig. 5(b)]. A random spatial shift was then selected within the library's range of motion, and a treatment image was synthesized at this shift. A random polynomial phase shift (sixth-order) and a

FIG. 2. Comparison of modeling errors in sagittal liver images of a healthy volunteer. (a) (top) A library image (sum-of-squares across coils) and (bottom) image phase for a single coil. The phase is smooth in the liver center, but varies rapidly near the ribs. The highlighted region in (a) is shown across library images in (b) to illustrate the up-down motion of the liver. (c) shows residual temperature errors and standard deviations after estimation with the three methods in eight subsequently acquired images. The hybrid model is more accurate than the other two methods in both the inner liver and over the liver/rib interface region.

Gaussian hot spot (peak magnitude π , $\sigma=0.035$ voxels) were incorporated into the treatment image phase. Note that the hot spot is nearly indistinguishable from the background phase variation [Fig. 5(d), white arrow]. The penalized-likelihood hybrid algorithm was run for 25 iterations, and used $\lambda=10^{-3}$.

To investigate the algorithm's sensitivity to λ , the magnitude and size (σ) of the Gaussian hot spot were varied and the simulated data were repeatedly processed using the hybrid algorithm with λ values between 10^{-9} and 1 in integer powers-of-ten increments, and 25 iterations. The estimation

error of each θ result was recorded. In the hot spot magnitude simulation, the size was held fixed at $\sigma=0.035$ voxels, and the magnitude was varied between 0 and π . In the hot spot size simulation, the magnitude was held fixed at π radians, and the size was varied between $\sigma=0$ and 0.1 voxels. In the hot spot location simulation, a Gaussian hot spot (peak magnitude π , $\sigma=0.035$ voxels) was moved along a diagonal line from the lower right corner to the upper left corner of the liver (shown in Fig. 7). The resulting treatment images were then processed using the hybrid algorithm with 25 iterations and $\lambda=10^{-3}$.

TABLE I. Convergence of the iterative penalized-likelihood hybrid algorithm. For one image in each experimental and simulated data set, the number of iterations required to reach convergence (defined as a relative cost decrease between consecutive iterations of less than 0.1%) is shown, as well as the cost at convergence and 25 iterations relative to the initial estimate. Using this definition of convergence, for all data sets the algorithm converged within 20 iterations.

Data set	Iters. to convergence	Relative cost at convergence	Relative cost at 25 iters.
Cardiac	4	3.3×10^{-12}	3.3×10^{-12}
Liver	19	5×10^{-11}	5×10^{-11}
Phantom	3	0.77	0.77
Simulation	2	0.14	0.13

IV. RESULTS

IV.A. Algorithm convergence and execution time

Table I lists the number of iterations required for the algorithm to converge, when convergence is defined as a relative cost function decrease between consecutive iterations of less than 0.1%. In most data sets, the penalized-likelihood algorithm required less than five iterations; in all data sets, convergence was reached in less than 20. The liver data set required the largest number of iterations; this is likely due to its small library size which led to larger inaccuracies in the initial θ and c , and subsequently more significant updates of these parameters were required than for other data sets. Table I also lists the relative cost at convergence and at the 25th iteration. These figures can vary significantly, depending on the accuracy of the initial estimate. In all data sets, little improvement in cost was achieved by continuing iterations past convergence.

Table II lists execution times for each data set. Due to larger image sizes, the experimental data sets exhibited longer compute times than the simulation. The liver data set had the longest execution time due to the relatively large image size and multicoil acquisition. In general, execution time is dominated by w and θ updates, although the liver data set exhibited relatively short w update time due to its small library size, and relatively long θ and c update times due to multicoil acquisition. The cardiac data set exhibited the longest w update time due to its relatively large baseline library.

IV.B. Model validation experiments

Figure 1(a) shows library images (magnitude and phase) over the first heartbeat. Phase in the septum is smooth, while

a large, steep phase shift appears in the LV free wall next to the lung, due to large main field inhomogeneities there. Figure 1(b) shows residual temperature errors within the septum and LV free wall for the three methods. The images show that while the referenceless model achieves small errors in the septum, it is unable to fully capture the pretreatment phase in the LV free wall, resulting in temperature errors that can exceed 10 °C. In general, the multibaseline model achieves lower errors than referenceless, however, large errors remain near the lung and diaphragm, where the main field shifts dynamically with respiration. In contrast, the hybrid model is more accurate throughout the myocardium.

A comparison of modeling errors in the liver of a healthy volunteer is shown in Fig. 2. The liver is in the center of the reduced FOV sum-of-squares reconstructed magnitude images, with ribs to the left and right [Fig. 2(a), top]. The highlighted region is shown across the library images [Fig. 2(b), top]. Note the up and down motion. The single coil phase image in Fig. 2(a) (bottom) shows smooth phase in the liver center, but rapidly varying phase near the rib interfaces. Phase motion due to respiration is shown in Fig. 2(b) (bottom). The temperature maps in Fig. 2(c) show residual phase after model fitting to eight images after the library acquisition using referenceless, multibaseline, and hybrid techniques. The referenceless model performs well in the center of the liver, but poorly close to the ribs, where rapid anatomical phase variations cannot be regressed away by smooth polynomial functions. The multibaseline model is effective in removing rapid anatomical phase variations, but smoother errors remain due to respiration and motion of the organ in the main field. In contrast, the hybrid model reliably captures both effects in all images, leaving residual phase errors at or below the noise level.

TABLE II. Execution times of the iterative penalized-likelihood algorithm. The total execution time is shown for one image in each experimental and simulated data set, as well as the portion of that time spent updating each of the three model parameters. With the exception of the liver data set, updates to w and θ comprise the majority of the compute time.

Data set	Total execution time	Total w time	Total θ time	Total c time
Cardiac	26.9 s	25.4 s (94.3%)	1.4 s (5.2%)	0.11 s (0.39%)
Liver	42.2 s	2 s (4.7%)	34.6 s (82.0%)	4.8 s (11.4%)
Phantom	17.2 s	10.1 s (59.0%)	6.3 s (36.8%)	0.62 s (3.6%)
Simulation	5.1 s	3.0 s (59.6%)	1.88 s (26.9%)	0.13 s (2.6%)

FIG. 3. Modeling errors as a function of baseline library size [(a) and (c)] and polynomial order [(b) and (d)], in the heart [(a) and (b)] and liver [(c) and (d)]. Compared to the multibaseline and referenceless models, the hybrid model requires significantly smaller baseline libraries and polynomial orders to achieve a given level of temperature error.

Comparisons of modeling error as a function of baseline library size are shown in Fig. 3(a) (heart) and Fig. 3(c) (liver). In both organs, the hybrid model requires significantly fewer baseline images than the multibaseline model to achieve a given level of error. For example, in the heart, to reach an error level lower than the hybrid model's error for 75 baseline library images, the multibaseline model requires 175 images. In the liver, to reach an error level lower than the hybrid model's error for ten baseline library images, the multibaseline model requires 20 images. Comparisons of modeling errors as a function of polynomial order are shown in Fig. 3(b) (heart) and Fig. 3(d) (liver). Up to a large polynomial order of 10, the referenceless model was unable to achieve the same level of error as the hybrid model with polynomial order of 1, indicating that the hybrid model requires much smaller polynomial orders to reach a given level of temperature error.

IV.C. Phantom heating experiment

Figure 4(c) plots the temperatures estimated by the three methods in the center of the hot spot over the course of the phantom HIFU sonication. Because the phantom's pretreatment phase is smooth [shown in Fig. 4(e)], the referenceless method performs well, and the hybrid and referenceless temperature curves largely coincide over the course of heating and cooling. The temperature maps estimated by the two

methods [Fig. 4(d)] are also very similar. However, the multibaseline method fails at two points in time. The first error appears in the multibaseline method's temperature curve approximately 8 s into the sonication and was caused by phantom motion that was slightly out of the range covered by the baseline library (not shown). The second and more significant error appears when the water bottle was dropped, around 30 s into the sonication. Treatment image phase pre- and postfield disturbance is shown in Fig. 4(e), as well as the phase difference between the two. The phase difference is a smooth, nonconstant function of space. Because the multibaseline method was not capable of regressing away this smooth phase shift, its temperature estimate erroneously increased by approximately 5 °C after this point, and remained approximately this much larger than the other two methods' estimates throughout the remainder of the sonication. The temperature map estimated by the multibaseline method at 42 s [Fig. 4(d)] contains a smooth error caused by the phase shift.

IV.D. Heating simulations

Figure 5 shows that the hybrid algorithm estimates the simulated hot spot with a high degree of accuracy. The RMSE in this simulation was 1.4×10^{-3} radians, while the

FIG. 4. Phantom HIFU heating experiment. (a) Images acquired at approximately the largest positive and negative displacements in the motion cycle (b). (c) Temperatures measured in the hot spot center by the three methods, and also by the multibaseline method with phase drift compensation. (d) Temperature maps estimated by the three methods 42 s into the sonication. (e) Left images: treatment image phase before and after the phase disturbance; right image: the difference between the two phase images.

maximum error was 9.7×10^{-3} radians, both of which correspond to errors well below 1°C for the experimental liver imaging parameters of $\text{TE}=16.2$ ms at 3T.

Figure 6 illustrates the penalized-likelihood hybrid algorithm's sensitivity to the sparsity regularization parameter λ . Figure 6(a) shows that, as would be expected of an ℓ_0 penalty, the algorithm is entirely insensitive to changes in the hot spot magnitude, and $\lambda=10^{-4}$ is the best parameter for all magnitudes. Figure 6(b) shows that while the estimation error increases for hot spots larger than $\sigma=0.05$, the best value of λ is not strongly dependent on this parameter. Figure 6(c) plots the error across a finer grid of λ values for hot spot parameters $\sigma=0.035$ and magnitude π . The error is not highly sensitive to λ variations of about an order of magnitude around $\lambda=10^{-4}$.

Figure 7(b) plots estimation error versus hot spot displacement along a diagonal in the rectangular object. This plot and the images in Fig. 7(c) show that the error is largely insensitive to hot spot location within the object, although the error increases slightly when the hot spot reaches the object's edge due to reduced image magnitude there.

V. DISCUSSION

A hybrid thermometry method has been proposed that combines multibaseline subtraction thermometry and refer-

enceless thermometry. The central advantage of the hybrid method compared to multibaseline subtraction is its robustness to main field shifts and other sources of smooth phase errors between baseline and treatment images. The central advantage of the hybrid method over referenceless thermometry is its ability to estimate temperature in the presence of rapidly varying anatomical phase.

It was shown that the hybrid pretreatment phase model is more accurate than multibaseline subtraction and referenceless models in *in vivo* liver and the heart data. Temperature uncertainties measured in the septum and left ventricular free wall were below 3 and 5°C , respectively, while in the liver, the hybrid method was able to regress both rapidly varying anatomical background phase, and smooth background phase shifts. For a given level of temperature error, the hybrid method generally requires a smaller baseline library than multibaseline subtraction and smaller polynomial orders than referenceless. A smaller baseline library is desirable because it reduces the library acquisition time, while a smaller polynomial order is desirable because it reduces the likelihood that the hot spot will bias the polynomial fit and cause temperature underestimation.¹⁹ The model and algorithms naturally extend to multicoil image reception, as was demonstrated in the *in vivo* liver study. The method was also compared to multibaseline and referenceless thermometry in

FIG. 5. Simulation of hybrid referenceless thermometry with a thermal treatment-induced phase shift. A treatment image (a) was generated within the range of 50 library positions [(b) with library image numbers indicated on the zeroth and 45th profile]. To the rapidly varying background anatomical phase (c), a Gaussian hot spot was added [(d) white arrow]. The hybrid method estimates an accurate image (e). The ℓ_0 penalty forces most of the estimated temperature map to zero, and leaving some residual phase (f) in the object that is not regressed by any parameter. (g) The Gaussian hot spot is estimated with a high degree of accuracy and small error (h) throughout the FOV.

a phantom HIFU sonication, where it was shown that it was robust to a sudden smooth main field change that occurred midway through the sonication.

To a degree, the hybrid method also inherits some of the limitations of multibaseline subtraction and referenceless thermometry methods. Like multibaseline subtraction, when motion occurs in a direction not covered by the baseline library, the method will become inaccurate. However, in practice, if navigator or registration information is available to track this motion, then the hybrid method may allow a previously acquired multibaseline library to be reused in the new orientation, since smooth phase shifts between the orientations would be robustly removed by the method's polynomial regression component. The method should also tolerate larger excursions from the motion range covered by the baseline library than the multibaseline method, since the referenceless component will partially compensate resulting phase errors. These will be topics of future investigation.

Like referenceless thermometry, the hybrid model and method assume that the hot spot is localized to a small region of the treated organ. As a hot spot becomes large and diffuse, both referenceless and hybrid thermometry can become inaccurate, since the hot spot becomes increasingly well modeled by the polynomial basis.¹⁹ The result in this case is that heat will be underestimated. The converse is also true: for a fixed hot spot size, the higher the polynomial order, the higher the risk that the basis will fit the hot spot. However, in this work it was demonstrated experimentally that the hybrid method requires significantly lower polynomial orders than referenceless thermometry to achieve a given level of accuracy. Furthermore, the hybrid method fits the polynomials to larger image regions than referenceless. Therefore, the likelihood that a hot spot will be fitted by the polynomial is reduced in hybrid thermometry, compared to referenceless methods.

The described penalized-likelihood hybrid algorithm requires the user to select the number of iterations, a smooth basis set (contained in the matrix A), and the sparsity regularization parameter λ . While a fixed number of iterations was used here (25 in all experiments and simulations), it was demonstrated that one could effectively detect convergence and stop iterations when the relative difference between consecutive cost function values falls below a certain threshold. Guided by previous experience with referenceless thermometry in the liver,¹⁵ a sixth-order polynomial basis set was selected for A in all experiments and simulations, and no attempt was made to optimize this parameter. However, as the purpose of this article was to validate and compare the hybrid method to multibaseline subtraction and referenceless methods, rather than to achieve minimal temperature error for a given treatment scenario, the choice of A was not as critical as it would be in a clinical application. The selection of A has been treated previously in referenceless thermometry;^{15,16} the same approaches could be used with the hybrid method with little modification. For example, as in Ref. 15, one could mask out the target treatment region in a pretreatment image not in the baseline library, apply the hybrid method with a very large λ value to force the estimated temperature-induced phase shift map to zero, and select the polynomial order that minimizes the residual phase error in the target region. This could be repeated with several library images using a leave-one-out strategy, and the polynomial order with the best average error would be selected. Numerical simulations demonstrated that the best choice of the sparsity regularization parameter λ in the penalized-likelihood hybrid algorithm is not a strong function of hot spot magnitude or size. Furthermore, the error is not highly sensitive to variations in λ over about one order of magnitude. In this work, λ was set empirically in scenarios containing heat, by trying several logarithmically distributed values in integer power-of-ten increments, and choosing the value that produced the most qualitatively pleasing result. In practice, one could set λ automatically by adding a synthetic hot spot to an image not in the baseline library in the target treatment location prior to treatment, testing a handful logarithmically distributed λ values, and selecting the one that

FIG. 6. Sparsity regularization parameter (λ) sensitivity simulations. [(a) and (b)]. The best value of λ , as measured by RMSE between the true and estimated hot spots, is largely independent of hot spot magnitude and size. The white lines indicate the minimum RMSE across tested λ values. (c) RMSE versus λ for a hot spot with peak magnitude π and $\sigma=0.035$ voxels. The optimal value of λ is about 10^{-4} ; the estimation error is not highly sensitive to λ variations of about an order of magnitude around this point.

minimizes the estimation error. Because the penalized-likelihood cost function [Eq. (3)] is nonconvex in $\mathbf{c}$ and $\boldsymbol{\theta}$, the algorithm also requires initial estimates that are reasonably close to the global minimum. Although the sensitivity of the algorithm was not studied in detail here, practical experience has shown that the method is robust to deviations in the initial polynomial and $\boldsymbol{\theta}$ up to about π radians. Furthermore, no scenario has yet been encountered in which the initialization strategy described in Appendix B failed to provide suitable initial estimates.

Note that additional regularization terms could be easily appended to Eq. (3) to incorporate other *a priori* information about $\boldsymbol{\theta}$ or $\mathbf{c}$. One notable example would be a quadratic roughness penalty²² to penalize rapid variations in $\boldsymbol{\theta}$, which would reflect one's expectation that the heated voxels will be largely contiguous with each other, and that the phase change map will not contain sharp edges in most treatment sce-

narios. This penalty may also improve the accuracy of estimates in regions of low image magnitude, or where temperature-induced changes in T_1 and T_2 reduce image magnitude, and hence degrade phase SNR inside a hot spot. The hybrid image model could also be extended to fat and water-separated acquisitions,²⁵ where common baseline weights and polynomial would be determined using both fat and water images, and the temperature-induced phase shift would be estimated from the water image only. This could improve the method's accuracy, since the polynomial fit to the fat image would not be biased by the hot spot.

The reported compute times for the MATLAB hybrid algorithm implementation were between one and two orders of magnitude too large for real-time application, depending on problem size. Conversion to C should dramatically reduce these times, and the algorithm itself could be sped up in several ways. As was reported, generally the most costly operations at each iteration are updates to the baseline weight vector $\mathbf{w}$ and to the temperature-induced phase shift map $\boldsymbol{\theta}$. Because elements of $\boldsymbol{\theta}$ are not coupled to each other in the penalized-likelihood cost function, updates to this parameter could be parallelized across space, for example, by dividing the image FOV into blocks and performing $\boldsymbol{\theta}$ updates in parallel for each block. Updates to $\mathbf{w}$ require the solution of a quadratic program; one could reduce the computational burden of this step by reducing the size of the quadratic program. For example, navigator information or knowledge of the respiratory phase of the treatment image could be exploited to dramatically reduce the size of the baseline library to contain only images near the current phase. Finally, in multicoil acquisitions, polynomial coefficient vector ($\mathbf{c}$) updates could be performed in parallel across coils.

The hybrid method requires no hot spot tracking, gating, or navigator acquisitions to produce accurate temperature estimates. Although in practice some form tracking is likely to be necessary to steer focused ultrasound beams in moving organs, independence of the temperature estimation algorithm from tracking is a desirable feature. For example, should tracking fail, heat may be deposited in an incorrect location, but the hybrid method would still accurately depict the true heating location if it were contained in the imaging volume. The temperature maps estimated by the hybrid

FIG. 7. Simulation of sensitivity to hot spot location. (a) The unshifted hot spot and the two hot spots with largest displacement, superimposed. The dashed line indicates the object's boundary. (b) Estimation error versus displacement. The error is relatively flat until the hot spot starts to leave the object. (c) Error images of the unshifted spot and two positive displacements. The +10.8 voxel displacement error image is of similar magnitude to the unshifted error, while the +20.3 voxel image contains larger errors near the edge of the object.

method could be used for tracking or to validate tracking approaches. On the other hand, if reliable hot spot tracking is available, it could be exploited to improve the robustness of the hybrid algorithm to large hot spot sizes and to decrease the method's sensitivity to the choice of λ . This would be done by excluding spatial locations within the hot spot from the ℓ_0 penalty, in a manner analogous to masked reweighted ℓ_1 referenceless thermometry.¹⁹

VI. CONCLUSIONS

A new hybrid image model and penalized-likelihood algorithm for temperature estimation from PRF-shift MR thermometry data have been introduced and validated. It was demonstrated that the hybrid image model accurately captures sources of PRF-shift image phase in the liver and left ventricular myocardium *in vivo*. It was also demonstrated in a phantom HIFU experiment and simulations that the method accurately estimates temperature rises.

ACKNOWLEDGMENTS

This work was supported by NIH Grant Nos. R21 EB007715 and CA121163.

APPENDIX A: PARAMETER UPDATE EQUATIONS

Because there are no analytical solutions for the minimizers of Eq. (3), we use a block alternating minimization approach in which we cycle through $\mathbf{w}$, $\mathbf{c}$, and $\boldsymbol{\theta}$, minimizing Eq. (3) with respect to one of the three parameters while holding the other two constant. For fixed $\mathbf{c}$ and $\boldsymbol{\theta}$, the optimal $\mathbf{w}$ is obtained by solving the following quadratic program:

$$\begin{aligned} & \text{minimize} \quad \frac{1}{2} \sum_{j=1}^{N_s} \left| y_j e^{-i(\{\mathbf{A}\mathbf{c}\}_j + \theta_j)} - \sum_{b=1}^{N_b} x_{b,j} w_b \right|^2 \\ & \text{subject to} \quad \sum_{b=1}^{N_b} w_b = 1 \\ & \quad \quad \quad \mathbf{w} \geq 0. \end{aligned} \quad (\text{A1})$$

For fixed $\mathbf{w}$ and $\mathbf{c}$, the problem of minimizing Eq. (3) with respect to $\boldsymbol{\theta}$ is nontrivial due to the nonlinearity of the complex exponential and the ℓ_0 penalty on $\boldsymbol{\theta}$. To implement the ℓ_0 penalty, an iteratively reweighted least-squares strategy is used,²³ in which a vector of weights $\mathbf{v}(\boldsymbol{\theta})$ is determined periodically as

$$v_j(\theta_j) = (\theta_j^2 + \epsilon)^{-1}, \quad (\text{A2})$$

where ϵ is a damping parameter that is initialized to $\epsilon = 10^{-4}$ at the beginning of each $\boldsymbol{\theta}$ update and is decreased via division by 10 at each $\boldsymbol{\theta}$ update subiteration. Then the ℓ_0 norm in Eq. (3) is replaced by

$$\frac{1}{2} \boldsymbol{\theta}' \text{diag}\{\mathbf{v}\} \boldsymbol{\theta}, \quad (\text{A3})$$

which is quadratic and differentiable in $\boldsymbol{\theta}$. To ensure that the first term in Eq. (3) is decreased monotonically, an optimization transfer approach is used in which a majorizing parabola surrogate for that term is minimized. Following an analogous derivation to Ref. 22 leads to the following iteration that is guaranteed to minimize Eq. (3) (after the substitution [Eq. (A3)] is made):

tion that is guaranteed to minimize Eq. (3) (after the substitution [Eq. (A3)] is made):

$$\boldsymbol{\theta}^{(n+1)} = \left\{ \boldsymbol{\theta}^{(n)} - \text{diag} \left\{ \frac{1}{d_j^{(n)} + \lambda v_j} \right\} \nabla \Psi(\boldsymbol{\theta}^{(n)}) \right\}_+, \quad (\text{A4})$$

where $\{x\}_+ \triangleq \max\{x, 0\}$ enforces non-negativity of the temperature shift map (assuming heating rather than cooling); defining $\phi_j \triangleq \{\mathbf{A}\mathbf{c}\}_j + \theta_j + \angle(\sum_{b=1}^{N_b} x_{b,j} w_b) - \angle y_j$, d_j is given by

$$d_j^{(n)} = \left| y_j \left(\sum_{b=1}^{N_b} x_{b,j} w_b \right) \right| \frac{\sin(\{\phi_j^{(n)}\}_{2\pi})}{\{\phi_j^{(n)}\}_{2\pi}}. \quad (\text{A5})$$

Here, $\{\phi_j^{(n)}\}_{2\pi} = |\phi_j^{(n)} - l2\pi|$, where l is the integer that satisfies $|\phi_j^{(n)} - l2\pi| \leq \pi$.

A similar derivation yields an update for $\mathbf{c}$ when $\mathbf{w}$ and $\boldsymbol{\theta}$ are fixed,

$$\mathbf{c}^{(n+1)} = \mathbf{c}^{(n)} - (\mathbf{H}^{(n)})^{-1} \mathbf{A}' \nabla \psi(\mathbf{c}^{(n)}), \quad (\text{A6})$$

where

$$\mathbf{H}^{(n)} = \mathbf{A}' \text{diag} \left\{ \left| y_j \left(\sum_{b=1}^{N_b} x_{b,j} w_b \right) \right| \frac{\sin(\{\phi_j^{(n)}\}_{2\pi})}{\{\phi_j^{(n)}\}_{2\pi}} \right\} \mathbf{A}. \quad (\text{A7})$$

APPENDIX B: ALGORITHM INITIALIZATION

Estimating temperature using the penalized-likelihood cost function in Eq. (3) is complicated by the nonconvexity of the function with respect to $\mathbf{c}$ and $\boldsymbol{\theta}$. Thus, the penalized-likelihood hybrid algorithm must be initialized with $\mathbf{c}$ and $\boldsymbol{\theta}$ reasonably close to the optimal values. To achieve this, in this article the penalized-likelihood hybrid algorithm is initialized using a baseline fit performed on magnitude images only, so that the initial baseline estimate is robust to smooth main field shifts and the hot spot. The initial baseline estimate is obtained by solving the following quadratic program:

$$\begin{aligned} & \text{minimize} \quad \frac{1}{2} \sum_{j=1}^{N_s} \left| |y_j| - \sum_{b=1}^{N_b} |x_{b,j}| w_b \right|^2 \\ & \text{subject to} \quad \sum_{b=1}^{N_b} w_b = 1 \\ & \quad \quad \quad \mathbf{w} \geq 0. \end{aligned} \quad (\text{B1})$$

This is nearly the same expression as Eq. (A1); the differences are the absolute value operators around the treatment image signal y_j and the baseline library images $x_{b,j}$. Once the initial weights $\mathbf{w}$ are obtained by solving Eq. (B1), the residual phase between the estimated baseline image and the treatment image is calculated as

$$\phi_j = \angle y_j - \angle \left(\sum_{b=1}^{N_b} x_{b,j} w_b \right), \quad (\text{B2})$$

and finite difference-based reweighted ℓ_1 referenceless thermometry¹⁹ is applied to ϕ to obtain initial $\mathbf{c}$ and $\boldsymbol{\theta}$. That algorithm finds the vector $\mathbf{c}$ that approximately minimizes the cost function,

$$\Psi(c) = \|\tilde{\Phi}^x - \tilde{A}^x c\|_0 + \|\tilde{\Phi}^y - \tilde{A}^y c\|_0, \quad (\text{B3})$$

where $\tilde{\Phi}^x$ and $\tilde{\Phi}^y$ are the wrapped finite-differenced phase images, with differences taken in the x and y dimensions, respectively, and $\tilde{A}^x$ and $\tilde{A}^y$ are matrices containing finite-differenced polynomial basis functions. Generally very few (zero in this article) reweights are required to obtain good initial c and θ for the penalized-likelihood algorithm. In the multicoil case, reweighted ℓ_1 thermometry is performed separately for each coil, and the individual θ estimates are combined using an image magnitude-weighted mean. In this work, the penalized-likelihood multibaseline subtraction implementation is initialized using this technique with a zeroth-order polynomial fit; the penalized-likelihood referenceless implementation²¹ is initialized using the reweighted ℓ_1 method with no prior baseline phase subtraction.

^{a)}Electronic mail: wgrissom@gmail.com

¹V. Rieke and K. Butts Pauly, "MR thermometry," *J. Magn. Reson. Imaging* **27**, 376–390 (2008).

²Y. Ishihara, A. Calderon, H. Watanabe, K. Okamoto, Y. Suzuki, K. Kuroda, and Y. Suzuki, "A precise and fast temperature mapping using water proton chemical shift," *Magn. Reson. Med.* **34**, 814–823 (1995).

³J. De Poorter, C. De Wagter, Y. De Deene, C. Thomsen, F. Stahlberg, and E. Achten, "Noninvasive MRI thermometry with the proton resonance frequency (PRF) method: In vivo results in human muscle," *Magn. Reson. Med.* **33**, 74–81 (1995).

⁴R. D. Peters, R. S. Hinks, and R. M. Henkelman, "Ex vivo tissue-type independence in proton-resonance frequency shift MR thermometry," *Magn. Reson. Med.* **40**, 454–459 (1998).

⁵R. Salomir, B. D. de Senneville, and C. T. Moonen, "A fast calculation method for magnetic field inhomogeneity due to an arbitrary distribution of bulk susceptibility," *Concepts Magn. Reson., Part B* **19B**, 26–34 (2003).

⁶W. M. Gedroyc, "New clinical applications of magnetic resonance-guided focused ultrasound," *Top. Magn. Reson. Imaging* **17**, 189–194 (2006).

⁷R. O. Illing, J. E. Kennedy, F. Wu, G. R. ter Haar, A. S. Protheroe, P. J. Friend, F. V. Gleeson, D. W. Cranston, R. R. Phillips, and M. R. Middleton, "The safety and feasibility of extracorporeal high-intensity focused ultrasound (HIFU) for the treatment of liver and kidney tumours in a western population," *Br. J. Cancer* **93**, 890–895 (2005).

⁸S. Morikawa, T. Inubushi, Y. Kurumi, S. Naka, K. Sato, K. Demura, T. Tani, and H. A. Haque, "Feasibility of respiratory triggering for MR-guided microwave ablation of liver tumors under general anesthesia," *Cardiovasc. Intervent Radiol.* **27**, 370–373 (2004).

⁹M. Lepetit-Coiffé, B. Quesson, O. Seror, E. Dumont, B. Le Bail, C. T. W. Moonen, and H. Trillaud, "Real-time monitoring of radiofrequency ablation of rabbit liver by respiratory-gated quantitative temperature MRI," *J. Magn. Reson. Imaging* **24**, 152–159 (2006).

¹⁰C. Weidensteiner, N. Keroui, B. Quesson, B. D. de Senneville, H. Trillaud, and C. T. W. Moonen, "Stability of real-time MR temperature mapping in healthy and diseased human liver," *J. Magn. Reson. Imaging* **19**, 438–446 (2004).

¹¹K. K. Vigen, B. L. Daniel, J. M. Pauly, and K. Butts, "Triggered, navigated, multi-baseline method for proton resonance frequency temperature mapping with respiratory motion," *Magn. Reson. Med.* **50**, 1003–1010 (2003).

¹²A. V. Shmatukha and C. J. G. Bakker, "Correction of proton resonance frequency shift temperature maps for magnetic field disturbances caused by breathing," *Phys. Med. Biol.* **51**, 4689–4705 (2006).

¹³B. D. de Senneville, C. Mougenot, and C. T. W. Moonen, "Real-time adaptive methods for treatment of mobile organs by MRI-controlled high-intensity focused ultrasound," *Magn. Reson. Med.* **57**, 319–330 (2007).

¹⁴S. Roujol, M. Ries, B. Quesson, C. Moonen, and B. D. de Senneville, "Real-time MR-thermometry and dosimetry for interventional guidance on abdominal organs," *Magn. Reson. Med.* **63**, 1080–1087 (2010).

¹⁵V. Rieke, K. K. Vigen, G. Sommer, B. L. Daniel, J. M. Pauly, and K. Butts, "Referenceless PRF shift thermometry," *Magn. Reson. Med.* **51**, 1223–1231 (2004).

¹⁶K. Kuroda, D. Kokuryo, E. Kumamoto, K. Suzuki, Y. Matsuoka, and B. Keserci, "Optimization of self-reference thermometry using complex field estimation," *Magn. Reson. Med.* **56**, 835–843 (2006).

¹⁷N. McDannold, C. Tempany, F. Jolesz, and K. Hynnen, "Evaluation of referenceless thermometry in MRI-guided focused ultrasound surgery of uterine fibroids," *J. Magn. Reson. Imaging* **28**, 1026–1032 (2008).

¹⁸A. B. Holbrook, J. M. Santos, E. Kaye, V. Rieke, and K. B. Pauly, "Real-time MR thermometry for monitoring HIFU ablations of the liver," *Magn. Reson. Med.* **63**, 365–373 (2010).

¹⁹W. A. Grissom, M. Lustig, A. B. Holbrook, V. Rieke, J. M. Pauly, and K. B. Pauly, "Reweighted ℓ_1 referenceless PRF shift thermometry," *Magn. Reson. Med.* (in press).

²⁰D. Kokuryo, T. Kaihara, E. Kumamoto, S. Fujii, and K. Kuroda, "Method for target tracking in focused ultrasound surgery of liver using magnetic resonance filtered venography," *Conf Proc IEEE Eng Med Biol Soc*, 2007, pp. 2614–2617.

²¹W. A. Grissom, K. B. Pauly, M. Lustig, V. Rieke, J. M. Pauly, and N. McDannold, "Regularized referenceless temperature estimation in PRF-shift MR thermometry," *Biomedical Imaging: Macro to Nano*, 2009. 6th IEEE International Symposium on Biomedical Imaging, 2009, pp. 1235–1238.

²²A. Funai, J. Fessler, D. Yeo, V. Olafsson, and D. Noll, "Regularized field map estimation in MRI," *IEEE Trans. Med. Imaging* **27**, 1484–1494 (2008).

²³R. Chartrand and W. Yin, "Iteratively reweighted algorithms for compressive sensing," *Proc. Int. Conf. Acoustics, Speech, Signal Processing (ICASSP)*, 2008, pp. 3869–3872.

²⁴K. S. Nayak, C. H. Cunningham, J. M. Santos, and J. M. Pauly, "Real-time cardiac MRI at 3 tesla," *Magn. Reson. Med.* **51**, 655–660 (2004).

²⁵B. J. Soher, C. Wyatt, S. B. Reeder, and J. R. MacFall, "Noninvasive temperature mapping with MRI using chemical shift water-fat separation," *Magn. Reson. Med.* **63**, 1238–1246 (2010).