Biometric Navigation with Ultrasound

Citation

Schwartz, Benjamin Matthew. 2012. Biometric Navigation with Ultrasound. Doctoral dissertation, Harvard University.

Permanent link

http://nrs.harvard.edu/urn-3:HUL.InstRepos:10536969

Terms of Use

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. Submit a story .

Accessibility

Dr. Nathan J. McDannold   Benjamin M. Schwartz

Biometric Navigation with Ultrasound

Abstract

We have designed and demonstrated a new class of medical navigation methods that use the fingerprint-like biometrically distinct ultrasound echo patterns produced by different locations in tissue. As an example of this new biometric navigation approach, we have constructed and tested a system that uses ultrasound data to achieve prospective motion compensation in MRI, especially for respiratory motion during interventional MRI procedures in moving organs such as the liver. The ultrasound measurements are collated with geometrical information from MRI during a training stage to form a mapping table that relates ultrasound measurements to positions. During prospective correction, the system makes frequent ultrasound measurements and uses the map to determine the corresponding position.

Results in motorized linear motion phantoms and freely breathing animals indicate that the system performs well. Apparent motion is reduced by up to 97.8%, and motion artifacts are reduced or eliminated in 2D Spoiled Gradient-Echo images. The motion compensation is sufficient to permit MRI thermometry of focused ultrasound heating during respiratory-like motion, with results similar to those obtained in the absence of motion. This new technique may have applications for MRI thermometry and other dynamic imaging in the abdomen during free breathing.

We have also extended this technique to situations in which external position information during training is unavailable or incomplete, by extending the concept of Simultaneous Localization and Mapping to include determining the topology of a dense motion path through a gaussian random field.

Dr. Nathan J. McDannold   Benjamin M. Schwartz

In the course of these investigations, we have also developed modified forms of referenceless MRI thermometry and Kalman filtering, specially adapted to optimize accuracy under our experimental conditions.

Contents

Listing of figures

ix

For my parents, Laura and Josh.

Acknowledgments

In the beginning, there was nothing.

Nigh Omnipotent – Alex Lee Martinez

1

Background

SINCE TIME IMMEMORIAL, doctors and patients alike have dreamed of treating disease from outside the body, to reproduce the invaluable curative power of surgery without the dangers of surgical invasion. Today non-invasive treatments are an essential part of the medical system, and some, such as x-ray beam therapy [11, 49] and ultrasound lithotripsy [12], have even supplanted surgical techniques that were previously the standard of care. However, countless invasive surgical procedures are still performed

every day, over 125,000 in the United States alone.[26] Every invasive treatment is an opportunity to advance the state of medicine, if the need for surgical invasion can be removed.

To expand the range of treatments that can be performed non-invasively, one of the most promising avenues is Focused Ultrasound Surgery. [51, 89] These techniques use high-power ultrasound, often over a kilowatt of acoustic power, that enters the body distributed over a large area of skin surface. The ultrasound wavefront is arranged so as to focus down to a small target volume where treatment is desired. At the target location, ultrasound intensities can be as high as 100 watts/mm^2 [79], but the power flux is confined to a volume set by the ultrasound wavelength. At typical treatment frequencies (150 kHz to 1.5 MHz), the wavelength is 1–10 mm, allowing highly specific targeting. Typical focused ultrasound surgery systems operate at a single fixed frequency set by the mechanical resonance of the piezolectric transducer that produces the wave.

At high intensities, focused ultrasound can have several different relevant bioeffects. Perhaps the simplest effect is hyperthermia. As human tissue is an imperfect conductor of ultrasound, a portion of the power in the ultrasound wave is absorbed and converted to heat at each location along the ultrasound beam path. Thus, these systems can produce highly localized heating at the focus, with localization limited by the ultrasound spot size and the tissue's heat diffusion characteristics. When the heat is maintained at a single location for sufficient duration, the tissue dies due to thermally induced coagulation and subsequent apoptosis or necrosis [88]. The amount of time required depends on the temperature, as higher temperatures kill

2

more quickly [73].

At extremely high intensities, nonlinear acoustic effects become biologically important. Most crucially, when the peak negative pressure exceeds a threshold (dependent on frequency and tissue characteristics), the stretching force is sufficient to separate the molecules of the medium and produce a vacuum, or cavity, which quickly fills with gas [30]. Alternatively, this may be regarded thermodynamically as a kind of spontaneous boiling due to the low pressure. This phenomenon is called *cavitation*, and it can be highly destructive to tissue [79]. Cavitation has been used as a component of some focused ultrasound treatments in humans [43].

In recent years, research in focused ultrasound treatment has included methods that combine focused ultrasound with one or more injectable agents. For example, a burgeoning experimental field combines focused ultrasound with microbubble contrast agents, which are injectable gas bubbles that circulate in the blood and normally serve to improve the diagnostic capabilities of ultrasound imaging (especially in the heart). In the extreme environment of a high-power ultrasound focus, microbubbles have been observed to produce clinically relevant bioeffects, including tissue destruction [59] and blood-brain barrier disruption [56]. In addition to microbubbles, other injectables such as drug-laden nanoemulsions been investigated in conjunction with focused ultrasound.[67] These combinations have been the subject of many animal trials, but have not yet been tested in humans.

Currently, Focused Ultrasound systems have regulatory approval from the Food and Drug Administration in the United States for treatment of uterine fibroids.[25] These treatments have achieved patient outcomes and safety

competitive with the current standard of care, an invasive embolization procedure.[80] Pre-approval human trials have been conducted for treatment of *glioblastoma multiforme* (a deadly cancer of the brain) [58], chronic neuropathic pain [53], essential tremor [63], prostate cancer [85], bone pain due to metastasis [29], breast cancer [38], and tumors in the liver [40, 60, 92].

In traditional invasive surgery, treatment guidance and monitoring is normally performed by the surgeon's own hands and eyes. When performing non-invasive treatments, more care is required to ensure that treatment occurs in the desired locations, and to the desired extent, without over-or under-treatment. In the case of focused ultrasound surgery, guidance and monitoring serve to ensure that the focus point is in the desired location, that its intensity is sufficient to achieve the desired biological effects, and that nearby areas that should not be affected remain untouched. To meet these goals, implementors typically employ either ultrasound or MRI imaging. For guidance, the structures of interest must be visible in the images, and the relative geometry of the image volume and the ultrasound transducer must be known. For monitoring, the effects of treatment must be visible in the image.

Ultrasound imaging offers the advantage of lower cost, and hence wider availability, than MRI. Ultrasound imaging is also generally faster, with widely available systems operating at 30 frames per second. Ultrasound imaging also imposes different and perhaps looser requirements on the geometry and construction of the focused ultrasound transducer assembly. However, MRI images are often more effective for both guidance and monitoring. In many instances, MRI images more clearly show features such as tumor

4

locations that are of interest during treatment. Tissue changes in response to treatment can be seen well in MRI, especially with contrast enhancement [39]. Finally, ultrasound imaging is not generally possible through the skull.

One particular advantage of MRI in the context of focused ultrasound surgery is its ability to produce sensitive, accurate images showing the temperature at each location.[42, 61] The most commonly used technique is called Proton Resonance Frequency (PRF) shift thermometry.[41] This approach relies on the temperature sensitivity of the chemical shift of protons in water, -0.01 ppm/℃.[36] By measuring the PRF throughout the image volume during each sonication, it is possible to observe the course of temperature elevation associated with each sonication. Thermometry of this kind is also useful for guidance, as it can detect small temperature changes (as small as 1℃) that are not medically significant. Thus, a common strategy is to use the focused ultrasound transducer at a low power to produce a slight heating, visible on MRI, in order to confirm that the focal location is as desired. Then a second sonication may be performed in the same location at higher power to achieve the desired effect (e.g. thermal ablation) [58]. (Other MRI imaging methods, such as MR-ARFI [57], have also been used to verify the location of an ultrasound focus spot.)

MRI thermometry has proven tremendously useful, but it is also limited by its need for high-quality MRI images. Acquiring a complete temperature image can take several seconds, and any motion during this time will corrupt the image, potentially creating artifacts that overwhelm the temperature information. This high sensitivity to motion has presented a serious challenge for MRI thermometry, and hence MRI-guided focused ultrasound, in the

organs of the abdomen, where large motions are present due to respiration. For example, the publication record does not yet include any trials of MRI-guided focused ultrasound of the liver in humans. (Both trials noted above were performed under ultrasound imaging guidance, and hence without the benefit of thermometry imaging.)

In order to enable MRI-guided focused ultrasound of the liver during free breathing, several key challenges must be met. Ultrasound energy must be directed through the ribcage (*transcostal*) to the desired target location, without unacceptable heating or attenuation at the ribs. The focal point must move to follow the motion of the target volume. Finally, a thermometry technique is required that can operate correctly in MRI images of a moving volume.

In existing human studies, no special effort has been made to avoid delivering ultrasound power to the ribs. As a result, patients have experienced painful superficial burns, visible on the skin.[40] These burns were considered acceptable in the trial context, but a growing body of literature describes methods to reduce the amount of unnecessary ribcage heating by using phased array transducers.[5, 10, 35, 82] Manufacturers such as Insightec have also begun developing phased array transducers suitable for transcostal applications.

Phased array transducers can easily be configured to change the location of their focus over time in response to a position input. To enable this treatment, a motion tracking system that can provide such a position estimate is required. A variety of position estimation systems have been developed for the express purpose of steering a focused ultrasound array, deriving position

6

estimates from MRI images [3, 71, 76], longitudinal speckle tracking [52, 64], 2D ultrasound imaging [4], or breath bellows [46]. The technique presented here, biometric ultrasound navigation, is another entrant in the category of motion tracking systems suitable for focused ultrasound guidance.

Once a steerable focused ultrasound array with real time motion tracking is in place, the problem remains to produce MRI thermometry images that are not corrupted by the presence of motion. Traditional MRI thermometry will not suffice, due to its reliance on subtraction from a reference frame. Proposed solutions to the problem of temperature estimation include referenceless methods [68], multiple-baseline approaches [86], and combinations of these [20, 32].

Respiratory motion is a major cause of image degradation in many applications of abdominal MRI, including purely diagnostic radiology. The artifacts caused by respiratory motion are often alleviated by the use of MRI navigators.[23] In typical MRI navigator techniques, short imaging blocks that serve to assess the current respiratory phase are interleaved between the blocks that produce medically relevant images. For example, the navigator block may be a pencil beam excitation that crosses the diaphragm [90], or a bright-blood single-shot EPI image of the liver region [71].

The position information derived from the navigator can be applied *retrospectively* to compensate during image reconstruction for the effects of motion [23], or *prospectively* to change the excitation and readout parameters for the next imaging block [54]. Navigator techniques may also be classified as *gating* if they simply reject unusable data or *motion correcting* if they modify each acquisition in proportion to the measured position offset.[75]

7

The position estimates from an MRI navigator can also be used to control another real-time system, such as a steerable focused ultrasound ablator.[69]

(The term *position*, here and throughout, is used to mean a representation of the complete geometrical state of the anatomy, a vector in $\mathbb{R}^n$ whose dimensionality depends on the complexity of the motion model in use. A simple linear shift model is parameterized by a one-dimensional position, whereas a complex deformation model could be characterized by a position in 10 or more dimensions.)

Navigators often achieve good artifact reduction, but they also have significant costs. Navigator echoes typically slow down the imaging process, can interfere with steady-state magnetization, and require laborious and difficult pulse-sequence engineering that may need to be repeated for each combination of navigator type and imaging sequence.

We have developed an alternative to MRI navigator echoes in the context of abdominal imaging. The proposed technique employs a single MRI-compatible ultrasound transducer, placed against the abdominal skin, which produces a pencil-beam ultrasound field that is oriented approximately in the dorsal direction. For maximum accuracy, the transducer is positioned so that its beam passes near the center of the region of interest.

A number of papers have demonstrated that pulse-receive ultrasound devices can be operated in the MRI environment.[77][34][15][81][19] Previous studies that used ultrasound to compensate for motion in MRI have arranged for their ultrasound data to indicate the position directly. In one approach, a pencil-beam ultrasound transducer is oriented so that the direction of motion is along the axis of the beam.[19] The position may then be computed

directly from any shift observed in the echo delay. As the authors of [19] note in their conclusion, this technique is not expected to work *in vivo* because organs such as the liver predominantly move in the cranio-caudal direction, so no externally placed transducer can be aligned with the motion as required. As a topic for future work, they suggest that tracking such motions with a single ultrasound probe may require a training-based method like the one we demonstrate here.

To solve the problem of tracking motion along an inaccessible axis, the direct shift tracking technique has been extended to three or more transducers.[64] The transducers are widely spaced so that their beams may be oriented toward the focus at large relative angles. The direction of motion can be determined from the shift observed in each transducer. This technique has the advantage of direct displacement measurement, but requires multiple-transducer transmit-receive capability, is intrinsically limited to simple translation measurement, and may drift as errors accumulate due to velocity integration in the position estimate. (This difficulty could potentially be addressed using techniques like the one described in Chapter 5.) It has not been demonstrated in conjunction with MRI.

To permit measurement of non-translational motion and avoid the problems associated with cumulative estimators, another technique employs a linear ultrasound transducer array that produces 2-D ultrasound images.[34] The position is indicated directly by shifts and rotations observed in these images. This approach has recently been demonstrated *in vivo* for motion compensation in cardiac imaging [24], and also for MRI-guided focused ultrasound in a liver phantom [4]. This technique is capable of tracking motion

along the craniocaudal axis, but requires a substantial investment in ultrasound equipment and electronics. It is also limited to motion within the ultrasound imaging plane, and will not detect small displacements in the through-plane direction, which may be important in interventional applications. To achieve 3-D sensitivity, direct image-based tracking of organ position has been extended to use 4-D imaging with ultrasound matrix probes.[6] However, this technique has not yet been demonstrated in conjunction with MRI, and requires an even greater expenditure on ultrasound equipment.

The approach we have developed works differently, because the ultrasound data do not directly indicate the organ position. Instead, it is sufficient that the ultrasound echoes observed at different respiratory phases are distinctly identifiable, as theorized and verified in Chapter 2. The organ position is indicated by an MRI *training navigator*, and ultrasound data are acquired concurrently with each MRI position measurement. From these synchronized training data, the system constructs a table that maps observed ultrasound lines to the positions indicated by MRI. During motion compensation, the system makes frequent ultrasound echo measurements and uses the table of training data to infer the current position. The system is thereby able to estimate the current position without using any MRI time and without altering the steady state magnetization. The pulse sequence need only be modified if prospective correction is required and the pulse program does not already accept position updates from external sources.

We describe the ultrasound measurement as a *biometric navigator*, by analogy with biometric identifiers such as fingerprints and iris scans. Like an iris scan or a fingerprint, an ultrasound echo contains a pattern produced

by stochastic biological processes during gestational development. In both cases, a training process is required to relate the raw biometric data (an iris scan/ultrasound echo) to the desired information (the name of a person/position of an organ). In both cases, the process relies on uniqueness: each individual or position must exhibit a distinct biometric signature for each system to work as designed.

A related approach has recently been presented, using the signal produced by a respiratory bellows to identify the current liver position.[46] These methods are both indirect: they require a training period to determine the correspondence between the measured quantity and the actual anatomical position. Both methods are therefore constrained by the contents of the training data; they cannot track motion with more accuracy than the navigator used to train them, and they can only estimate positions along axes measured by the navigator. They are distinguished, however, by the dimensionality of the measured quantity. A bellows measures a single one-dimensional value, which therefore is not biometrically unique. Without biometric uniqueness, the system cannot detect when the training data are no longer applicable, which may occur if the patient moves laterally or the organs undergo any unanticipated shift. One-dimensional measurements are also restricted to reciprocal motion models unless strong assumptions are made about the true motion cycle (as in cardiac MRI with ECG gating).

Our implementation of this design for biometric ultrasound navigation is called ULTRACK, and is described in detail in Chapter 3. To test UL-TRACK, we constructed motorized motion phantoms appropriate for both ultrasound and MRI imaging, moving with amplitude and frequency approx-

imately consistent with human respiratory motion. As we are interested in tracking liver motion during thermal ablation in freely breathing patients (similar to [46] and [3]), we also tested the accuracy of the system's position measurements and its ability to reduce motion artifacts in 2-D MRI images, including MRI thermometry images produced by the Proton Resonance Frequency Shift method.[55] Chapter 6 contains a detailed description of our implementation of PRF shift thermometry, which includes optimizations particular to this application. To test the method's feasibility *in vivo*, we verified its performance in a freely breathing rabbit. The *in vivo* results, presented in Chapter 4, show that the system performs well, with good time efficiency, little additional position error, and dramatic reductions in visual artifacts.

To further improve the system's position error, we have developed a method for generating filters to enhance the position estimation accuracy, described in Chapter 7. The system produces Kalman-like filters that are automatically tuned to simultaneously reduce noise and compensate for latency.

Although the system, as described so far, requires an MRI navigator in order to train the system, it is possible to provide significant motion compensation and gating functionality in the absence of an MRI navigator, by employing a more sophisticated analysis of the ultrasound data. Algorithms for biometric navigation in the absence of a training navigator are presented in Chapter 5.

He propelled a hand-held array across the top of the desk. "Put your palm on that, we need a biometric. You're not cleared for the medical levels yet, understand?"

<div align="right">

Judas Unchained – Peter F. Hamilton

</div>

2

Models of Biometric Navigation and Dissimilarity

BIOMETRIC NAVIGATION IS DISTINGUISHED from other position measurement methods by its use of a unique identifier derived from the tissue at each location. This identifier does not itself indicate a position, but its biometric uniqueness ensures that once a position is associated with an identifier, the position will not be indicated spuriously. The uniqueness of the identifier

representing each location is due to the large entropy of the high-dimensional space from which biometric patterns are randomly drawn.

In the case of ultrasound navigators, entropy is provided by reflections from macroscopic anatomical structures, but also by sound scattered off of countless microstructures. The resulting interference pattern is termed *speckle*, and is well approximated by a broadband Gaussian random field that does not change with time unless the relevant microstructures are moved or altered.[1]

Most cellular tissues in the body, including the liver, produce a strong speckle signal in backscatter ultrasound. The analysis in this section assumes that the only content of the signal is uniform speckle. This assumption is certainly untrue; the existence of diagnostic ultrasound imaging disproves it. However, we expect that additional scattering from macroscopic structures will typically further improve the uniqueness properties of the signal, by adding more distinct content at each location. If non-uniformity in the signal proves disruptive, it may be possible to apply some processing (such as dynamic range compression) to improve apparent uniformity.

The uniqueness property only applies for measurements that are unrelated, in the sense that they represent patterns produced by non-overlapping sets of scatterers. Clearly two measurements made at the same location, and otherwise in the same configuration, will be identical, and measurements made very close together will be very similar.

We may anticipate that the length scale above which a transverse motion produces a new, unrelated scattering pattern is set by the shape of the ultrasound beam, which is the volume within which scatterers contribute

to the received speckle pattern. If two measurements are made sufficiently far apart that the beam profiles are largely non-overlapping, then we may expect that the two echo patterns will show no relationship. Conversely, if the two measurements are made close together, so that the shift is small relative to the scale of variation in the beam, then we may expect that the two echo patterns will be very similar.

A typical simple ultrasound transducer has a flat, circular active surface that is several times large in diameter than the wavelengths at which it operates. Transducers of this type are sometimes termed *pencil-beam* transducers because they produce a beam that is approximately collimated near the tranducer surface, diverging only gradually at greater distances. For transverse motion, the similarity length scale is set by the width of this column.

This simple intuitive model of echo pattern similarity at different distances is sufficient to understand how biometric navigation might be possible. Suppose that a table of ultrasound echoes is recorded that covers the range of observed motion at a granularity much smaller than the similarity length scale, and also suppose that each of these echoes is associated with a known position. Then when a new echo is recorded, it will be similar only to those echoes that correspond to nearby positions, and will be unrelated (i.e. dissimilar) to the other echo patterns representing positions that are further away. The system may therefore infer that the present position is close to the position associated with the similar echoes in the table. This method is discussed at greater length in Section 3.1.

2.1 Information theoretic arguments for uniqueness

We may estimate the entropy of the ultrasound navigators by considering only the speckle and applying the Shannon-Hartley channel capacity theorem for a noisy channel:[78]

$$C = B \log_2 \left(1 + \frac{S}{N} \right).$$ (2.1)

In our experiments a typical ultrasound transducer has bandwidth of at least $B = 1.5$ MHz and signal-to-noise ratio of at least $S/N = 3$. (Note that although speckle is often regarded as noise, here the time-invariant speckle is *signal*, and its magnitude sets the value of S in this calculation.) From these parameters we compute a channel capacity of $C = 3$ Megabits per second, or 4 bits/mm in pulse-echo ultrasound. (For comparison, biometric iris scans exhibit an effective entropy of 3.2 bits/mm^2.[16]) A target region 20 mm in size would have an entropy of 80 bits, or a collision probability of 10^{-24}. This strong expectation of uniqueness motivates our study of ultrasound navigators for position tracking.

We use the term *false positive* to describe a situation in which two signals are deemed to match even though the signals are actually drawn from distant locations. In a probabilistic model of echo patterns, the chance that any individual's echo patterns exhibit a collision that could produce a false positive is never exactly zero, but it can be negligibly small.

To estimate the probability of a false positive, we return to the example above, a signal space representing 20 mm of tissue at a bandwidth of 1.5 MHz and $S/N = 3$. This signal contains 80 bits of random information, providing

high uniqueness. We will sacrifice some of this uniqueness by employing the shift-invariant distance metric in Section 2.2, so we now include an estimate of the entropy lost to shift invariance. A bandwidth of 1.5 MHz corresponds to a longitudinal decorrelation length of 1 mm, indicating that the set of all shifts of each test signal corresponds approximately to 20 different independent test signals. This is between 2^4 and 2^5, so it corresponds to a loss of 4–5 bits of information, leaving about 75 bits of unique information.

By the birthday theorem,[47] a collision would be expected when there are more than $2^{75/2} = 2 \times 10^{11}$ samples in the set. Our ultrasound beams typically have a width of at least 1 mm, so this would correspond to a displacement path spanning 2×10^8 meters. A typical human has a respiratory displacement path spanning about 5×10^{-2} meters) [17], well below the birthday paradox threshold. Over that path length, the collision probability is approximately 3×10^{-20}, indicating that there is less than one chance in a billion that such a collision would occur for any human being presently alive.

These extreme numbers are highly approximate estimates, not accounting for any of the nonidealities of a real (finite) implementation, for which perfect information capture performance is mathematically impossible.[78] However, the very low probability of false positives suggests that there is a significant engineering safety factor: even imperfect implementations will be unlikely to suffer from false positives and similar errors.

2.2 Dissimilarity function

To make use of this intuitive model we must formalize a definition of similarity. We do so by defining a dissimilarity function $D(\boldsymbol{U}, \boldsymbol{V})$ between two ultrasound signals $U(t)$ and $V(t)$ as

$$D^2(\boldsymbol{U}, \boldsymbol{V}) = \min_{\Delta t} \int_{-\infty}^{+\infty} (W(t)U(t) - W(t + \Delta t)V(t + \Delta t))^2 \, dt. \qquad (2.2)$$

By this definition, $D(\boldsymbol{U}, \boldsymbol{V})$ is the minimum Euclidean distance between $\boldsymbol{U}$ and any shifted copy of $\boldsymbol{V}$, where both $\boldsymbol{U}$ and $\boldsymbol{V}$ have been windowed by multiplication with a window function $\boldsymbol{W}$. For a pencil-beam transducer, a window in time corresponds approximately to a range of depths, or a total volume approximately in the shape of a finite cylinder. Minimizing across a range of time shifts Δt is important in order to isolate transverse displacement (across the ultrasound beam) from longitudinal displacement (along the beam axis). The ultrasound signals are rapidly varying functions of t, so without this minimization a small shift in time (corresponding to a very small displacement along the direction of the beam) might result in a large difference between two otherwise identical signals.

The sensitivity of the difference to small shifts in time is set by the highest frequency containing significant energy in the signal. Thus, one alternative to minimizing over many shifts would be to produce a new signal with a lower maximum frequency. Because our signals are naturally band-limited by the physical properties of our ultrasound transducers, it is often possible to demodulate without a large loss of information. However, this approach has not been very attractive because much of the information content in the sig-

nal is in the phase, so a simple magnitude-only demodulation (i.e. an amplitude envelope) is not sufficient. The computational complexity of broadband phase-preserving demodulation is high enough that it does not represent an important savings over our optimized shift search (Section 3.2.1).

The value of Δt that minimizes $D(\boldsymbol{U}, \boldsymbol{V})$ represents an estimate of the longitudinal shift, which may be valuable information when motion along this axis is of interest. It is even possible to envision a variety of techniques that combine this shift with the transverse dissimilarity for tracking of motions oblique to the beam path. However, throughout this thesis, we will assume that all interesting motion is transverse to the primary beam axis, and discard the information provided by Δt. This greatly simplifies the analysis of system behavior, and also corresponds well to the natural geometry of our primary application.

In an actual implementation, the continuous ultrasound signal will necessarily be sampled at some rate. Discretizing Equation 2.2 to operate on sampled signals produces a new formula

$$d^2(\boldsymbol{u}, \boldsymbol{v}) = \min_s \sum_i \left(w_i u_i - w_{i+s} v_{i+s} \right)^2 \tag{2.3}$$

where $\boldsymbol{w}$ is a rectangular window corresponding to the selected range of samples. Discretization is necessarily an approximation, but the error introduced is small as long as the signal is well oversampled, i.e. most of the signal's energy is present in frequencies that are small relative to the Nyquist rate. A detailed description of our optimized implementation of Equation 2.3 for real-time applications is available in Section 3.2.1.

19

2.3 Gaussian Random Field models of dissimilarity

Intuitively, the ultrasound signal observed at a given location is determined by the set of scatterers that fall within the ultrasound beam profile. A transverse motion much smaller than the beam's width only slightly changes the set of scatterers, and so should only slightly change the received signal. This property is important for our application; if small changes could cause unbounded variation in the ultrasound signal, then no finite number of samples would suffice to capture the variation present in some region.

We can formalize this intuition by determining how the expected dissimilarity of two ultrasound lines depends on the transverse distance that separates them. To do so, we model the ultrasound echo pattern as a stationary Gaussian Random Field (GRF). This model is known to apply well for ultrasound speckle due to the central limit theorem, although it is not an accurate description of larger-scale anatomical structures.[1]

A stationary zero-mean GRF is fully described by a domain vector space S_D, a range vector space S_R, and an autocovariance function

$$\mathcal{C} : S_D \rightarrow S_R \times S_R.[2] \qquad (2.4)$$

In this model of transverse motion, S_D is the two-dimensional position space $\mathbb{R}^2$, S_R is the space of ultrasound signals, windowed to represent the depths of interest, and $\mathcal{C}$ represents the whole spatiotemporal sensitivity of the ultrasound transducer.

To model the behavior of the dissimilarity of ultrasound signals made at various transverse separations, we suppose for simplicity that the shift

search finds only the geometrically correct alignment, so that $\mathbf{d}^2$ is exactly the Euclidean metric. Consider two ultrasound signals $\boldsymbol{U}$ and $\boldsymbol{V}$ recorded at a two-dimensional transverse separation $\vec{x}$. Each signal is then a random column vector drawn from a Gaussian distribution on $\mathrm{S_R}$ with covariance matrix $\mathcal{C}(\vec{0})$, and the covariance between the two vectors is $\mathcal{C}(\vec{x})$, i.e.

$$\langle \boldsymbol{U}\boldsymbol{U}^* \rangle = \langle \boldsymbol{V}\boldsymbol{V}^* \rangle = \mathcal{C}(\vec{0}), \tag{2.5}$$

$$\langle \boldsymbol{U}\boldsymbol{V}^* \rangle = \langle \boldsymbol{V}\boldsymbol{U}^* \rangle = \mathcal{C}(\vec{x}). \tag{2.6}$$

We define the difference between these vectors as $\boldsymbol{\Delta} \equiv \boldsymbol{U} - \boldsymbol{V}$, a new gaussian random vector with mean zero and covariance

$$\mathcal{C}_{\boldsymbol{\Delta}}(\vec{x}) = \langle \boldsymbol{\Delta}\boldsymbol{\Delta}^* \rangle = \langle (\boldsymbol{U} - \boldsymbol{V})(\boldsymbol{U}^* - \boldsymbol{V}^*) \rangle = 2\mathcal{C}(\vec{0}) - 2\mathcal{C}(\vec{x}). \tag{2.7}$$

With this model for the difference between ultrasound signals, we now consider two special cases: signals measured at very distant locations, and signals measured at very close locations.

2.3.1 Dissimilarity of distant ultrasound signals (no false positives)

When two ultrasound echoes are acquired at widely separated locations, they are essentially uncorrelated, because each signal is generated by a different random set of scatterers within the tissue. More formally, we may assert that

$$\lim_{|\vec{x}| \to \infty} \mathcal{C}(\vec{x}) = 0, \tag{2.8}$$

21

because the covariance function is determined by the ultrasound beam profile, which is spatially localized. If two signals U and V are acquired at sufficiently wide locations that $\mathcal{C}(\vec{x}) \approx \mathbf{0}$, then their difference $\boldsymbol{\Delta}$ is a Gaussian random vector with mean zero and covariance $2\mathcal{C}(\vec{0})$. Although the mean difference is zero, the probability that the difference is actually near zero can be low if the effective dimensionality (i.e. degrees of freedom) is high.

The effective dimensionality of an ultrasound echo is approximately given by its time-bandwidth product. For the example numbers above (20 mm$\rightarrow$ 27μs, 1.5 MHz) this is approximately 40 degrees of freedom, meaning that magnitude of the difference is distributed according to a χ distribution with a parameter of approximately $k = 40$. This distribution is approximately normal, with its mean greater than the standard deviation by a factor of

$$\frac{\mu}{\sigma} = \sqrt{2k-1} = 8.9. \tag{2.9}$$

The chance of a sample from a normal distribution being more than 6 standard deviations below the mean is less than 10^{-9}, and is lower still for a χ distribution. Therefore, we may expect to find a threshold value, well above zero, such that any pair of ultrasound echoes whose dissimilarity is below the threshold were almost certainly acquired at nearby locations with overlapping sound fields. With high confidence, we may infer a close distance when two measurements show low dissimilarity. We term this property the guarantee of *no false positives*.

2.3.2 Dissimilarity of nearby ultrasound signals (no false negatives)

When two measurements are sufficiently close together, we would also like to have confidence that we will not make the opposite error of Section 2.3.1 and incorrectly judge them to be far apart. To see that this is the case, we first note that the ultrasound beam profile, and therefore the autovariance $\mathcal{C}$ is smooth below some transverse length scale. In the case of our disc-shaped transducers, the beam profile's smallest length scale (outside the evanescent wave zone) is set by the ultrasound wavelength and disc radius.[14] Therefore, below this length scale, $\mathcal{C}$ is smooth and differentiable. As the beam profile is radially symmetric, and there is no distinction between the two transverse axes of displacement, our model is isotropic, and the covariance function depends only on the distance $|\vec{x}|$.[2] We may therefore approximate the autocovariance by a function of scalar distance $\mathbf{C}(|\vec{x}| > 0)$, with a Taylor expansion near the origin:

$$\mathcal{C}(\vec{x}) = \mathbf{C}(0) + \frac{1}{2}\mathbf{C}''(0)|\vec{x}|^2 + O(|\vec{x}|^4). \tag{2.10}$$

Combined with the previously derived expression for the covariance of the difference vector $\mathcal{C}_{\mathbf{\Delta}}$ (Equation 2.7), we may conclude that for small $\vec{x}$,

$$\mathcal{C}_{\mathbf{\Delta}}(\vec{x}) = -\mathbf{C}''(0)|\vec{x}|^2 + O(|\vec{x}|^4). \tag{2.11}$$

Using the covariance, we may compute the expected dissimilarity:

$$\langle ||\mathbf{\Delta}||^2 \rangle = \langle \mathbf{\Delta}^* \mathbf{\Delta} \rangle = \mathrm{tr}(\mathcal{C}_{\mathbf{\Delta}}(\vec{x})) \propto |\vec{x}|^2 + O(|\vec{x}|^4). \tag{2.12}$$

Thus, the mean squared dissimilarity is expected to be linearly proportional to the squared distance, for small distances. By an argument on the number of degrees of freedom, as in Section 2.3.1, we expect the standard deviation of $||\mathbf{\Delta}||^2$ to be relatively small, so that values are tightly clustered.

For the proposed dissimilarity-based method to work effectively in practice, the essential requirement is to be able to tell reliably whether a newly measured ultrasound signal was acquired near to the position of one that has been recorded previously. At a minimum, it must be possible to distinguish reliably between an ultrasound echo acquired close to the reference, and one acquired far from the reference. Since the dissimilarity of nearby points becomes arbitrarily small as the points get closer, and the dissimilarity of far away points is effectively bounded below by a constant, it follows that there exists a discriminatory threshold T^2 and two distances $a > b > 0$ such that if the true distance is greater than a, then the dissimilarity d^2 is always greater than T^2, and if the true distance is less than b, then the d^2 is always less than T^2. (See Figure 2.2 for an example of data providing these properties.) This latter property allows a guarantee of *no false negatives*.

2.4 Experiment: Distance-Dissimilarity Relationship

The purpose of this experiment was to determine whether the expected behaviors described in Section 2.3 are observed in practice.

2.4.1 Materials: Phantom tracking transducer

The phantom tracking transducer used was a broadband piston-type ultrasound transducer with a nominal center frequency of 5 MHz and diameter of

4 mm. It was the center element of an annular array produced by Acoustic Technologies Laboratories. It was not manufactured for MRI-compatibility, and as such did produce a significant dropout of MR signal in its immediate vicinity due to the induced magnetic field inhomogeneity.

The phantom tracking transducer was driven by a pulser-receiver (Olympus 5072PR) pulsing at its lowest energy setting (13 μJ) and a damping impedance of 50 Ω. The receiver was configured for a bandpass of 1–10 MHz and +30dB of amplification.

The ultrasound field of the phantom tracking transducer was mapped in a tank of degassed water with a 0.2-mm diameter needle hydrophone (Onda HNC-0200) mounted on a computer-controlled, three-axis positioning system (Velmex VP-9000). (Experiments testing the accuracy of displacement tracking in the structured phantom target were performed by attaching the transducer to the same positioner.) A maximum ultrasound intensity of 16 mJ/m^2 (pulse intensity integral) was measured at a focal distance of 36 mm, with a half-intensity width of 2.0 mm and length of 55 mm. At the focus, the maximum spectral energy density occurred at 5.3 MHz, with a half-power bandwidth of 1.6 MHz ($Q = 3.2$). For additional characterization data, see Figure 2.1.

2.4.2 Methods

We attached the tracking transducer to a computer-controlled positioning device, aimed at the structured phantom in a water tank. The positioner scanned 1000 locations in each of two transverse grid patterns (coarse and fine), recording the ultrasound echo at each location. The dissimilarity func-
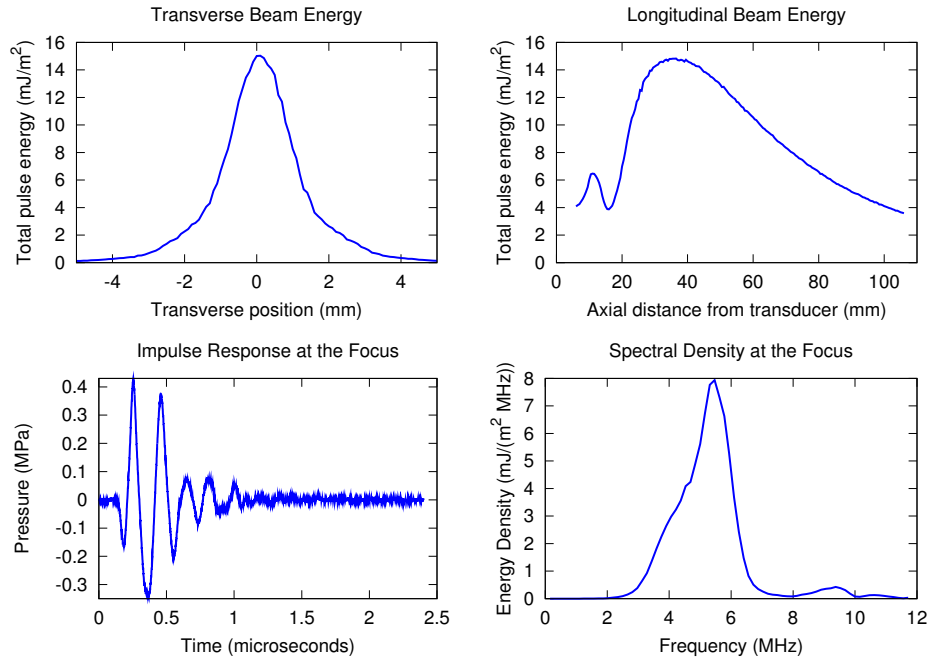
Figure 2.1: Characterization of the phantom tracking transducer, as described in Section 2.4.1.

tion was then computed for all pairs of ultrasound echoes in each data set (10^6 pairs in total), and compared to the true transverse distance between the two locations in the grid. This arrangement was also used to simulate the response of motion compensation to a lateral shift (i.e. motion perpendicular to the expected trajectory).

2.4.3 RESULTS

The results of the dissimilarity assessment are plotted in Figure 2.2. At short distances, below 0.5 mm, the squared dissimilarity is linearly related to the squared distance. At large distances, above 2 mm, the dissimilarity

26

Figure 2.2: The dissimilarity measured between all pairs of ultrasound signals separated by a given distance. The center line shows the mean dissimilarity observed at that distance, the shaded region indicates one standard deviation above and below, and the outer lines show the minimum and maximum. Data for nearby points are shown in (a) with quadratically scaled axes, and data for larger distances are shown in (b) with linear scaling.

saturates, and becomes constant.

These behaviors are as expected from a simple model of ultrasound echogenicity, in which the signal changes smoothly on scales smaller than the ultrasound beam width, but becomes uncorrelated at distances larger than the beam width (2.0 mm for this transducer). The offset at zero distance is expected due to factors such as thermal and electrical noise that decrease the similarity of successive measurements even if the location is unchanged.

Importantly, the data suggest that an appropriately chosen dissimilarity threshold may reject all pairs separated by more than 2 mm, and accept all pairs separated by less than 0.5 mm. This is sufficient to enable reliable matching as required by the subsequent experiments.

3

ULTRACK: a prospective Biometric Ultrasound position tracking system

TO PUT THE ULTRASOUND DISSIMILARITY MEASURE OF CHAPTER 2 TO USE, we constructed a system called ULTRACK that provides prospective motion correction of MRI images. In order to validate the effectiveness of the ULTRACK motion compensation system we conducted a series of experiments in motorized motion phantoms, designed to resemble the geometry

and motion parameters of potential medical applications of the ultrasound biometric navigation technique for MRI motion compensation. The results show that the system performs well, suggesting that the principle of biometric navigation may be applicable to real-world motion compensation problems.

3.1 Basic Biometric Tracking Algorithm

Having concluded in Chapter 2 that the dissimilarity measure will not produce false positives or false negatives, we may outline a basic algorithm that uses dissimilarity for position tracking. This algorithm requires two inputs: the biometric signatures and an additional source of position information. It operates in two stages: training and tracking. The algorithm relies on one core assumption: the motion is essentially repetitive, in the sense that it crosses and recrosses the same locations, not deviating by further than the length scale set by the biometric signature.

In the training stage, the system accepts synchronized input, forming pairs of simultaneous biometric and position data. These pairs are accumulated over the course of the training stage, which should be long enough to capture most or all of the distinct patterns that are seen along the motion path.

When the training period concludes, the system *summarizes* the training data, producing a smaller set of paired biometric signatures and positions. Summarization may proceed by selecting a subset of the training data according to some criterion, or it may include a more complex synthesis of the training data, producing biometric signatures that are not identical to any single input measurement. One reasonable heuristic for summarization is to

select a set of input measurements that are approximately equally separated (according to the position data) and span the range of observed motion. The reduced set of signatures and positions is called a *mapping table*, because it serves to map biometric signatures to their corresponding positions.

Once the mapping table is ready, the system may transition to the tracking stage. In this stage, only biometric input is received. Each input is compared, using the dissimilarity measure, to the biometric signature component of each entry in the mapping table. Whichever entry produces the lowest dissimilarity is termed the *best match*, and its position component is transmitted as the response to the input signature.

If the dissimilarity between the input and the best match is too high, then by the property of no false negatives, the input signature must not be very close to any element of the mapping table. While the best match is still the best estimate of the current position, the confidence in this estimate is now very low. In some systems it may be valuable to convey this information, for example by replacing the transmitted position by a flag indicating that no reliable position estimate could be computed. This kind of failure to match can occur if the repetitivity assumption is violated (i.e. the system has moved to a position not explored during training), or if the summarization discarded too much information.

3.2 BIOMETRIC ULTRASOUND POSITION TRACKING SYSTEM: UL-TRACK

We constructed an ultrasound-MRI motion compensation system that implements the basic biometric tracking algorithm where the biometric signatures

are ultrasound echoes from a single transducer, and the position measurements are derived from one-dimensional MRI navigators. (There is no substantial barrier to the use of multi-dimensional navigators, but development of that capability was not undertaken.) At the core of the system is a software program called ULTRACK that manages all ultrasound data collection and MRI motion compensation. Real-time components of ULTRACK are written in C, and other components are written in Python using the SciPy numerical analysis framework.[44]

The system is designed to operate in two stages. First, in the training stage, an MRI navigator that indicates all desired position information is configured to run continuously. The MRI control electronics produce an external trigger pulse at the beginning of each navigator block. This pulse triggers the acquisition of an ultrasound echo by the pulser-receiver and digitizer (National Instruments PXI-5124). Each ultrasound echo is recorded to disk for later analysis, as are the corresponding MRI data.
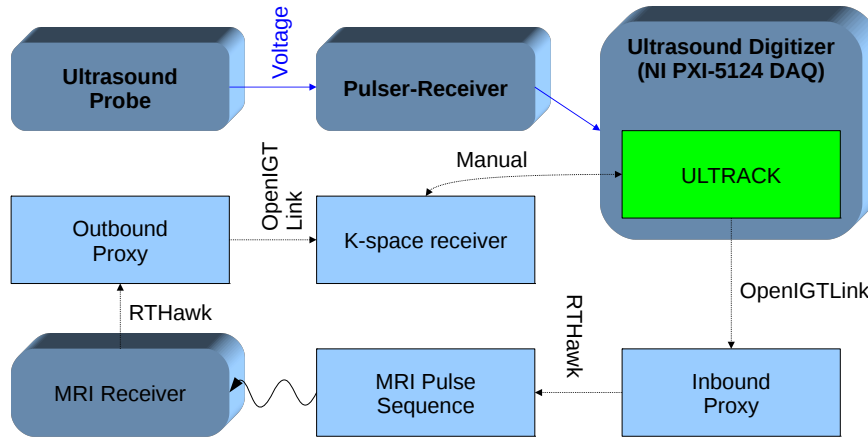
Figure 3.1: Diagram of the motion compensation system including ULTRACK

The training period must be long enough to ensure that the entire range of motion is densely sampled, so it typically spans multiple motion cycles. Once all the training data are recorded, the ultrasound and MRI data are condensed into a *mapping table* containing a set of ultrasound measurements and their corresponding positions. For reasons of computation speed, the maximum number of entries in the mapping table may be limited. In this case, the training software must select a representative subset of measurements for use in the mapping table.

During the correction stage, an identical triggering arrangement is used, and the pulse sequence is configured to produce the desired images. UL-TRACK compares each incoming ultrasound line against every entry in the mapping table, using the discretized dissimilarity function. The current position estimate is determined to be the position associated with the most similar entry.

Once a position estimate has been determined, ULTRACK transmits the position update to the pulse program running on the MRI scanner. The pulse program changes its imaging parameters as needed to center on the new position. The position information transmitted by ULTRACK includes a three-dimensional vector indicating the center of the imaging volume and a rotation matrix indicating the slice coordinate system. The pulse sequence was modified to check for changes in these parameters before every excitation pulse, and adjust oscillator frequencies and gradient amplitudes accordingly.

In the current implementation, ULTRACK transmits positions, and receives MRI data, using the OpenIGTLink protocol.[83] The pulse program receives positions and transmits data using the RTHawk realtime library for GE scanners.[72] These two systems are connected by a proxy server, running programs that convert messages between these two protocols. The servers communicate over ethernet.

3.2.1 Optimized Computation of Discretized Dissimilarity

Biometric navigation, when used in real-time applications such as prospective motion compensation, requires fast computation of dissimilarity between newly measured biometric data and entries in the mapping table. For biometric ultrasound, this means a fast method of computing the dissimilarity function described by Equation 2.3. To find a fast implementation, we first reformulate Equation 2.3 to expand the quadratic and extract constant

terms:

$$d^2(\boldsymbol{u}, \boldsymbol{v}) = \min_s \sum_i w_i^2 u_i^2 + w_{i+s}^2 v_{i+s}^2 - 2w_i u_i w_{i+s} v_{i+s} \tag{3.1}$$

$$= \min_s \left(\sum_j w_j^2 u_j^2 \right) + \left(\sum_k w_k^2 v_k^2 \right) + \sum_i -2w_i u_i w_{i+s} v_{i+s} \tag{3.2}$$

$$= \left(\sum_j w_j^2 u_j^2 \right) + \left(\sum_k w_k^2 v_k^2 \right) - 2 \max_s \sum_i w_i u_i w_{i+s} v_{i+s}. \tag{3.3}$$

To achieve maximum speed, we choose a simple rectangular window function

$$w_i = 1 \text{ if } a \le i < a + N \text{ else } 0, \tag{3.4}$$

and define new windowed ultrasound vectors

$$x_i = w_{i-a} u_{i-a}, \tag{3.5}$$

$$y_i = w_{i-a} v_{i-a}. \tag{3.6}$$

This simplifies our formula to

$$d^2(\boldsymbol{u}, \boldsymbol{v}) = \left(\sum_{j=0}^{N-1} x_j^2 \right) + \left(\sum_{k=0}^{N-1} y_k^2 \right) - 2 \max_s \sum_{i=-\min(s,0)}^{N-\max(s,0)-1} x_i y_{i+s}. \tag{3.7}$$

For convenience we define

$$P = \sum_{j=0}^{N-1} x_j^2, \tag{3.8}$$

$$Q = \sum_{k=0}^{N-1} y_k^2, \tag{3.9}$$

$$R_s = \sum_{i=-\min(s,0)}^{N-\max(s,0)-1} x_i y_{i+s}, \tag{3.10}$$

so that

$$d^2(u,v) = P + Q - 2 \max_s R_s. \tag{3.11}$$

If computed naively, R will require an independent sum over products at each value of s, requiring $\Theta(N^2)$ computation time. To perform the computation faster, we first note that R is the cross-correlation vector of $\boldsymbol{x}$ and $\boldsymbol{y}$:

$$R_s = (\boldsymbol{x} \star \boldsymbol{y})_s. \tag{3.12}$$

By the Fourier cross-correlation theorem, Equation 3.12 is equivalent to

$$\boldsymbol{R} = \mathfrak{F}^{-1}(\mathfrak{F}(\boldsymbol{x})^* \cdot \mathfrak{F}(\boldsymbol{y})), \tag{3.13}$$

where $\mathfrak{F}(\cdot)$ represents a discrete Fourier transform of length at least $2N - 1$. With a Fast Fourier Transform, this allows d^2 to be computed in $\Theta(N \log N)$ time.

In our implementation, the dissimilarity function is used to compare a newly acquired ultrasound signal $\boldsymbol{v}$ with each vector $\boldsymbol{u}$ from the mapping table. Because all the vectors $\boldsymbol{u}$ are known before real-time computation

starts, we may compute and store P and $\mathfrak{F}(\boldsymbol{x})^*$ ahead of time. We also only compute Q and $\mathfrak{F}(\boldsymbol{y})$ once for each new signal, and reuse these values in each of the dissimilarity computations. If there are K entries in the mapping table, then these choices reduce the total number of P and Q computations required for each search from $2K$ to 2, and the total number of forward and inverse transforms from $3K$ to $K + 1$. We use the FFTW package [27] to compute these transforms quickly, and perform the other computations in C++.

Because each of the K computations of d^2 is independent, they may be performed in parallel. We use the OpenMP multithreading abstraction to spread these K computations across CPU cores. As expected, we observe near-perfect linear scaling with the number of cores. On an 8-core 2.5 GHz Xeon E3560, with $N = 4096$ and $K = 256$, we observed an average table search time of 2.5 ms, i.e. 400 Hz, or 100,000 computations of d^2 per second. This is significantly faster than required by our current experiments, indicating that larger mapping tables may be feasible.

There are many possibilities for future optimization, should it be necessary, such as excluding unlikely (i.e. distant) reference vectors from consideration or using extremely parallel hardware such as GPUs. This problem should be well-suited for GPU acceleration because of the very small amount of data transfer, highly parallel search, and use of power-of-2 FFTs. The performance optimization techniques used here may also be applicable to variations on this distance metric. For example, it is possible to avoid windowing both inputs, or to use a non-rectangular apodization, without losing the key performance advantages achieved here.

3.3 MATERIALS

3.3.1 STRUCTURED PHANTOM TARGET

The structured phantom target consisted of polyacrylamide gel, cast into a mold containing a few arbitrary non-MRI-visible plastic objects that provided visible structure in MRI. The gel contained 6% powdered silica by

Figure 3.2: A cartoon of the motorized phantom, showing the structured phantom tied to a string that rolls it over a water bath containing the ultrasound transducer.

mass, which created a strong ultrasound speckle. Ultrasound measurements were performed in a region of the phantom containing speckle only; the plastic objects were used only for experiments that served to evaluate the tracking in MRI.

3.3.2 HEATING PHANTOM TARGET

The heating phantom target contained polyacrylamide gel with 3% powdered silica, in order to produce scattering and also match the absorption observed in human soft tissues.[62] A focused ultrasound transducer was attached to the phantom so that its focus was at a stationary point inside the gel.

Figure 3.3: A photo of the heating phantom in the water bath.

3.3.3 MRI-COMPATIBLE MOTORIZED LINEAR MOTION SYSTEM

To model human respiratory motion, we constructed a motorized system that imparted motion to the gel phantoms described in Sections 3.3.1 and 3.3.2. The motor, which was located outside the MRI room, was coupled to the target by a rope running through the waveguide and around several MRI-compatible pulleywheels (K'Nex). The target was attached to a dolly and constrained to move along one direction, parallel to the MRI bore. The gel was partially immersed in a water bath that served to conduct ultrasound from the transducer, which was oriented vertically.

The motor moved the phantom at a controllable speed. The motion was not sinusoidal, but was approximately periodic and reciprocal. The phantom

could be operated at two different motion amplitudes covering a distance of approximately 6.5 cm or 7.5 cm. The motion amplitude was chosen to exceed the maximum reported liver motion observed during deep breathing (5.7 cm).[17]

3.4 Experiment: Assessing the accuracy and latency of ULTRACK

This experiment served to measure the accuracy and latency of ULTRACK's MRI motion compensation by applying the compensation to the navigator itself. If ULTRACK were providing perfectly ideal motion compensation, then the motion-compensated navigator would appear stationary, even though the phantom is moving.

3.4.1 Methods

In this experiment, the imaging pulse sequence was Spoiled Gradient-Echo with TE/TR=4.2/8.9 ms, imaging a 22 cm square coronal slice with 10 mm thickness at a resolution of 256×256. The frequency encode direction was aligned with the MRI bore. The MRI navigator used for training purposes consisted of the same sequence with phase encoding disabled.

The training period consisted of 10^4 acquisitions, requiring 89 seconds to complete and covering approximately 9 cycles of the motion phantom. During the training stage, the phantom was set to its larger motion amplitude, akin to asking a patient to breathe deeply during training, in order to ensure that the entire range of motion was observed.

Analyzing the navigator to determine the position of the phantom was

trivial because the slice consists of a rigid body moving linearly in air. Once the software had extracted positions from each MRI acquisition, it used this information to select 256 measurements that were evenly spaced in position and spanned the whole observed range of motion.

After training was complete, ULTRACK was configured for motion compensation, and the phantom was changed to its lower motion amplitude. The MRI navigator pulse sequence was run again, this time under prospective compensation from ULTRACK. The positions were extracted again from these data by the same algorithm.

3.4.2 Results

The results of Experiment 3.4 are shown in Figure 3.4. Using prospective motion compensation, ULTRACK reduced the observed standard deviation of position from 25.3 mm to 0.57 mm. This indicates that ULTRACK preserved 97.8% of the MRI navigator's positioning accuracy.

The residual position error in Figure 3.4 exhibits periodic behavior due to the latency of correction. These data were acquired and reconstructed without predictive filtering, so any delay contributes to the positioning error during prospective compensation. For retrospective motion correction, the delay can easily be removed. By minimizing the observed positioning error, we determined that the total motion compensation delay in this experiment was 27 ms. Retrospectively compensating for this delay reduced the observed standard deviation to 0.35 mm.
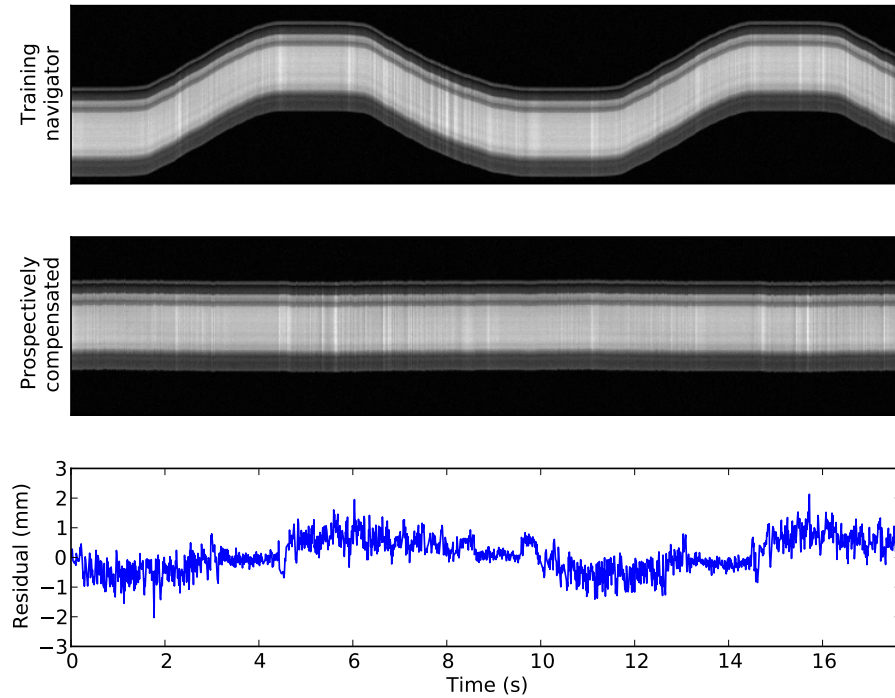
Figure 3.4: The top image shows the MRI training data used to compute the mapping table. The middle image shows the result when the same pulse sequence is performed under prospective navigation from this mapping table. The lower plot shows the residual position variation in the middle image. The variation is the sum of error in the ULTRACK motion compensation and in the MRI navigator itself.

3.5 Experiment: Determining the ability of ULTRACK to reduce artifacts in 2D images

This experiment served to determine whether ULTRACK can successfully achieve its goal of reducing motion artifacts in three different types of 2D images: coronal, axial, and oblique. The oblique plane was generated by rotating the coronal slice by approximately 15 degrees from all three cardinal axes (double oblique).

3.5.1 METHODS

Each image type was tested using the same procedure as in Experiment 3.4, except that phase encoding was enabled during motion compensation. Additionally, a control image set was acquired for each imaging plane with ULTRACK disabled, to determine the effects of motion without compensation.

3.5.2 RESULTS

The results indicate that ULTRACK's motion compensation performance is nearly sufficient to eliminate all motion artifacts from the moving object. In the in-plane motion test (Figure 3.5), the residual artifacts are greatly suppressed, and not easy to see. One source of remaining artifacts

Figure 3.5: Prospective motion compensation performance for in-plane motion. The top row shows images during uncorrected motion. The bottom row shows images acquired during corrected motion, and the middle row shows the corresponding applied position correction over time. All images are originally 256×256, but have been cropped to show the region of interest. The tracking was performed in a homogeneous region of the phantom. Plastic objects were placed behind this region for visualization in MRI.

is position-dependent geometric distortion due to nonideality of the MRI gradient fields. In static images this distortion is corrected by the gradient warping procedure during reconstruction, but no such correction has been applied here. Correction of dynamically varying gradient warping is possible, but it requires the use of complex, computationally intensive reconstruction techniques.[65]

In the through-plane test (Figure 3.7), the appearance of the moving object is substantially improved compared to the uncorrected images, in which the phantom is often invisible because it has moved out of the imaging plane entirely. Unrelated artifacts are still visible due to turbulence in the water tank and the presence of ferromagnetic material in this ultrasound

43

Figure 3.6: The profile of a single line through each of 31 motion compensated coronal images, taken from the same data set as Figure 3.5. The line is parallel to the direction of motion.

transducer.

In the oblique test (Figure 3.8), ULTRACK greatly reduced the appearance of motion artifacts by applying corrections along the phase, read, and slice axes.

3.6 Experiment: Determining whether ULTRACK permits MRI thermometry of a moving phantom

This experiment served to determine whether ULTRACK's motion compensation accuracy is sufficient to enable MRI thermometry of moving objects, using the heating phantom with attached focused ultrasound transducer.

3.6.1 Methods

The temperature rise was estimated by Proton Resonance Frequency Shift analysis of data from a standard sequence used for this purpose: Spoiled Gradient-Echo (SPGR) with TE/TR=11.7/23.8 ms, and a bandwidth of
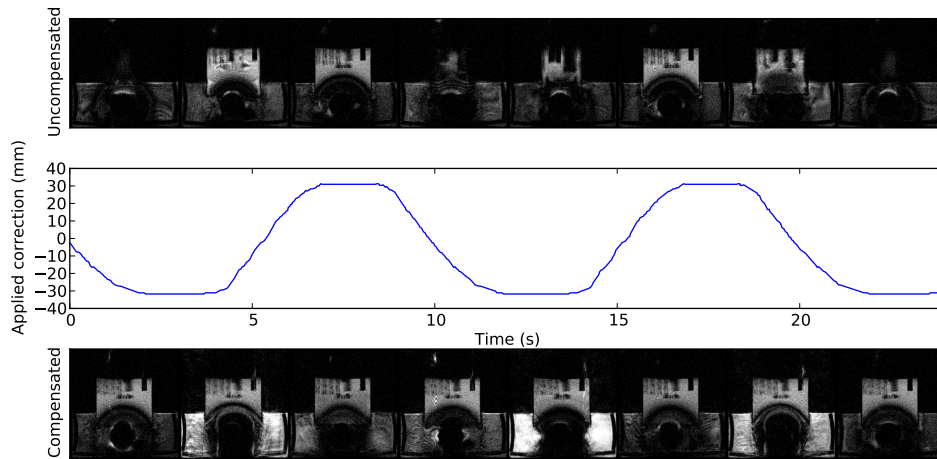
Figure 3.7: Prospective motion compensation performance for through-plane motion, structured as in Figure 3.5. The dark blooming artifact was caused by ferromagnetic components of the transducer, and the surrounding brightness variation is due to turbulence and saturation effects in the water bath.

3.57 kHz, imaging a 12 cm square coronal slice with a thickness of 3 mm at a resolution of 128×128. The slice was selected to pass through the focal spot of the heating transducer with an orientation orthogonal to its central axis. Compared to Experiment 3.5, motion compensation in this experiment is much more challenging due to this sequence's longer TR and TE, thinner slices, and extreme sensitivity to phase variation.

After waiting for 7 images without heating, 8.3 Watts of electrical power was applied for 60 seconds. This procedure was performed 16 times, while the phantom was alternately stationary or moving with ULTRACK motion compensation. The motion phantom was configured for a period of 18 seconds and a maximum velocity of 1.3 cm/sec.

Temperature changes were computed by referenceless thermometry similar to [68], with a basis consisting of a 5th-order polynomial (the background phase) and a Gaussian spot with variable amplitude and radius (the focal
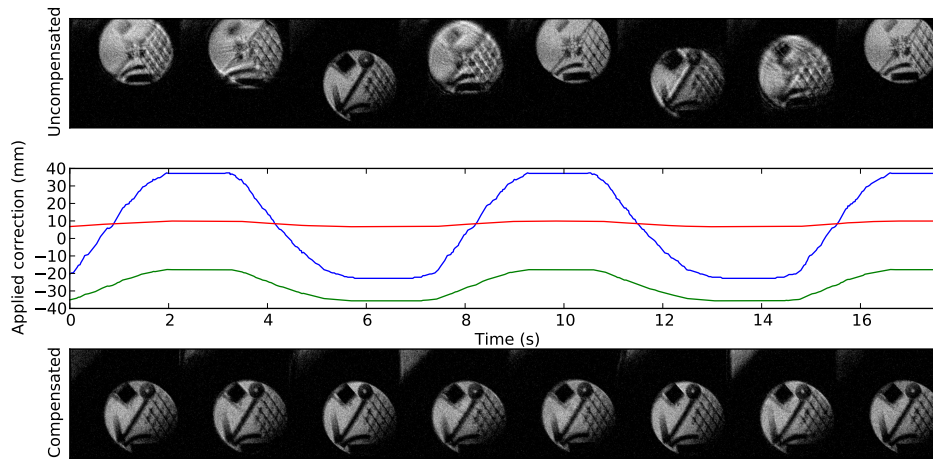
Figure 3.8: Prospective motion compensation performance for an oblique plane, structured as in Figure 3.5. In some points along the motion trajectory the oblique imaging plane includes a portion of the water bath, visible in the upper left corner of the compensated images. The three curves represent displacements along the phase, read, and slice axes.

heating). Referenceless thermometry was chosen because it removes the need for subtraction. This reduces the sensitivity to spurious drift and fluctuation, and also increases the number of independent temperature measurements that can be obtained in an experiment of fixed duration. This maximizes the opportunity to detect any systematic error induced in the temperature measurements.[20] The thermometry method is detailed in Chapter 6.

3.6.2 RESULTS

The results indicate that MRI thermometry is possible under ULTRACK guidance. As shown in Figure 3.9, the temperature curves observed over time are similar. A linear fit between the stationary and moving mean peak temperatures had a slope of 0.92 and a R^2 value of 0.96, indicating that any systematic error introduced by motion compensation is small. During

46
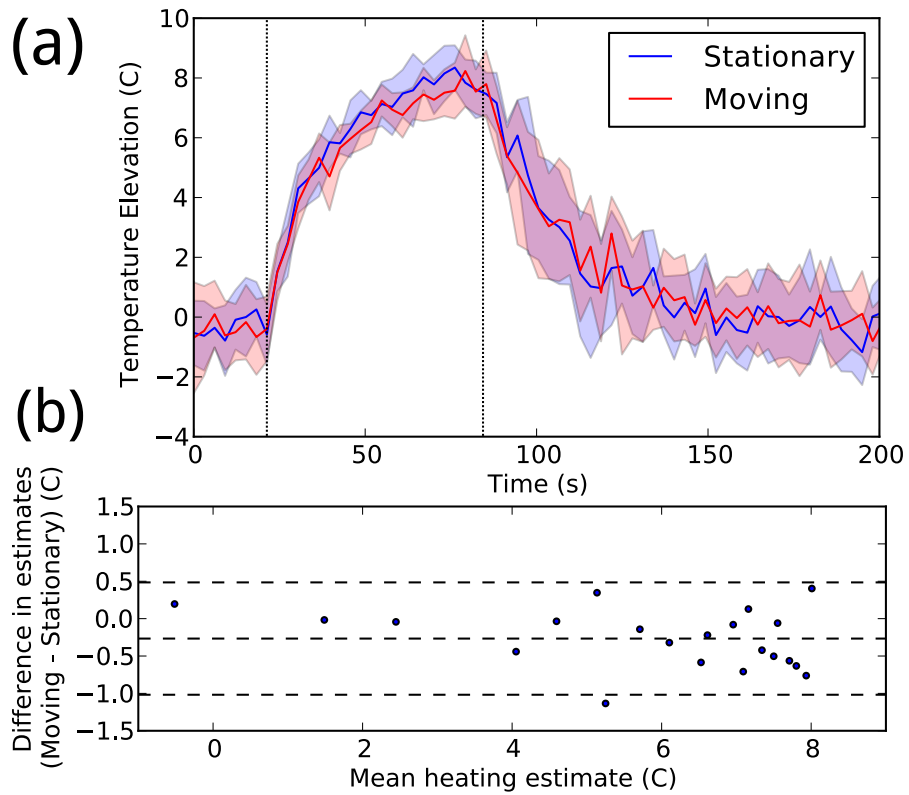
Figure 3.9: (a) Temperature data reconstructed with and without motion. In each color, the central line shows the mean peak temperature timeline for eight experiments, and the shaded area shows one standard deviation above and below. (b) Bland–Altman plot comparing the two sets of temperature measurements. Each point represents a single time during heating (from 21–81 seconds), averaged over all repetitions.

heating, the mean standard deviation in peak temperature was 0.83℃ while stationary and 0.78℃ while moving, indicating that motion compensation did not introduce significant additional temperature noise.

Figure 3.10: (a) Average temperature image at 79 seconds (near peak heating) after background phase subtraction. Both images are shown in the same temperature scale. (b) Three consecutive magnitude images of the heating phantom undergoing motion without compensation, using the same pulse sequence as used for temperature images. The black box indicates the image region in which thermometry was performed.

The geometry of the temperature image is also preserved, as seen in Figure 3.10a, although a slight broadening is evident. Without motion compensation this pulse sequence produces images, shown in Figure 3.10b, that are corrupted by severe motion artifacts and displacements that would likely preclude any thermometry.

4

On the practicality of biometric ultrasound navigation *in vivo*

IN THE PRECEDING CHAPTERS WE HAVE CONSIDERED HIGHLY IDEALIZED MODELS, both mathematical and mechanical, of the clinical scenario we wish to address. One critical idealization has been that the ultrasound patterns repeat over many respiratory cycles. The motion phantoms used in Chapter 3 are internally inert and mounted on rails, so they have extremely high

repeatability. The dynamic nature of living tissue, and the tendency of living organisms to move and shift over time, both cast doubt on the ideal of repeatability. If the repertoire of ultrasound patterns is entirely different after just a few breath cycles, then each mapping table will be effective only briefly before a new training stage is required. If frequent retraining were necessary, it could require such a large fraction of imaging time as to negate any benefits from biometric ultrasound navigation.

In this chapter we describe two experiments to discern the behavior of biometric ultrasound navigation in the absence of idealized repeatability. In the first experiment, we simulate the effect of a lateral shift, as if a patient had moved slightly to the left or right, using a phantom. In the second experiment, we perform biometric ultrasound navigation on MRI images of a live, freely breathing rabbit.

4.1 Experiment: Lateral shift sensitivity simulation

Most current methods of anatomical motion tracking, such as MRI navigators, breath bellows, ultrasound cross-correlation, and image-based ultrasound methods, cannot detect motion outside of the degrees of freedom that they are designed to track. For example, for 2D image-based methods, an anatomical shift perpendicular to the image plane is unlikely to be detected until its magnitude is far greater than motions that can be detected in-plane.

As noted in Section 3.1, the biometric tracking method relies on an assumption that the motion remains repetitive to within the length scale set by the width of the ultrasound beam. In respiratory motion compensation, the repetitive motion is predominantly along a single axis, but the patient

50

may make an unexpected motion perpendicular to this axis (i.e. *lateral*), either suddenly or by gradual drift. In some motion tracking applications (such as MRI-guided Focused Ultrasound Surgery) it is important to detect and report when even a small lateral displacement has occurred, while in other applications it is preferable that tracking along the main axis not be affected by lateral motions.

4.1.1 METHODS

To assess the sensitivity of our motion compensation system to lateral motions, we acquired data to simulate the effect of a lateral shift on the accuracy of motion tracking. As in Experiment 2.4, the phantom tracking transducer (Section 2.4.1) was mounted on a computer-controlled positioning device and aimed at the structured gel phantom (Section 3.3.1). The positioner was programmed to acquire data that closely resembled the experiments in Chapter 3, with 256 equally spaced samples in a 70 mm row. We acquired several parallel rows separated by a lateral distance of 0.2 mm. One of these rows was selected as the reference (i.e. a mapping table), and the position of each echo was estimated as equal to the position of the most similar echo in the reference row.

4.1.2 RESULTS

The results are shown in Figure 4.1.

The results show that for lateral shifts of up to 1 mm, the simulated tracking error parallel to the reference row was always less than 1 mm. At a lateral shift of 1.2 mm the parallel tracking error was less than 1 mm
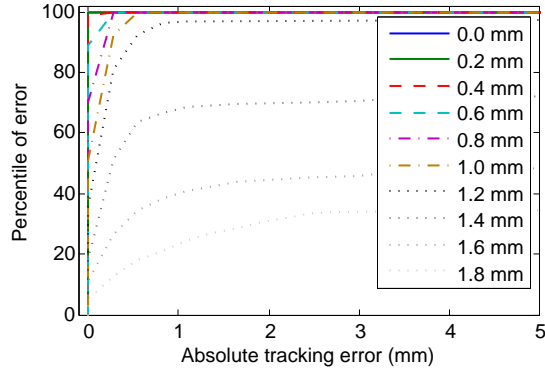
Figure 4.1: Cumulative distribution function of tracking error at 10 different lateral shifts. Motion tracking was simulated using a data set acquired using a robotic positioner.

for 96% of samples, but only 21% of samples were this accurate at a shift of 1.8 mm. These results suggest that, for this phantom, transducer, and geometrical arrangement, lateral shifts of up to 1 mm outside the reference set are unlikely to cause significant tracking errors. This result is consistent with the dissimilarity saturation observed in Figure 2.2, indicating that the sensitivity to lateral shift is also determined by the shape of the transducer's ultrasound field.

4.2 Experiment: Biometric Ultrasound Navigation in a Freely Breathing Rabbit

The animal experiments were approved by our institutional animal committee. Imaging was performed on a male rabbit (4 kg). Before the experiments, the animal was anesthetized with ketamine and xylazine, but was not intubated and was allowed to breathe freely. For optimal ultrasound coupling, the rabbit's ribcage area was shaved and coated in ultrasound gel. The an-

imal was placed head first and right lateral decubitus on a circular MRI surface coil, with the animal tracking transducer in the center underneath a 1 cm-thick ultrasound gel pad.

4.2.1 MATERIALS

4.2.1.1 MRI

Animal imaging was performed on a 3T MRI (GE), using a circular surface coil and the clinical Fast SPGR gradient-echo pulse sequence. Real-time scanner control was not available, so experiments were performed by retrieving raw acquisition data from the scanner and simulating prospective correction in purpose-built software after acquisition was complete.

4.2.1.2 ULTRASOUND

The animal tracking transducer was a broadband piston-type transducer with a nominal center frequency of 5 MHz and a diameter of 8 mm. It was manufactured from MRI-compatible materials by Imasonic SAS. It was driven by the same pulser-receiver, set identically except for a damping impedance of 25 Ω.

The transducer was driven by the same electronics as the phantom transducer (Section 2.4.1), and characterized using the same positioning system, tank, and hydrophone. For the animal tracking transducer, a maximum ultrasound intensity of 31 mJ/m^2 was measured at a focal distance of 54 mm, with a half-intensity width of 1.9 mm and length of 52 mm. At the focus, the maximum spectral energy density occurred at 5.8 MHz, with a half-power bandwidth of 3.0 MHz ($Q = 1.9$). For additional transducer characterization

53

data, see Figure 4.2.
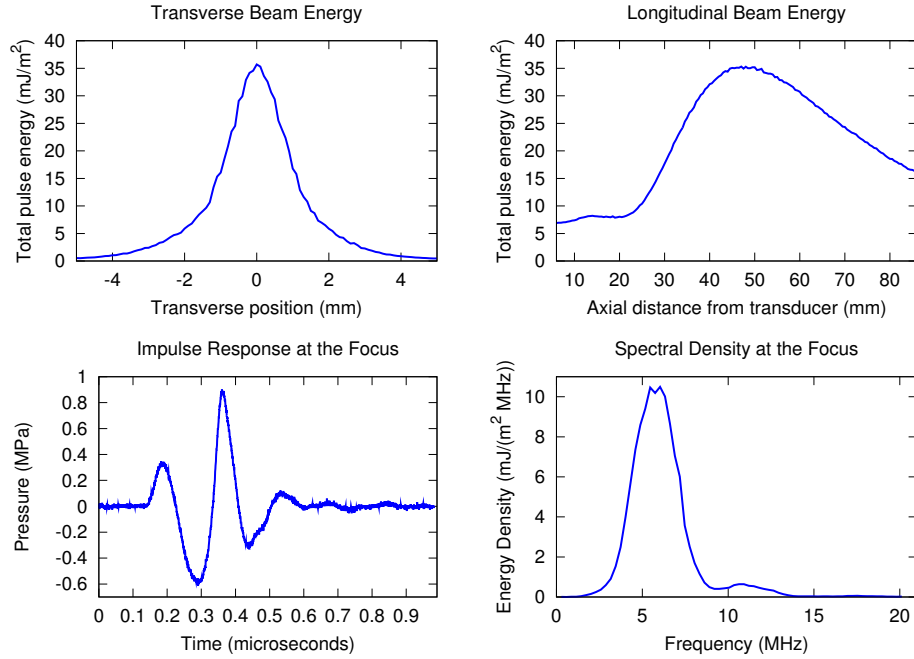
Figure 4.2: Characterization of the animal tracking transducer, as described in Section 4.2.1.2.

4.2.2 METHODS

The purpose of this experiment was to confirm whether biometric ultrasound navigation can be accurate and effective *in vivo*, by analyzing simultaneous ultrasound and MRI data acquired from a freely breathing anesthetized rabbit. For all imaging, the MRI scanner executed a Spoiled Gradient Echo sequence with a resolution of 256×256, 5 mm slice thickness, 31.25 kHz receive bandwidth, and frequency encode along the S/I axis. MRI navigator data were acquired by the same pulse sequence used for imaging, but with phase encoding disabled. The mapping tables contained 256 entries and

were computed from a training period lasting 14 seconds.

To assess the accuracy of the method *in vivo*, we acquired 143 seconds of test data consisting of paired navigator and ultrasound acquisitions, similar to the preceding training period. The images represented a coronal slice with a 17 cm square field of view and TE/TR=3.6/7.8 ms. Using the mapping table from the training period, we computed position estimates from the ultrasound data and compared them to the position computed from the navigator, which served as a reference. The paired position measures are plotted in Figure 4.3.
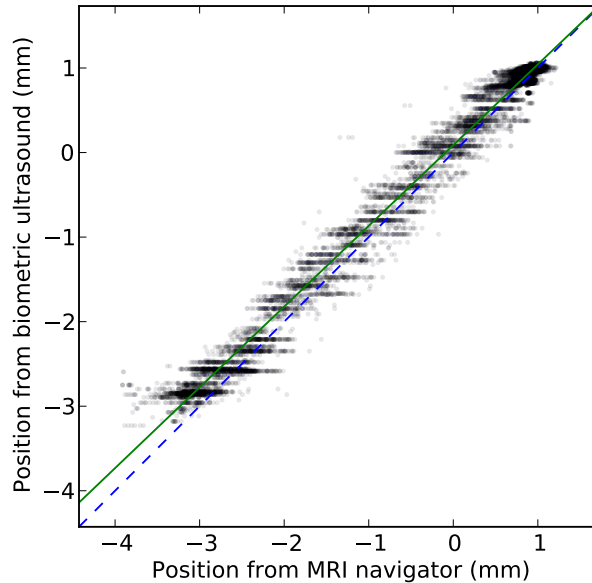
Figure 4.3: Scatter plot showing 17654 paired position measurements acquired in a rabbit during 143 seconds of free breathing. The dashed line represents perfect performance, and the solid line shows the least-squares linear fit.

The standard deviation of the reference position during the test period

55

was 1.41 mm. After subtracting the position estimate from ultrasound, the standard deviation of the residual was 0.16 mm, indicating that the biometric navigator captured 89% of the variation in the MRI navigator. A least-squares linear fit between the two position measures had a slope of 0.95 and a R^2 value of 0.98.

To assess the effectiveness of the method *in vivo*, we acquired a sagittal image series and training scan with an 11 cm field of view and TE/TR=4.1/11.5 ms. A mapping table was computed from the training scan and used to estimate the displacement at each ultrasound acquisition during the image series. Motion compensation of the MRI images was achieved retrospectively by applying a phase roll to each line of k-space to remove the displacement estimated from ultrasound. This set of MRI data, reconstructed with and without retrospective motion compensation, is visualized in Figure 4.4. The correction produced a clear improvement in image stability and reduction of motion artifacts.

These results show that biometric ultrasound navigation can provide accurate respiratory displacement estimates *in vivo*, using less than 9% of total imaging time for training. They also demonstrate that the method is effective at producing sharp, stable MRI images during free breathing.
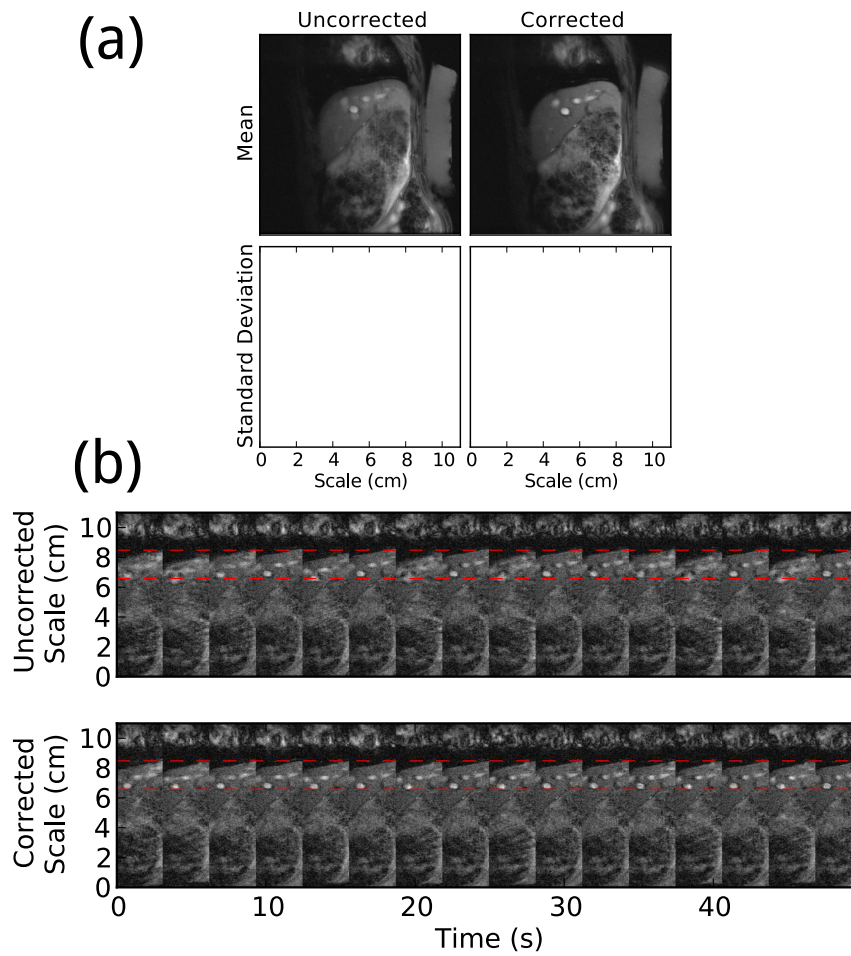
Figure 4.4: (a) Mean and standard deviation of 32 consecutive images of a freely breathing rabbit, with and without retrospective biometric ultrasound motion compensation. Each image pair is shown in the same contrast. (b) Time series showing the central section from 16 consecutive images, with and without motion compensation. Horizontal dashed lines indicate the main area of motion compensation, containing bright blood vessels just below the diaphragm. Cardiac motion artifacts at the top of each image are not corrected.

5

Biometric Bootstrapping in Pseudo-position Space

WE HAVE ASSUMED, in the preceding chapters, that a fast, accurate position monitoring system is already available as an input for any biometric navigation system. The basic biometric navigation algorithm, which has been the focus of all results so far presented, requires such a position measure as an input during the training stage.

In this chapter we consider what can be achieved using only a single ultrasound transducer oriented transverse to the direction of motion, when an external position input for training is unavailable or incomplete. Our method works by first extracting as much information as possible from the biometric signature alone, and then augmenting it with any external position input that may be available. Our results indicate that biometric navigation can provide a significant amount of information about motion patterns without requiring any additional equipment or data. To gain the benefits of spatial mapping without relying on external position data, the method operates in *pseudo-position space*, a Euclidean space that represents the abstract positional state space of the object being observed.

5.1 Algorithms for Pseudoposition mapping

The overall algorithm for motion model construction using pseudo-position mapping consists of three main stages: Pseudo-SLAM, bin reconstruction, and motion model extraction. Pseudo-SLAM operates only on the biometric signatures themselves, while the subsequent stages rely on the presence of an external position information source. It uses pairwise distance measurements to find a pseudoposition assignment pseudo-positions based solely on the biometric data. Bin reconstruction aggregates external training information corresponding to locations in pseudo-position space, in order to construct complete data for many pseudospace locations. Finally, motion model extraction analyzes and compares the aggregated training data, producing a mapping from pseudo-position space to real geometrical parameters.

5.1.1 Pseudo-SLAM

In the Basic Biometric Navigation Algorithm, an external position data source (such as an MRI imager) is used to tag every biometric signature (such as an ultrasound echo) with its corresponding position. These pairs can then be used to build up a table representing the relationship between biometric signatures and positions. Once the training is complete, the table can be used to determine the present location given a biometric signature.

In robotics, problems analogous to the computation of this table are simply called *mapping*, in the sense of cartography. An environmental data source (e.g. camera, LIDAR) is combined with an external position data source (e.g. GPS) to build up a representation of the robot's surroundings. The problem of determining the current position, given the map and an environmental measurement, is called *localization*.

Given only a series of biometric signatures and no external position data source, we may still wish to estimate a map. The problem is much harder than simple mapping, however, because the location of each biometric signature within the map is not initially known. In the field of robotics, problems of this class are termed *Simultaneous Localization and Mapping* (SLAM).[22]

In the usual conception of SLAM, the output map does not represent position in any pre-existing absolute reference space, because there is not enough information available. For example, a robot navigating an indoor maze may not be able to determine its latitude, longitude, or compass orientation. Thus, the map produced by SLAM is often regarded as degenerate up to rigid transformations.

When biometric signatures are the only available data source, we can

expect even more degeneracy in the resulting map. By the analysis in Chapter 2, it is possible to determine some information about the distance between two biometric signatures with good reliability, but there is almost no information about the direction in which this separation occurs. Even if all the distances among a set of points are known exactly, it is still impossible to distinguish between one set of points and its mirror image, so the degeneracy in the position coordinates includes not only rigid transformations but also reflections. The limited range of distance measurements on biometric signatures (a consequence of the saturation phenomenon discussed in Section 2.3) limits the accuracy of long-range measurements, so only the local structure can be reproduced reliably. As a result, the coordinate space will also be nearly degenerate up to smooth deformations that do not greatly alter the distances between nearby points. Because the resulting map space is likely to be so different from the true position space, we describe the map coordinate as a *pseudo-position*. Accordingly, we call SLAM-like problems on biometric signatures *Pseudo-SLAM*.

The pseudo-position space is a representation of the abtract phase space of the system, with the additional property that the function that maps its coordinates to real position space is smooth. There is no need for the pseudo-position and the real position to occupy a space of the same dimensionality, and indeed their dimensionalities are different in the current implementation of the algorithm.

5.1.1.1 Multi-Dimensional Scaling for Pseudo-SLAM

Our system for Pseudo-SLAM proceeds as follows. After acquiring the training data, approximately 10^4 samples, we first make pairwise similarity measurements between the biometric signatures. Our signatures take the form of ultrasound echoes, and so our similarity measurements are computed using the measure described in Equation 2.3. We express similarity in terms of an estimated distance

$$D(\boldsymbol{a}, \boldsymbol{b}) = \min_{s} \sqrt{\left(\sum_i (a_i - b_{i-s})^2\right) - n}. \tag{5.1}$$

Here n is a noise-compensation term, to account for the fact that, even if there has been no motion, successive acquisitions will not be identical. (This is the offset observed in Figure 2.2a.) The minimization over all shifts serves to discard information about small shifts parallel to the beam axis, which would otherwise overwhelm the effects of transverse displacements.

This distance measure is suitable for small displacements because the profile of the ultrasound beam is smooth and radially symmetric. Its spatial autocorrelation function is therefore approximately quadratic in a small region around the origin, which makes $D(\boldsymbol{a}, \boldsymbol{b})$ proportional to the true transverse displacement. For larger displacements, the ultrasound echoes become uncorrelated, and their observed numerical distance approaches a fixed ceiling. Therefore, we select a threshold T, lower than $D(\boldsymbol{a}, \boldsymbol{b})$ for all widely separated measurements but higher than $D(\boldsymbol{a}, \boldsymbol{b})$ for all closely spaced measurements. Such a threshold must exist for any biometric signature, and we expect it to exist in our case due to the analysis in Section 2.3. If $D(\boldsymbol{a}, \boldsymbol{b}) > T$,

then we regard the distance between $\boldsymbol{a}$ and $\boldsymbol{b}$ in pseudo-position space as unknown, but constrained to be at least T.

The function described for $D(\boldsymbol{a}, \boldsymbol{b})$ is too computationally intensive to calculate for all pairs of input biometric signatures in large training sets. An optimized implementation (Section 3.2.1) can compute 10^5 dissimilarities per second. On a typical training set of 10^4 signatures, there are $10^8/2$ distinct pairs, so computing all the dissimilarities would require about 500 seconds, which is unacceptably long. Moreover, even if the dissimilarity computation can be made faster, subsequent steps in the algorithm also take time proportional to the number of measured dissimilarities. Therefore, we avoid calculating all $10^8/2$ pairs, and instead only compute distances for a small, strategically chosen subset of pairs.

First, each signature is compared to the 60 readouts acquired before and after it. As long as the training data represents continuous motion sampled at a fast constant rate, the dissimilarity between successive signatures should be very small, leading to low apparent distance. Given only this information, the input should form a sort of chain in pseudo-position space, with a spacing representing the instantaneous velocity. Additionally, 60 evenly spaced readouts are selected as anchor points. Each anchor point is compared to all 10^4 readouts (except those to which distances have already been computed). Each anchor point tells us that a certain set of points in the chain should all be at the same location (because they have similar signatures), forcing the chain to fold back over itself. In total, about 10^6 pairwise dissimilarities are measured, or 2% of the total, although most of these are found to be greater than T.

Knowledge of pairwise distances is sufficient to reconstruct the relative positions of a set of measurements. The problem of identifying a spatial embedding of a set of points that is most nearly consistent with known pairwise distances is called Multi-Dimensional Scaling (MDS).[84]

The form of MDS used in our implementation, known as metric MDS, works by minimizing a quantity called the *stress*:

$$\mathcal{S} = \sum_{(i,j)\in P} K_{ij} \left(d(\vec{x}_i, \vec{x}_j) - D_{ij}\right)^2 \tag{5.2}$$

where P is the set of all pairs with known distances, $d(\cdot, \cdot)$ is the Euclidean distance function, $\vec{x}_i$ is the pseudo-position associated with the measurement i, D_{ij} is the known distance between samples i and j, and K_{ij} is a constant representing the confidence in the measurement of D_{ij}. If each measurement D_{ij} is subject to independent Gaussian noise with standard deviation $1/K_{ij}$, then minimizing $\mathcal{S}$ corresponds to a Maximum Likelihood Estimation (MLE). For our ultrasound measurements, we do expect D_{ij} to have an approximately Gaussian distribution, as discussed in Section 2.3. However, the MLE is only the correct Maximum *a posteriori* (MAP) estimate in a Bayesian framework if the prior on D_{ij} is uniform (i.e. flat). The flat prior is clearly a problematic choice, as it suggests that, in the absence of other information, all points are expected to be infinitely far apart.

As noted in Section 5.1.1, the distance measures available for the data set are only valid for small distances. Accordingly, we introduce a modified

64

stress:

$$
E_{ij} = \begin{cases} K_{ij} \cdot (d(\vec{x}_i, \vec{x}_j) - D_{ij})^2 & : D_{ij} \leq T \\ K_{ij} \cdot (d(\vec{x}_i, \vec{x}_j) - T)^2 & : D_{ij} > T \wedge d(\vec{x}_i, \vec{x}_j) < T \\ 0 & : \text{else} \end{cases} \tag{5.3}
$$

$$
\mathcal{S} = \sum_{(i,j) \in P} E_{ij} \tag{5.4}
$$

The new E_{ij} represents the squared distance error associated with the pair (i, j). If D_{ij} is larger than a threshold T, the pair is only regarded as contributing error when the $d(\vec{x}_i, \vec{x}_j) < T$.

The modified stress is designed to retain compatibility with the technique of *stress majorization.*[18] The standard stress function is not a polynomial function of the positions, and there is no analytic method for locating even a local minimum. Instead, one must use iterative methods. The most popular iterative methods work by, at each step, minimizing a quadratic expansion $\tilde{\mathcal{S}}$ of the stress function around the current set of pseudopositions. With the traditional metric stress function, this approximation is guaranteed to converge because $\tilde{\mathcal{S}}$ touches $\mathcal{S}$ at the current location and $\tilde{\mathcal{S}}$ is a convex upper bound on $\mathcal{S}$. These properties make $\tilde{\mathcal{S}}$ a *majorizing function.*

For the modified stress with constraint conditions, $\tilde{\mathcal{S}}$ is no longer necessarily an upper bound on $\mathcal{S}$, and so the existing convergence proof does not hold. We have designed our modified stress to retain a strong convergence guarantee, but a formal proof of this guarantee has not yet been formulated. In our implementation, minimizing $\tilde{\mathcal{S}}$ proceeds by the Conjugate Gradient method, as the corresponding matrix is both very large and very sparse.

MDS can determine the geometrical shape of the point set, but not the set's global orientation. Thus, the set of ultrasound echo distances is insufficient to determine the motion parameters of the observed organ. Our data is further insufficient because only 2% of the distances are actually computed, and of these most have only a known lower bound. Standard MDS algorithms rely on known distances between far away points in order to reconstruct large-scale structure correctly. Without known large distances, these algorithms tend to get stuck in false optima, with the set's large-scale structure folded in on itself.

One technique for determining structure from only small distances has been termed Local MDS (LMDS).[13] In LMDS, the standard stress function is supplemented by a repulsive force between every pair of points, similar to the dark energy or cosmological constant of modern general relativity. This force serves to stretch out the set's geometry, so that points are only near to each other if the known distances require them to be. Mathematically, this corresponds to a second modification to the stress:

$$E_{ij}^+ = E_{ij} - \alpha \cdot d(\vec{x}_i, \vec{x}_j)^2. \tag{5.5}$$

This term may be regarded as stretching out the points to form a smooth, low-dimensional manifold embedded in the output space, countering the tendency of noise in D_{ij} to cause the manifold to buckle and crinkle wherever a distance is locally underestimated. We experimented with local smoothness penalties on the chain of points, but found that they were too easily trapped in false minima.

One difficulty with the repulsion, as it is described in [13], is that it

66

eliminates the sparsity of the system, because E_{ij}^+ is non-zero for almost all i and j. This sparsity is crucial for computational efficiency. However, we can rescue the sparsity by reformulating the repulsion:

$$\mathcal{S}^+ = \mathcal{S} - \alpha \sum_{i,j} d(\vec{x}_i, \vec{x}_j)^2 \tag{5.6}$$

$$= \mathcal{S} - \alpha \sum_{i,j} \sum_k (x_{i,k} - x_{j,k})^2 \tag{5.7}$$

$$= \mathcal{S} - \alpha \sum_{i,j} \sum_k x_{i,k}^2 + x_{j,k}^2 - 2x_{i,k}x_{j,k} \tag{5.8}$$

$$= \mathcal{S} - \alpha \sum_i \sum_k n x_{i,k}^2 + \left(\sum_j x_{j,k}^2 \right) - 2x_{i,k} \sum_j x_{j,k} \tag{5.9}$$

$$= \mathcal{S} - \alpha \sum_k n \left(\sum_i x_{i,k}^2 \right) + n \left(\sum_j x_{j,k}^2 \right) - 2 \left(\sum_i x_{i,k} \right) \left(\sum_j x_{j,k} \right) \tag{5.10}$$

$$= \mathcal{S} - \alpha \sum_k 2n \left(\sum_i x_{i,k}^2 \right) - 2 \left(\sum_i x_{i,k} \right)^2 \tag{5.11}$$

$$= \mathcal{S} - 2\alpha n^2 \sum_k \frac{1}{n} \left(\sum_i x_{i,k}^2 \right) - \left(\frac{1}{n} \sum_i x_{i,k} \right)^2 \tag{5.12}$$

$$:= \mathcal{S} - 2\alpha n^2 \sum_k \frac{1}{n} \left(\sum_i x_{i,k}^2 \right) - \mu_k^2 \tag{5.13}$$

$$= \mathcal{S} - 2\alpha n \sum_k \sum_i x_{i,k}^2 - \mu_k^2. \tag{5.14}$$

$$= \mathcal{S} - 2\alpha n \sum_i d(\vec{x}_i, \vec{\mu})^2. \tag{5.15}$$

$$\tag{5.16}$$

This derivation proves that the repulsion stress between all pairs is equivalent

to a force that acts only between each vector $\vec{x}_i$ and the mean vector $\vec{\mu}$. Since the absolute location within pseudo-position space is arbitrary, we set $\vec{\mu} = \vec{0}$ without loss of generality. This reformulated repulsion allows $\Theta(n)$ computational complexity instead of $\Theta(n^2)$, which is important because the known distances are extremely sparse and n is on the order of 10^4.

In our implementation, we combine the repulsion of LMDS, reformulated for sparsity, with the modified stress function (Equation 5.4) to account for constraints. We call this approach Constrained LMDS.

An additional difficulty with LMDS and Constrained LMDS is the need to determine an appropriate strength for the repulsion force. A repulsive force that is stronger than the attractive forces between the measurement points can prevent the optimization from converging. If the repulsion strength exceeds some threshold, $\tilde{\mathcal{S}}$ will become nonconvex and the points will repel each other away to infinity. However, too weak a repulsion will not have the desired regularizing effect.

In Bayesian terms, the repulsion of LMDS represents an even more problematic and improper prior than the flat prior implicit in Equation 5.2. The flat prior corresponds to a probability distribution that, while not normalizable, is at least contained in the L^2 Hilbert space (i.e. square-integrable), and is representable as a normal distribution in the limit of infinite variance. This means that as long as the likelihood function is a proper probability distribution, by Bayes' theorem the posterior distribution will be proper as well. The repulsion prior, in contrast, resembles a normal distribution with *negative* variance, and there are many proper likelihood functions that, with this prior, will produce improper posterior distributions. In these cases,

the gradient of the posterior points toward configurations that are infinitely large. These situations correspond to a repulsion force that is too strong.

To ensure that the optimization does not diverge, we first determine the value of this upper bound on repulsion. In the limit of infinitely widely separated points, the constraint penalties are all zero, and we are left with

$$K_{ij}^{\infty} = K_{ij} \text{ if } D_{ij} < T \text{ else } 0 \tag{5.17}$$

$$E_{ij}^{\infty} = K_{ij}^{\infty} \cdot (d(\vec{x}_i, \vec{x}_j) - D_{ij})^2 - \alpha \cdot d(\vec{x}_i, \vec{x}_j)^2 \tag{5.18}$$

$$= (K_{ij}^{\infty} - \alpha)d(\vec{x}_i, \vec{x}_j)^2 + O(d(\vec{x}_i, \vec{x}_j)). \tag{5.19}$$

Thus, the stress is dominated by the quadratic distance term, which can be expressed in linear algebraic terms (for each pseudoposition axis separately) as

$$L_{ij} = \begin{cases} -K_{ij}^{\infty} & : i \neq j \\ \sum_k K_{ik} & : i = j \end{cases} \tag{5.20}$$

$$C_{ij} = \begin{cases} -1 & : i \neq j \\ n - 1 & : i = j \end{cases} \tag{5.21}$$

$$\mathcal{S}^{\infty} = \boldsymbol{x}^*(\mathbf{L} - \alpha\mathbf{C})\boldsymbol{x} + O(||\boldsymbol{x}||) \tag{5.22}$$

If we regard $\mathbf{K}$ as representing the weights of an undirected graph whose nodes are the measurement timepoints, then $\mathbf{L}$ is called the *graph Laplacian* matrix. Like all Laplacian matrices, the smallest eigenvalue (λ_0) is zero, corresponding to the eigenvector $[11...1]$. This represents the insensitivity of the stress to a translation of the entire point set.

The second smallest eigenvalue of $\mathbf{L}$ (λ_1) is called the *algebraic connectivity* of the graph, and the corresponding eigenvector ($\boldsymbol{e_1}$) is called the *Fiedler vector*.[8] For any eigenvector $\boldsymbol{e_k}$ of $\mathbf{L}$ and $k > 0$, we have $\mathbf{C}\boldsymbol{e_k} = n\boldsymbol{e_k}$ due to the construction of $\mathbf{C}$, and therefore

$$(\mathbf{L} - \alpha\mathbf{C})\boldsymbol{e_k} = (\lambda_k - n\alpha)\boldsymbol{e_k}. \qquad (5.23)$$

If $\mathbf{L} - \alpha\mathbf{C}$ has all positive eigenvalues (apart from $\lambda_0 = 0$), then the optimization of $\mathcal{S}^\infty$ is stable, because $\mathcal{S}^\infty$ is dominated by a convex function. However, if $\mathbf{L} - \alpha\mathbf{C}$ has any negative eigenvalue, then the optimization is unstable, because the stress maybe made arbitrarily low by increasing the amplitude of the corresponding eigenvector. Therefore, the optimization is stable if and only if $\alpha < \lambda_1/n$, i.e. the repulsion is weaker than the algebraic connectivity. In our implementation, we compute this eigenvalue using LOBPCG [48] as implemented by Scipy [44], and heuristically choose $\alpha = \lambda_1/5n$.

The optimization problem in question resembles folding a chain, representing the pseudo-position over time, back and forth on top of itself, so that appropriate sections of the string are adjacent to each other. (Any resemblance to the protein folding problem is purely coincidental.) The problem can easily be captured in local minima, preventing it from reaching the correct solution. We make two additional tweaks to the algorithm to improve the avoidance of local minima.

One way the system can be trapped in a minimum is if the string of positions is tied in a knot. This can easily happen if the initial positions are randomized, and the system can tie itself in knots even from other initial

positions. To avoid the formation of knots, we operate in a pseudo-position space with dimensionality of at least 4, because it is known from knot theory that a finite chain cannot be knotted in a Euclidean space with dimensionality greater than 3.[93]

To encourage convergence to the global optimum, we choose the initial conditions based on the graph Laplacian eigenvectors. Each of the M dimensions is assigned to one of the M smallest nonzero eigenvectors, as computed by LOBPCG. This initial condition contains no information about the distances computed between individual points, only whether a distance between each pair is known at all. Each axis is initially scaled according to a simple heuristic such that the energy associated with each dimension is equal to 1. All axes are then rescaled so that the overall scale is approximately consistent with the desired distances between points.

In summary, the Constrained LMDS algorithm is:

1. Take as input: $n \times n$ sparse matrices $\mathbf{D}$ (distance between pairs) and $\mathbf{K}$ (strength of each known pairwise connection); repulsion factor α; pseudo-space dimensionality m; distance threshold T.

2. Construct $\mathbf{K}^\infty$ containing only the attractive interactions.

3. From $\mathbf{K}^\infty$ compute the graph Laplacian matrix $\mathbf{L}$.

4. Using LOBPCG, compute the first m nonzero eigenvalues ($\boldsymbol{\lambda}$, $m \times 1$) and eigenvectors ($\mathbf{V}$, $n \times m$) of $\mathbf{L}$.

5. Compute $\mathbf{V}'$ such that $V'_{ij} = V_{ij}/\lambda_j$.

6. Compute β that minimizes $\mathcal{S}(\beta\mathbf{V}')$, considering only the attractive

71

interactions. Set $\mathbf{P_0} \leftarrow \beta \mathbf{V}'$, representing the initial pseudopositions. Set $k = 0$.

7. Compute a representation of $\tilde{\mathcal{S}}_k$, which majorizes $\mathcal{S}$ around $\mathbf{P_k}$.

8. Compute $\mathbf{P_{k+1}}$ by minimizing $\tilde{\mathcal{S}}_k$ using the method of conjugate gradients.

9. Increment k. While $k < k_{\text{limit}}$, go to step 7

5.1.1.2 Bin reconstruction

The output from Pseudo-SLAM is a mapping of acquisition time points into a generic state space of arbitrary dimension, which we call the pseudo-position space. Nearby points in this space correspond to similar positional states, and local distances are proportional to physical distance. If the only available information is the biometric signatures, then the pseudo-position is the closest we can come to a reconstruction of the true position. However, if some external position data is available, then we can determine the relationship between pseudo-position space into real space. *Bin reconstruction* is the first step in that process. It is a method by which data representing real position is assembled at a specific location in pseudo-position space.

The purpose of pseudo-position mapping is to permit position estimation from a biometric signature, even when complete position information is not available for each measured signature. For example, suppose the position information consists of N sequential types of acquisitions that must be combined to form a position estimate, and a biometric signature is available for each measurement. Then during the training period, each biometric signa-

ture only has a fraction of the data required to compute a position estimate. If significant motion occurs during the course of the N acquisitions, then their combination will not be a correct position estimate, if it is even interpretable.

Our motivating application is to use 2D Gradient-Echo MRI as the training source for biometric navigation. An $N \times N$ MRI image contains N lines of k-space acquired sequentially, often over the course of multiple seconds. Movement during this time will causing blurring and other motion artifacts, making position information difficult to discern.

To determine a mapping from pseudo-position space into true positional state, we first select a large number (40, in the current implementation) of seed points, selected from among the sample points and approximately uniformly separated in pseudo-position space. From each seed point, we synthesize a complete position measurement by selecting, for each of the N types of acquisitions, the one that is nearest in pseudo-position space. Thus, we build up a complete position measurement (e.g. an MRI image) whose components may have been acquired at widely separated times, but were all acquired in similar geometric conformations. Conversely, a single acquisition may participate in multiple nearby measurements.

In MRI, this kind of acquisition reordering and binning is termed *cine*, in the sense of building up a movie that represents a cyclical process. It is most commonly applied to the heartbeat, with the cardiac phase (determined from EKG) serving as a sort of pseudo-position.[91] Thus, for some applications it may be that these binned images are the desired output product, and no further analysis is required.

5.1.1.3 Motion model extraction

Once the binning process has produced complete position measurements corresponding to various locations in pseudo-position space, the system uses these measurements to deduce the mapping to real space. The exact mechanism for this depends on the nature of the position measurements. In some cases each binned measurement may be analyzed independently, while other methods (especially those that rely on image registration) require comparison between different reconstructed measurements. In any case, the result should be, for each complete measurement, a vector in the output position space.

Each element of the output position vector may then be tabulated with the pseudo-position coordinates of the images' seed points, forming a table that represents a pseudo-position input vector and a scalar output. Fitting each table with a low-order multivariate polynomial, and combining these polynomials into a single vector-valued function, produces a smooth mapping from pseudo-position space to the desired output position space. This mapping may be computed for each acquisition timepoint based on its pseudo-position, with the result that each biometric signature (and position acquisition) is tagged with a position estimate.

Once position estimates are available, the system may be configured to do either retrospective or prospective motion compensation. In the retrospective case, the biometric signatures are no longer relevant, and motion compensation is accomplished by performing a motion-aware reconstruction of the images used for training, using the new position estimate associated with each acquisition. This only works if the images for which motion com-

74

pensation is desired are the same ones that served as position measurements during motion model extraction. This might be true when the application requires retrospective motion compensation of dynamic images, the images also show the required motion information, and the changes to be observed are not large enough to corrupt the motion information.

For prospective compensation, the position-data acquisitions are no longer needed and may be discarded. The remaining data, in the form of paired biometric signatures and position estimates, is of exactly the form produced by the basic biometric navigation algorithm (Section 3.1). The same procedure applies here. The system summarizes the database of pairs into a mapping table that relates biometric signatures to output positions.

5.2 EXPERIMENT: RETROSPECTIVE COMPENSATION OF RIGID 2D MOTION USING THE PSEUDOPOSITION METHOD

To test the pseudoposition method described in Section 5.1, we constructed an MRI/Ultrasound motion phantom exhibiting rotational and translational motion. Unlike the linear translation motion phantom in Section 3.3.3, this phantom's motion could not be captured by a single linear MRI navigator, and fully phase-encoded images showed substantial motion artifacts because the phantom could move through a substantial fraction of its range during a single image acquisition. Without an MRI image that accurately captures the current position, it is not possible to construct a mapping table by the method used in Chapter 3.

5.2.1 Methods

The motorized motion system (Section 3.3.3) was configured to rotate the structured gel phantom (Section 3.3.1) about an off-center axis perpendicular to the coronal plane. Images were made in the coronal plane, so that there was no or minimal motion through the plane.

The MRI scanner (GE Signa 1.5T) acquired a 22×22 cm field of view with a 10 mm slice thickness at a resolution of 256×256 using a gradient-echo (SPGR) pulse sequence with TE/TR=4.2/8.9 ms. (These are the same imaging parameters as in Experiment 3.5.) A total of 10^4 readouts were acquired, corresponding to 39 images.

Each gradient echo excitation readout was accompanied by a synchronized ultrasound echo made using the phantom tracking transducer (Section 2.4.1). This ultrasound data served as input to a Pseudo-SLAM implementation operating in 4 dimensions. Because the motion was known to be reciprocal, after pseudo-SLAM the coordinates were projected onto their principal axis (the largest eigenvector), reducing the dimensionality to 1. Bin reconstruction was then performed in the linear pseudo-position space, seeded at the location of each MRI line that passes through DC in k-space. One reconstructed image was chosen as the coordinate reference, and all the others were registered to it using 2D rigid registration (3 degrees of freedom) implemented in Scipy.[44]. This resulted in a 3-component vector at each of 34 different seed locations. A mapping from pseudo-position to real geometry was then estimated by fitting each component of the vector as a second-order polynomial (quadratic) in the pseudo-position. This polynomial was then used to compute a rigid transform for each input timepoint.

To simulate some aspects of prospective motion compensation, the 10^4 input lines were summarized into a 256-entry mapping table. A table search was then used to produce a degraded transform estimate for each timepoint. The sequence of 34 images was then reconstructed again in the order in which they were acquired, with each line of k-space transformed in accordance with its position estimate. This resulted in non-uniformly sampled k-space, which was reconstructed by the DING regridding algorithm.[28]

5.2.2 RESULTS

Figure 5.1 shows the dissimilarities computed from the biometric ultrasound. Visual inspection shows a square grid rotated at 45 degrees, consistent with periodic reciprocal motion.

Figure 5.1: Matrix of dissimilarities D_{ij}, computed among the 10^4 biometric signatures. Blue entries were not computed, green entries are too dissimilar to estimate a distance, and the remaining entries show an estimated distance. For visualization, only entries with i, j multiples of 10 are shown.

The exact motion through pseudo-position space over time, and also the linear geometry of the optimized path, is shown in Figure 5.2. Images reconstructed around various seed points are shown in Figure 5.3, and the resulting map between pseudo-positions and real positions is plotted in Fig-

Figure 5.2: Pseudo-positions after optimization by Pseudo-SLAM. The upper scatter plot shows the four-dimensional distribution of points in pseudo-position space, with dimensions 3 and 4 portrayed by the red and green color channels. The lower line plot shows all 4 pseudospace vector components over time.

Figure 5.3: Cine image series, produced by bin reconstruction at each seed point.

79

ure 5.4. Finally, images motion-compensated by the estimated real positions are compared with uncompensated images in Figure 5.5

Figure 5.4: Quadratic fit between the pseudo-position coordinate of each seed point and one component of the rigid transformation.

Figure 5.5: Three consecutive images with and without motion compensation. The dark line that is sometimes visible in the motion-compensated images represents the cropping filter applied after RF chopping.[7, 422] It is dark due to local suppression of background noise, and it is visible when the transformation rotates it into the field of view.

The pseudo-position method successfully reduced the apparent motion and artifact in this experiment. Some artifacts and motion are still visible in the reconstructed images. There are many possible remaining sources of error, including inaccuracies in the position estimates and regridding reconstruction. One particularly notable source of error is gradient nonlinearity. In MRI imaging, spatial coordinates are defined by the strength of the gradient fields at each location, and where the gradients deviate from uniform linearity, the coordinate grid is distorted from rectilinear. An object moving rigidly through a nonuniform gradient field will appear to be subject to additional non-rigid deformations. The rigid 2D motion model in this experiment cannot capture such deformations, and so cannot correct for them.

She tapped out temperature. At ambient, which was no surprise. Then she squeezed the trigger for the sample function—no data—and asked for hotspots. Nothing: still at ambient temperature. Whatever this was, it was absorbing the energy and not ablating; not even warming up.

"In The Hall of the Mountain King" – Jerry Pournelle and S.M. Stirling

6

A high-precision method for referenceless Proton Resonance Frequency thermometry of focal heating

EVERY MECHANISM that couples the temperature of tissue to its MRI-visible properties creates a method of measuring temperature with MRI. Temperature dependences have been observed for at least three such properties: T_1

[61] (with correction for spin density variation by Curie's Law [37]), diffusion rate [50], and water proton resonance frequency [41]. Of these three, the last (often referred to as *PRF shift* thermometry) is the most preferred and widely used in Focused Ultrasound Surgery research and practice.[21, 55, 66]

PRF shift thermometry relies on the temperature dependence of the apparent resonant frequency of protons in water when the magnetic field B_0 is constant. In the range of interest, the shift is well approximated by a linear relationship of -0.01 ppm/°C. The shift has been known since the 1950s [74], but it was only with the advent of high-field MRI scanners that it could applied to produce spatial temperature maps *in vivo*.

The principal difficulty with PRF shift thermometry is that temperature is not the only factor that contributes to the proton resonance frequency. Spatial variations in the B_0 field, due both to the patient's susceptibility distribution and the scanner's unavoidable non-uniformity, lead to considerable differences in proton resonance frequency at different locations, even at a uniform temperature. Thus, the term PRF *shift* is used to indicate that temperature causes some additional variation from the non-uniform baseline.

In the original, and still widely used, conception of PRF shift thermometry, the baseline variation in resonance frequency was accounted for by subtraction.[41] Subtraction systems must acquire a baseline image at a reference temperature, giving the resonance frequency at each location before heating occurs. Subsequent images giving frequency can be subtracted from the reference, producing a spatial map of the temperature-induced frequency shift, and thereby an image of the temperature distribution.

Subtraction-based PRF shift thermometry is highly effective, but it can

be difficult to apply in certain circumstances. It is particularly problematic when the object in which to perform thermometry exhibits significant motion. Motion alters the magnetic field, because it represents a change in the magnetic field. In fact, even motion of objects outside the MRI field of view can still cause changes in B_0 field that significantly impair the accuracy of the temperature measurements.

To improve thermometry performance during motion, one proposed solution extends this method to multiple baseline images, each representing a different point in the motion cycle.[86] This requires the system to acquire a complete set of baseline images, representing a closely-spaced set of points that span the range of observed motions, before thermometry can begin. It also requires positions to be known both during baseline acquisition and during later temperature measurements.

This technique's requirement of extensive baseline acquisition and continuous position monitoring has motivated the development of alternative solutions for PRF shift thermometry during motion. One class of solutions is termed *referenceless* thermometry because they do not require the user to acquire a separate reference image. Instead, they rely on assumptions about the shape of the magnetic field.

The original formulation of referenceless thermometry works by making two approximate assumptions: temperature changes are confined to a known region of the image (the thermometry region), and the baseline proton resonance frequency (i.e. the local magnetic field strength) is a polynomial of known order (i.e. smooth) within a larger region that includes the thermometry region.[68] With these assumptions, each image contains enough

information to compute a surrogate reference, by removing the portion of the image in which heating occurs and replacing it by a polynomial fit to the surrounding area. The difference between the polynomial fit and the observed frequency provides the temperature estimate. In Bayesian terms, this technique adopts a strong prior assumption about the nature of the background field in each image, but a very weak prior on the foreground heating (just a region constraint).

In some cases we know more about the expected pattern of heating, and so we can employ stronger assumptions. Many heating methods, such as laser and focused ultrasound, produce a heating pattern that is known to be small and localized (pointlike), at least in one plane. A newer method, reweighted-L1 referenceless thermometry, takes advantage of this knowledge by replacing the region constraint prior with a sparsity prior.[31] This method assumes that most of the image shows no temperature change at all, but does not constrain the rare locations that do show a change. By employing a stronger prior than standard referenceless thermometry, this method achieves higher temperature-to- noise ratio (TNR) than standard referenceless thermometry. However, its stronger priors also restrict the range of applications in which it will provide correct output; broad heat distributions may not be detected at all.

For the purpose of validating motion compensated thermometry, we have introduced a form of referenceless thermometry that resembles the reweighted-L1 approach, but uses an even stronger prior on the foreground heating. Our prior assumes that the true heating pattern is a symmetric Gaussian spot centered on a known location. This model is motivated by

86

the central limit theorem as applied to our application (focused ultrasound surgery). When imaging an ultrasound heating focus in a plane orthogonal to the beam axis, the observed heating spot is described by a convolution of the ultrasound focus pattern, heat conduction spreading, and MRI point spread function, centered on the ultrasound transducer's focal point. (According to simple models, they are respectively an Airy disc [14], Gaussian, and sinc function [7].) The convolution of several such symmetric functions with similar variances is well approximated by a Gaussian function.

6.1 Mathematical formulation of Gaussian spot thermometry

For any fixed standard deviation (and polynomial order), the Gaussian spot and polynomial background form a linearly independent basis for the PRF. For simplicity of analysis, we model the noise in the PRF as Gaussian, which becomes exact in the limit of high SNR. With this signal, we have a forward model for the measured PRF $\hat{f}$ at each location i in the image:

$$f_i = \sum_j B_{ij} c_j, \tag{6.1}$$

$$dP[\hat{f}_i = x] = \frac{1}{\sigma_i \sqrt{2\pi}} \exp\left(\frac{-(x - f_i)^2}{2\sigma_i^2}\right) dx. \tag{6.2}$$

Given a basis set B consisting of the polynomial background and Gaussian foreground, the Maximum Likelihood Estimate for the coefficients c will be

the one that minimizes an objective function of the form

$$F(c) = \sum_i \frac{1}{\sigma_i^2} \left(\hat{f}_i - \sum_j B_{ij}c_j \right)^2 \tag{6.3}$$

or in matrix terms,

$$F(c) = ||\boldsymbol{w}^*(\hat{\boldsymbol{f}} - \mathbf{B}\boldsymbol{c})||^2 \tag{6.4}$$

where $w_i = 1/\sigma_i$ is the weight assigned to location i.

Common imaging sequences used for MRI thermometry are single-echo sequences, meaning that they produce an image consisting of one real and imaginary value for each location in the image.[7] The phase of each value v is computed as

$$\phi_i = \arctan(\Im(v_i), \Re(v_i)). \tag{6.5}$$

MRI images are contaminated by a significant amount of random error, which is usually well-modeled as independent and identically distributed Gaussian noise. However, the amount of error in the phase depends on the magnitude of the value. If one pixel has a small magnitude and another has a large magnitude, the same amount of noise may overwhelm the phase of the former but have no effect on the latter. An exact accounting of the error probability induced in the phase is complicated due to its nonlinear relationship with the input, but a simple geometrical argument gives the amount of phase noise in the limit of magnitudes $|v|$ much greater than the

absolute noise level n: [33]

$$\sigma_i = \frac{n}{|v_i|} \tag{6.6}$$

$$w_i = \frac{|v_i|}{n}. \tag{6.7}$$

The overall scale of $\boldsymbol{w}$ does not affect the location of the minimum of F, so for the purposes of the optimization we can ignore the value of n, or equivalently, set it to 1.

As a matter of physics, the phase of a pixel is controlled by

$$\phi_i = f_i \cdot \text{TE} \bmod 2\pi \tag{6.8}$$

For a fixed TE, several different values of f may produce the same values of ϕ, and two close values of f may produce ϕ values that differ by 2π. This phenomenon is termed *phase wrapping* in the literature. To maintain the problem's simple linear structure despite the nonlinear phase wrapping, we employ the finite-difference method also used in [31]. This solution works by applying the fit to the horizontal and vertical phase differences between neighboring pixels, which are wrapped into the range $[-\pi, \pi)$:

$$\mathbf{D} = \begin{bmatrix} \mathbf{D}_x \\ \mathbf{D}_y \end{bmatrix} \tag{6.9}$$

$$\phi' = \text{wrap}\left(\mathbf{D}\phi\right). \tag{6.10}$$

where D_x and D_y are the horizontal and vertical first difference operators. As long as neighboring pixels do not have phase differences larger than π, this

method produces a correct value for ϕ'. To incorporate this into the objective function, we move the entire computation into the difference space:

$$\boldsymbol{f}' = \phi'/\text{TE}, \tag{6.11}$$

$$\mathbf{B}' = \mathbf{D}\mathbf{B}, \tag{6.12}$$

$$F'(c) = |\boldsymbol{w}'^*(\hat{\boldsymbol{f}}' - \mathbf{B}'\boldsymbol{c})|^2. \tag{6.13}$$

Choosing a correct weighting $\boldsymbol{w}'$ is not as simple. Our current implementation uses a harmonic mean of the two contributing weights, e.g. for horizontal differences:

$$w_i' = \frac{1}{\frac{1}{w_{\text{left}(i)}} + \frac{1}{w_{\text{right}(i)}}}. \tag{6.14}$$

This weighting serves principally to ensure that a difference receives low weighting if either of its inputs is very noisy. In future implementations a harmonic Euclidean mean might be preferable

$$w_i' = \frac{1}{\sqrt{\frac{1}{w_{\text{left}(i)}^2} + \frac{1}{w_{\text{right}(i)}^2}}} \tag{6.15}$$

because it would match the way that variances of Gaussians add:

$$\sigma_{a+b}^2 = \sigma_a^2 + \sigma_b^2. \tag{6.16}$$

Because the difference operator $\mathbf{D}$ removes any spatial-DC offset in $\hat{\boldsymbol{f}}$, minimizing F' does not produce a correct value for the DC component. Specifically, suppose we compute a residual phase

$$\boldsymbol{r} = \text{wrap}(\hat{\boldsymbol{f}} - \mathbf{B}\boldsymbol{c}'). \tag{6.17}$$

90

Then where c' minimizes F', r will still contain an arbitrary constant offset.

To ease fixing this problem, we take one small novel step of structuring the rows of $\mathbf{B}$ as products of horizontal and vertical Laguerre polynomials, which are orthonormal on the square region of interest. The DC offset is then represented exclusively by the first row of $\mathbf{B}$, and so we may fix the offset by simply computing the appropriate weighted mean of r:

$$c_{\mathrm{DC}} = \frac{w^* r}{w^* \mathbf{1}} \tag{6.18}$$

A more sophisticated approach, not reliant on the orthogonal polynomial structure, could resolve the DC problem by a second optimization on a new objective function

$$G(\Delta c) = ||w^*(r - \mathbf{B}\Delta c)||^2, \tag{6.19}$$

resulting in a final estimated value of $c' + \Delta c$. This method would require more computation, but has the additional advantage of canceling out any noise amplification at high spatial frequencies, created by the finite difference operator. This effectively uses the image model once for phase unwrapping, and then again for actual fitting.

In previous referenceless thermometry techniques, the output of the optimization essentially resembles an image, providing temperature at each pixel location. A user who wanted to compute, for example, the peak temperature in the image would have to search the output image in order to locate the maximum value. This method is different because it incorporates an explicit, low-dimensional model of the foreground heating: the amplitude and width of a Gaussian spot. In our implementation we set the maximum value

of the Gaussian basis function to 1 by convention, so that the maximum foreground temperature rise is given directly by its coefficient in $\boldsymbol{c}$.

So far we have described the optimization process when the width of the heating spot is known. However, in our images we observe heating spots whose apparent width grows and changes over time. Thus, we run this optimization for a range of widths.

The natural way to select a width, in the absence of prior information about widths, is to select the width that produced the best fit, i.e. the lowest objective function value or maximum likelihood. However, when there is little or no actual heating, this approach produces problematic results. Empirically, we found that in the absence of heating, the system always selected either the maximum or minimum allowed width. We conceptualize this as fitting to either the smooth background B_0 variations or to a localized phase noise. Either way, the foreground Gaussian is not fitting any actual foreground heating, and its coefficient provides no useful information about the amount of heating.

To remedy this problem, we adopted a modified search criterion, based on the observation that if a Gaussian spot is actually present within the image, a basis function of the correct width produces a better fit than one that is too large or too small. Instead of simply choosing the radius ρ that produced the smallest fitting error, we choose the value that produced the lowest local minimum along ρ of fitting error. If no local minimum is present, we conclude that no focal heating could be detected, and return a foreground heating amplitude of zero. An example of this search is plotted in Figure 6.1.

Figure 6.1: A plot of the Gaussian spot radius search for one image acquired during Experiment 3.6. The best fit was achieved at the minimum radius (lower arrow), but our modified search criterion instead selected the radius corresponding to the unique local minimum at a radius of 4.2 pixels (upper arrow).

This modification to the search has the effect of generating a *dead zone*, i.e. a threshold of heating below which the deviation cannot be detected. The size of the dead zone varies with the width of the heating pattern, becoming larger as the heating spot becomes wide enough to be represented well by the polynomial basis.

Figure 6.2: An illustration of the internal structure of each thermometry method, using a single data set. Each column is presented in a fixed contrast.

An explanatory illustration of the thermometry methods is presented in Figure 6.2. The data are derived from a single temperature image from Experiment 3.6, acquired using the heating phantom (Section 3.3.2) in the absence of motion. The image is 3 cm wide, and all the referenceless methods were configured to use a 5th-order polynomial background phase fit.

A comparison of the methods on all the stationary data from Experiment 3.6 is shown in Figure 6.3. The results suggest that this method is more reproducible than standard thermometry, i.e. it appears to have higher temperature-to-noise ratio for the same input image. This is not surprising, given the strong prior assumption about the spatial structure of heating, which creates an effect similar to local averaging. The reduction in variability is especially strong during active heating, when the temperature spot is highly localized.

Figure 6.3: A comparison of the four different thermometry methods on the stationary data from Experiment 3.6. The band around each curve is half a standard deviation wide, in order to illustrate the relative variability of each method.

All of the referenceless methods show a systematic underestimation of the temperature change, especially in the cooldown period when the heat has spread out across length scales comparable to the background field varia-

tions. This method's underestimation appears to be no worse than that of existing methods.

In general, referenceless thermometry has lower sensitivity to spurious drift and fluctuation than subtraction methods.[20] For experiments that aim to detect small systematic temperature difference between two groups of images, and in which this method's geometrical assumptions are valid, this method's low noise sensitivity may make it preferable to the other referenceless methods.

7

A method for producing optimal Kalman-like filters for system inversion

KALMAN FILTERS ARE UBIQUITOUS in control systems engineering, where they convert past measurements of a system into an ideal estimate of its current and future states. The traditional construction of a Kalman filter requires a physics model, a measurement model, and a noise model. The physics model describes how the future state of the system is determined

97

by its present state, the measurement model describes how measured values are related to the physical values, and the noise model describes the sources of uncertainty in the physics and measurement. The Kalman filter requires that the physics and measurement models be linear operators, and the noise models be zero-mean Gaussian noise that is stationary and independent in time.[45]

Mathematically, Kalman filters are easiest to model in discrete time, where all events take place at multiples of some timestep Δt. In such a formulation, the physics model operates by the recurrence

$$x_k = \mathbf{F}_k x_{k-1} + w_k \tag{7.1}$$

where x_k is the state of the system at time $k\Delta t$, $\mathbf{F}_k$ is the linear physics model, and w_k is random Gaussian noise with covariance matrix $\mathbf{Q}_k$ as specified by the noise model. The matrices $\mathbf{F}$ and $\mathbf{Q}$ are given with subscripts because models are permitted in which these matrices change over time, although the change cannot depend on the values of x. Some versions of the Kalman filter also include a term representing the effect of outside control forces, but in this case we prefer to model systems over which we have no control.

The measurement model is similarly structured:

$$z_k = \mathbf{H_k} x_k + v_k \tag{7.2}$$

where z_k is the measurement made of x_k, $\mathbf{H_k}$ is the linear measurement model, and v_k is random Gaussian noise with covariance matrix $\mathbf{R_k}$. The

98

variation of **H** over time can be used to model situations in which different measurements are made at different times, or even times at which no measurement is made at all ($\mathbf{H} = \mathbf{0}$).

Kalman filters work by maintaining $\hat{\boldsymbol{x}}$, a running estimate of $\boldsymbol{x}$, and also **P**, a noise covariance representing the uncertainty in $\hat{\boldsymbol{x}}$. At each timestep these two variables are updated:

$$\mathbf{P_k} = \mathbf{F_k}(\mathbf{I} - \mathbf{P_k}\mathbf{H_k}^T(\mathbf{H_k}\mathbf{P_k}\mathbf{H_k}^T + \mathbf{R_k})^{-1}\mathbf{H_k})\mathbf{P_{k-1}}\mathbf{F_k}^T + \mathbf{Q_k}, \qquad (7.3)$$

$$\hat{\boldsymbol{x}}_{\boldsymbol{k}} = \mathbf{F_k}\hat{\boldsymbol{x}}_{\boldsymbol{k-1}} + \mathbf{P_k}\mathbf{H_k}^T(\mathbf{H_k}\mathbf{P_k}\mathbf{H_k}^T + \mathbf{R_k})^{-1}(\boldsymbol{z_k} - \mathbf{H_k}\mathbf{F_k}\hat{\boldsymbol{x}}_{\boldsymbol{k-1}}). \quad (7.4)$$

If all three models are correct, then $\hat{\boldsymbol{x}}_{\boldsymbol{k}}$ is the best possible estimate of the current state of the system given the available data.

Our interest in Kalman filters is motivated by a delay compensation problem. The motion compensation system described in Chapter 3 produces a series of position measurements that indicate the present position of a target that is being tracked. The measurements are acceptable for the required use case, but they are corrupted by a certain amount of noise, and are subject to a delay in transmission. We sought to compensate for this delay while also suppressing noise to the extent possible.

Our application represents a very special case of the general Kalman filter, because our system is metronomic (making measurements at a fixed interval) and not changing in time. The natural model for the measurement system therefore has **F**, **H**, **Q**, and **R** all time-independent constant matrices. This means that the recurrence relation for **P** will approach a constant asymptote $\mathbf{P}_\infty$ except in pathological cases, such as when the physics model **F** contain subsystems that do not interact with the measurement.

Once **P** has reached its asymptotic, optimal value, the recurrence relation on $\hat{\boldsymbol{x}}$ can be reorganized to have the form

$$\hat{\boldsymbol{x}}_{\boldsymbol{k}} = \mathbf{A}\hat{\boldsymbol{x}}_{\boldsymbol{k-1}} - \mathbf{K}\boldsymbol{z}_{\boldsymbol{k-1}} \tag{7.5}$$

where **A** and **K** are time-independent matrices that each depend on all of **F**, **H**, **Q** and **R**

In the preceding description of a Kalman filter we have assumed that a complete model of the system, including its random behaviors, is known prior to the construction of the filter, but no mention has been made of how these parameters can be determined. In fact, for our application it is of no importance whether the model parameters correctly describe the system. We are only interested in how accurately the model predicts future measurements; it is a black box, in the sense that we are concerned only with the relationship between its input and output.

Therefore, rather than build an *a priori* model of the system, we may instead deduce such a model from an example dataset. To do so, we can define an objective function

$$E(\mathbf{A}, \mathbf{K}, \mathbf{H}) = \sum_{k} |\mathbf{H}\hat{\boldsymbol{x}}_{\boldsymbol{k}} - \boldsymbol{z}_{\boldsymbol{k}}|^2 \tag{7.6}$$

representing the squared error between the predicted measurement values and the actual measurement values. Choosing matrices that minimize E produces a predictive filter that best matches the observed data. This filter represents the realization of any of a continuum of Kalman filters that have different system models but produce the same predictions given the same

measurements.

So far we have described our goal as delay compensation, to predict ahead by a time span equal to the amount of lag present in our system, so that position estimates are correct for the time at which they arrive, instead of for the time at which they were computed. However, delay compensation is only part of our aim. Our larger goal is to transform the output of the position estimation in whatever manner is necessary in order for it to match the position observed by the position receiver (in our case, an MRI machine) at the time that the position estimate arrives. As a technical matter, this is greatly simplified by the fact that each MRI position measurement is labeled, by the RTHawk framework [72], with the position estimate that was available at that time.

We may model the position receiver as a new measurement matrix $\mathbf{G}$, and produce an optimal filter by minimizing $E(\mathbf{A}, \mathbf{K}, \mathbf{G})$. The resulting filter is computed by the same recurrence (Equation 7.5), but the position estimate for the receiver is given by $\boldsymbol{z'} = \mathbf{G}\hat{\boldsymbol{x}}$. We term this arrangement a *Kalman cross-filter* because it serves not to predict the behavior of a single measurement system, but to predict the output of a second measurement system given the first as input.

Our application is especially simple because each position measurement vector $\boldsymbol{z}$ is actually a scalar value (z), so $\mathbf{K}$ is a column vector ($\boldsymbol{K}$). The entire predictive model is represented by $\mathbf{A}$, $\boldsymbol{K}$, and $\mathbf{G}$. If $\boldsymbol{x}$ has N entries, then in this representation the model has $N^2 + 2N$ degrees of freedom (compared to $1.5N^2 + 1.5N + 1$ in the original model). We started with first-order prediction filters, i.e. $N = 2$, which require 8 degrees of free-

dom. We optimized this model using a standard nonlinear optimization package to minimize E, discarding the first 10% of samples to give $\vec{\hat{x}}$ time to converge.[44] This approach produced filters that reduced E, but the numerical optimization process was exceedingly slow, and would sometimes fail to converge. One particular difficulty was that the optimizer would sometimes attempt to evaluate E over values of $\mathbf{A}$ that had an eigenvalue of greater than unity magnitude, leading to an exponential divergence behavior that could trigger floating-point overflow. To avoid this behavior, our implementation required an explicit computation of the eigenvalues of $\mathbf{A}$ before each computation of E.

To improve the performance of our optimization, we searched for a partial analytic solution that would improve the speed and reliability of the search. By matrix algebra, we may decompose the square matrix $\mathbf{A}$ into a product of complex square matrices $\mathbf{VDV}^{-1}$, with $\mathbf{V}$ unitary and $\mathbf{D}$ diagonal. If we define $\boldsymbol{y} = \mathbf{V}^{-1}\boldsymbol{x}$ as an alternative (but equivalent) representation of the internal state, then

$$\hat{\boldsymbol{y}}_{\boldsymbol{k}} = \mathbf{D}\hat{\boldsymbol{y}}_{\boldsymbol{k-1}} - \boldsymbol{L}z_{k-1} \tag{7.7}$$

where $\boldsymbol{L} = \mathbf{V}^{-1}\boldsymbol{K}$. Because $\mathbf{D}$ is diagonal, this is equivalent to a set of independent scalar first-order difference equations of the form

$$\gamma_k = d\gamma_{k-1} - lz_{k-1} \tag{7.8}$$

with γ, d, and l elements of $\hat{\boldsymbol{y}}$, $\mathbf{D}$, and $\boldsymbol{L}$. If we define

$$\sigma = \gamma/l \tag{7.9}$$

then we may trivially rewrite this equation as

$$\sigma_k = d\sigma_{k-1} - z_{k-1}, \tag{7.10}$$

A first-order difference equation with only a single parameter. If σ represents an element of a new vector $\boldsymbol{s}$, then $\boldsymbol{x}$ may be reconstituted by

$$\boldsymbol{x} = \mathbf{V}\boldsymbol{y} = \mathbf{V}\mathrm{diag}(\boldsymbol{L})\boldsymbol{s} \tag{7.11}$$

and therefore our desired cross-filtered output is given by

$$z' = \mathbf{G}\mathbf{V}\mathrm{diag}(\boldsymbol{L})\boldsymbol{s} = \mathbf{B}\boldsymbol{s} \tag{7.12}$$

where $\mathbf{B} = \mathbf{G}\mathbf{V}\mathrm{diag}(\boldsymbol{L})$ is a row vector if z' is scalar.

This formulation has reduced the number of apparent parameters from N^2+2N in $\mathbf{A}$, $\boldsymbol{K}$, and $\mathbf{G}$ to just $2N$ in $\mathbf{B}$ and $\mathbf{D}$, although these matrices may contain complex values. To further shrink the search space, suppose that $\mathbf{D}$ is fixed. Then $\boldsymbol{s_k}$ is also fixed, and all that remains is to find value of $\mathbf{B}$ that minimizes the sum of $|z'_k - \mathbf{B}\boldsymbol{s_k}|^2$. This is a linear least-squares problem, so it has a fast, straightforward analytic solution. Thus, a nonlinear optimization is only required over the contents of $\mathbf{D}$, representing N parameters. Because these parameters are exactly the eigenvalues of $\mathbf{A}$, we may ensure system stability by simply limiting the search domain.

Our present implementation is structured in exactly this fashion, with the additional simplification that $\mathbf{D}$ must contain only real values. This is equivalent to excluding underdamped (i.e. oscillatory) filters, a limitation

that we view as acceptable for our application.

7.1 Results

On a test data set, we computed optimized causal Kalman cross-filters of varying order N, and compared the RMS error to that achieved by the unfiltered data. The results are presented in Table 7.1.

Comparison	RMS error (mm)
Original	0.582359
Retrospective	0.371393
$N = 1$	0.565738
$N = 2$	0.383132
$N = 3$	0.382216
$N = 4$	0.381956
$N = 5$	0.381907

Table 7.1: RMS error produced by the different filtering methods on a test data set from Experiment 3.4.

These results show that a causal predictive filter with $N = 2$ captures 94% of the error reduction achievable by retrospective delay compensation. Increasing the model order beyond 2 further improved the match, but by amounts too small to be useful. The filter impulse responses are plotted in Figure 7.1. Note that the impulse responses all go below zero, and those with $N > 1$ remain there for an extended period.

It is possible that the numerical optimization was trapped in a local minimum, and that the global minimum would represent a more significant im-

Figure 7.1: Impulse response for the optimal filter discovered for the range of model orders that were tested.

Figure 7.2: Error achieved on the test data set for a range of filters with model order 2.

provement. The data in Figure 7.2 are somewhat reassuring in this regard, as they show no false minima on our test data for $N = 2$. However, higher dimensions may have more complicated optimization landscapes.

Future work on this topic should include additional theoretical analysis, such as determining if an analytic formula for ∇E exists, as this would further accelerate the numerical optimization. It would also be valuable to know how many nondegenerate local minima exist, and under what conditions. Finally, adding support for complex eigenvalues to the numerical optimization would greatly improve the generality of the technique.

All tasks, great or small, were of equal importance in the end.

     Too Many Curses – Alex Lee Martinez

8

Conclusions

ULTRASOUND BIOMETRIC NAVIGATION APPEARS TO BE A PROMISING TECHNIQUE for respiratory motion compensation and position monitoring. Possible applications include not only dynamic MRI (interventional and diagnostic) but also other imaging modalities, and even non-imaging interventions. Using relatively inexpensive hardware and simple software, we have implemented the technique and demonstrated that it achieves excellent artifact reduction in mechanical and animal models of human respiratory motion.

The accuracy is sufficient to enable temperature monitoring in a moving object, modeling the requirements of MRI-guided focused ultrasound surgery during free breathing. The speed, accuracy, and reliability of the technique may also be sufficient to serve as a safety-critical position monitor during medical interventions.

8.1 Limitations and Future Work

As described here, biometric ultrasound makes no use of the information conveyed by longitudinal shifts visible in the echo signals. This places it in stark contrast to methods such as the multiple-transducer system of [64], which derive position information exclusively from these shifts. It may be that the best tracking system would be one that uses a combination of these techniques, instead of ignoring either the longitudinal or transverse motion information provided by each transducer.

The ultrasound properties of human tissues are complex, and neither phantom nor animal experiments mimic them rigorously. When used in humans, ULTRACK may require a more sophisticated dissimilarity function. A suitable function might be derived from the displacement tracking techniques developed for ultrasound elastography and Doppler imaging.[87]

Biometric ultrasound, in all its variations, relies especially on the assumption that the ultrasound properties of a volume of tissue are highly repeatable over the course of an imaging session. To the extent that this assumption holds true, biometric navigation can function during arbitrary repetitive motions, even if the pattern includes nonrigid deformations that alter the speckle pattern (although only the motion components visible in the training

data can be reported, and some imaging modalities may not be able to make effective use of all the motion information that biometric navigation can produce). However, if the ultrasound appearance of the tissue is not stable over many breath cycles, the method may require such frequent re-training that it cannot offer a significant benefit.

The animal experiments show stable tracking performance over a period 10 times longer than the training period, indicating that the repeatability of ultrasound appearance and breathing patterns may be sufficiently high for this technique to function as desired *in vivo*. However, human breathing patterns, and the motions they produce, are complex.[9] They may vary from cycle to cycle, or drift over longer timescales, to a greater degree than was observed in our animal experiment. Stimuli experienced during an intervention such as focused ultrasound surgery might cause patients to alter their breathing patterns. Some applications of the technique would require the patients to be awake during the procedure, creating the possibility of voluntary motions that are not well modeled by the breathing of an anesthetized rabbit.

Motions outside of the training data, due to a gross patient shift or other unexpected motion, would result in ultrasound echoes that do not imply any position because they do not resemble any entry in the mapping table, effectively rendering the training invalid. Extensions to ULTRACK, such as occasional retraining when drift is detected or continuous adaptation of the mapping table, may help to improve performance in real use, but such extensions have not yet been tested.

High sensitivity to lateral or unexpected motions is not appropriate for

all imaging applications, but it is a valuable safety feature in focused ultrasound surgery, where it is crucial to detect any unexpected displacement instantly. The transducers used throughout this study were designed for high sensitivity to lateral shifts, consistent with their focal width of 2 mm. For applications in which lower sensitivity is desirable, we may be able to broaden the focus by changing the transducer's geometry. For example, elongated rectangular transducers might provide reduced sensitivity along their long axis, appropriate for applications where small displacements along that axis are not of interest.

The motion of the human liver during breathing is not exactly a linear shift, so motion compensation in the liver may benefit from a position model with more degrees of freedom.[70] Methods for motion compensation of multi-dimensional position models are discussed in Chapter 5, but these techniques have not yet been integrated into ULTRACK, nor otherwise demonstrated in a prospective configuration. The mathematical complexity of these advanced methods will demand extensive real-world testing to determine the system's failure modes, and how to detect failure.

8.2 CLOSING THOUGHTS

One core question remains unanswered: will biometric ultrasound navigation really work on humans? A satisfactory answer to this question will, at a minimum, require a trial in human volunteers to verify that human breathing motion provides the required repeatability.

Even if a first test in humans provides positive results, a long and difficult road remains for widespread use of any new medical device. It must outclass

every alternative, not just in technical superiority but in a hazier combination of clinical utility, convenience, comfort, and cost. It must overcome barriers of institutional inertia and personal familiarity that favor existing solutions. It will need industrial support for manufacturing, marketing, and regulatory approval. In short, it is no sure thing.

References

[1] John G. Abbott and F.L. Thurstone. Acoustic speckle: Theory and experimental analysis. *Ultrasonic Imaging*, 1(4):303–324, October 1979. ISSN 0161-7346. doi: 10.1016/0161-7346(79)90024-5. URL `http://www.sciencedirect.com/science/article/pii/0161734679900245`.

[2] Robert J. Adler. *The geometry of random fields*. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons Inc, Chichester, New York, Brisbane, Toronto, 1981. ISBN 9780471278443.

[3] Andrew B. Holbrook, Charles L. Dumoulin, Juan M. Santos, Yoav Medan, and Kim Butts Pauly. Real time respiration based steering for high intensity focused ultrasound in the liver. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, volume 19, page 525, Montreal - Quebec, Canada, May 2011. International Society for Magnetic Resonance in Medicine. URL `http://submissions.miracd.com/ismrm2011/proceedings/files/525.pdf`.

[4] Vincent Auboiroux, Lorena Petrusca, Magalie Viallon, Thomas Goget, Christoph D Becker, and Rares Salomir. Ultrasonography-based 2D motion-compensated HIFU sonication integrated with reference-free MR temperature monitoring: a feasibility study ex vivo. *Physics in Medicine and Biology*, 57(10):N159–N171, May 2012. ISSN 0031-9155, 1361-6560. doi: 10.1088/0031-9155/57/10/N159. URL `http://iopscience.iop.org/0031-9155/57/10/N159`.

[5] J-F Aubry, M Pernot, F Marquet, M Tanter, and M Fink. Transcostal high-intensity-focused ultrasound: ex vivo adaptive focusing feasibility study. *Physics in Medicine and Biology*, 53(11):2937–2951, June 2008. ISSN 0031-9155, 1361-6560. doi: 10.1088/0031-9155/53/11/012. URL `http://iopscience.iop.org/0031-9155/53/11/012`.

[6] Muyinatu A Lediju Bell, Brett C Byram, Emma J Harris, Philip M Evans, and Jeffrey C Bamber. In vivo liver tracking with a high volume rate 4D ultrasound scanner and a 2D matrix array probe. *Physics*

in Medicine and Biology, 57(5):1359–1374, March 2012. ISSN 0031-9155, 1361-6560. doi: 10.1088/0031-9155/57/5/1359. URL `http://iopscience.iop.org/0031-9155/57/5/1359`.

[7] Matt A. Bernstein, Kevin Franklin King, and Xiaohong Joe Zhou. *Handbook of MRI Pulse Sequences*. Elsevier, September 2004. ISBN 9780120928613.

[8] Norman Biggs. *Algebraic Graph Theory*. Cambridge University Press, 1993. ISBN 9780521458979.

[9] J M Blackall, S Ahmad, M E Miquel, J R McClelland, D B Landau, and D J Hawkes. MRI-based measurements of respiratory motion variability and assessment of imaging strategies for radiotherapy planning. *Physics in Medicine and Biology*, 51(17):4147–4169, September 2006. ISSN 0031-9155. doi: 10.1088/0031-9155/51/17/003. URL `http://iopscience.iop.org/0031-9155/51/17/003`.

[10] Svetlana Bobkova, Leonid Gavrilov, Vera Khokhlova, Adam Shaw, and Jeffrey Hand. Focusing of High-Intensity ultrasound through the rib cage using a therapeutic random phased array. *Ultrasound in Medicine & Biology*, 36(6):888–906, June 2010. ISSN 03015629. doi: 10.1016/j.ultrasmedbio.2010.03.007. URL `http://www.umbjournal.org/article/S0301-5629(10)00117-1/abstract`.

[11] Richard J. Carella, Joseph Ransohoff, and Joseph Newall. Role of radiation therapy in the management of meningioma. *Neurosurgery*, 10(3), 1982. ISSN 0148-396X. URL `http://journals.lww.com/neurosurgery/Fulltext/1982/03000/Role_of_Radiation_Therapy_in_the_Management_of.6.aspx`.

[12] C G Chaussy and G J Fuchs. Current state and future developments of noninvasive treatment of human urinary stones with extracorporeal shock wave lithotripsy. *The Journal of urology*, 141(3 Pt 2):782–789, March 1989. ISSN 0022-5347. URL `http://www.ncbi.nlm.nih.gov/pubmed/2645437`. PMID: 2645437.

[13] Lisha Chen and Andreas Buja. Local multidimensional scaling for nonlinear dimension reduction, graph drawing, and proximity analysis. *Journal of the American Statistical Association*, 104(485):209–219, March 2009. ISSN 0162-1459. doi: 10.1198/jasa.2009.0111. URL `http://pubs.amstat.org/doi/abs/10.1198/jasa.2009.0111`.

[14] Richard S. C. Cobbold. *Foundations of Biomedical Ultrasound*. Oxford University Press, 2007. ISBN 9780195168310.

[15] Laura Curiel, Rajiv Chopra, and Kullervo Hynynen. Progress in multi-modality imaging: Truly simultaneous ultrasound and magnetic resonance imaging. *IEEE Transactions on Medical Imaging*, 26(12): 1740–1746, December 2007. ISSN 0278-0062. PMID: 18092742 PM-CID: 2862902.

[16] John Daugman. The importance of being random: statistical principles of iris recognition. *Pattern Recognition*, 36(2):279–291, February 2003. ISSN 0031-3203. doi: 10.1016/S0031-3203(02) 00030-4. URL http://www.sciencedirect.com/science/article/pii/S0031320302000304.

[17] S C Davies, A L Hill, R B Holmes, M Halliwell, and P C Jackson. Ultrasound quantitation of respiratory organ motion in the upper abdomen. *Br J Radiol*, 67(803):1096–1102, November 1994. doi: <p>10.1259/0007-1285-67-803-1096</p>. URL http://bjr.birjournals.org/cgi/content/abstract/67/803/1096.

[18] J. de Leeuw. Applications of convex analysis to multidimensional scaling. *Recent Developments in Statistics. Amsterdam: North Holland*, page 133–145, 1977. URL http://escholarship.org/uc/item/7wg0k7xq.

[19] Philippe Lourenço de Oliveira, Baudouin Denis de Senneville, Iulius Dragonu, and Chrit T. W. Moonen. Rapid motion correction in MR-guided high-intensity focused ultrasound heating using real-time ultrasound echo information. *NMR in Biomedicine*, 23(9):1103–1108, November 2010. ISSN 09523480. doi: 10.1002/nbm.1526. URL http://onlinelibrary.wiley.com/doi/10.1002/nbm.1526/abstract.

[20] Baudouin Denis de Senneville, Sébastien Roujol, Chrit Moonen, and Mario Ries. Motion correction in MR thermometry of abdominal organs: A comparison of the referenceless vs. the multibaseline approach. *Magnetic Resonance in Medicine*, 64(5):1373–1381, November 2010. ISSN 1522-2594. doi: 10.1002/mrm.22514. URL http://onlinelibrary.wiley.com/doi/10.1002/mrm.22514/abstract;jsessionid=6683A1053FFED419095D9E92186C7447.d03t02.

[21] B. Denis de Senneville, B. Quesson, and C. T. W. Moonen. Magnetic resonance temperature imaging. *International Journal of Hyperthermia*,

21(6):515–531, January 2005. ISSN 0265-6736, 1464-5157. doi: 10.1080/ 02656730500133785. URL `http://informahealthcare.com/doi/abs/ 10.1080/02656730500133785`.

[22] H. Durrant-Whyte, D. Rye, and E. Nebot. Localisation of automatic guided vehicles. In *Robotics Research: The 7th International Symposium (ISRR'95)*, page 613–625, 1996.

[23] R L Ehman and J P Felmlee. Adaptive technique for high-definition MR imaging of moving structures. *Radiology*, 173(1):255 –263, October 1989. URL `http://radiology.rsna.org/content/173/1/255. abstract`.

[24] David A Feinberg, Daniel Giese, D. Andre Bongers, Sudhir Ramanna, Maxim Zaitsev, Michael Markl, and Matthias Günther. Hybrid ultrasound MRI for improved cardiac imaging and real-time respiration control. *Magnetic Resonance in Medicine*, 63(2):290–296, February 2010. ISSN 1522-2594. doi: 10.1002/mrm.22250. URL `http: //onlinelibrary.wiley.com/doi/10.1002/mrm.22250/abstract; jsessionid=8596DD6DDFE8CDD2578426FB3A000A71.d04t01? userIsAuthenticated=false&deniedAccessCustomisedMessage=`.

[25] FDA Center for Devices and Radiological Health. Consumer updates - ExAblate® 2000 system - p040003. http://www.fda.gov/ForConsumers/ConsumerUpdates/ucm080704.htm, October 2004. URL `http://www.fda.gov/ForConsumers/ ConsumerUpdates/ucm080704.htm`. ExAblate® 2000 System - P040003.

[26] US Centers for Disease Control. NHDS - national hospital discharge survey homepage. http://www.cdc.gov/nchs/nhds.htm, 2009. URL `http://www.cdc.gov/nchs/nhds.htm`.

[27] M. Frigo and S.G. Johnson. The design and implementation of FFTW3. *Proceedings of the IEEE*, 93(2):216–231, 2005. ISSN 0018-9219. doi: 10.1109/JPROC.2004.840301.

[28] Refaat E. Gabr, Pelin Aksit, Paul A. Bottomley, Abou-Bakr M. Youssef, and Yasser M. Kadah. Deconvolution-Interpolation gridding (DING): accurate reconstruction for arbitrary k-Space trajectories. *Magnetic Resonance in Medicine*, 56(6):1182–1191, December 2006. ISSN 07403194. doi: 10.1002/mrm.21095. URL `http: //onlinelibrary.wiley.com/doi/10.1002/mrm.21095/full`.

115

[29] David Gianfelice, Chander Gupta, Walter Kucharczyk, Patrice Bret, Deborah Havill, and Mark Clemons. Palliative treatment of painful bone metastases with MR imaging–guided focused ultrasound1. *Radiology*, 249(1):355–363, October 2008. ISSN 0033-8419, 1527-1315. doi: 10.1148/radiol.2491071523. URL `http://radiology.rsna.org/content/249/1/355`.

[30] Parag R. Gogate and Abhijeet M. Kabadi. A review of applications of cavitation in biochemical engineering/biotechnology. *Biochemical Engineering Journal*, 44(1):60–72, April 2009. ISSN 1369-703X. doi: 10.1016/j.bej.2008.10.006. URL `http://www.sciencedirect.com/science/article/pii/S1369703X08003392`.

[31] William A Grissom, Michael Lustig, Andrew B Holbrook, Viola Rieke, John M Pauly, and Kim Butts-Pauly. Reweighted 1 referenceless PRF shift thermometry. *Magnetic Resonance in Medicine*, 64(4):1068–1077, October 2010. ISSN 1522-2594. doi: 10.1002/mrm.22502. URL `http://onlinelibrary.wiley.com/doi/10.1002/mrm.22502/abstract`.

[32] William A. Grissom, Viola Rieke, Andrew B. Holbrook, Yoav Medan, Michael Lustig, Juan Santos, Michael V. McConnell, and Kim Butts Pauly. Hybrid referenceless and multibaseline subtraction MR thermometry for monitoring thermal therapies in moving organs. *Medical Physics*, 37(9):5014–5026, 2010. doi: 10.1118/1.3475943. URL `http://link.aip.org/link/?MPH/37/5014/1`.

[33] Hákon Gudbjartsson and Samuel Patz. The rician distribution of noisy MRI data. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, 34(6):910–914, December 1995. ISSN 0740-3194. PMID: 8598820 PMCID: PMC2254141.

[34] Matthias Günther and David A. Feinberg. Ultrasound-guided MRI: preliminary results using a motion phantom. *Magnetic Resonance in Medicine*, 52(1):27–32, 2004. doi: 10.1002/mrm.20140. URL `http://dx.doi.org/10.1002/mrm.20140`.

[35] Y. Hertzberg and G. Navon. Bypassing absorbing objects in focused ultrasound using computer generated holographic technique. *Medical Physics*, 38(12):6407–6415, 2011. doi: 10.1118/1.3651464. URL `http://link.aip.org/link/?MPH/38/6407/1`.

[36] J. C. Hindman. Proton resonance shift of water in the gas and liquid states. *The Journal of Chemical Physics*, 44(12):4582–4592, June 1966.

ISSN 00219606. doi: doi:10.1063/1.1726676. URL `http://jcp.aip.org/resource/1/jcpsa6/v44/i12/p4582_s1?isAuthorized=no`.

[37] P. Huber, B. Stepanow, J. Debus, K. Jochle, M. Mory, J. Jenne, A. Werner, G. van Kaick, and W.J. Lorenz. Temperature monitoring of focused ultrasound therapy by MRI. In *Ultrasonics Symposium, 1994. Proceedings., 1994 IEEE*, volume 3, pages 1825 –1828 vol.3, November 1994. doi: 10.1109/ULTSYM.1994.401945.

[38] Kullervo Hynynen, Oliver Pomeroy, Darrell N Smith, Peter E Huber, Nathan J McDannold, Joachim Kettenbach, Janet Baum, Samuel Singer, and Ferenc A Jolesz. MR Imaging-Guided focused ultrasound surgery of fibroadenomas in the breast: A feasibility study1. *Radiology*, 219(1):176–185, April 2001. ISSN 0033-8419, 1527-1315. URL `http://radiology.rsna.org/content/219/1/176`.

[39] Kullervo Hynynen, Nathan McDannold, Nickolai A. Sheikov, Ferenc A. Jolesz, and Natalia Vykhodtseva. Local and reversible blood–brain barrier disruption by noninvasive focused ultrasound at frequencies suitable for trans-skull sonications. *NeuroImage*, 24(1):12–20, January 2005. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2004.06.046. URL `http://www.sciencedirect.com/science/article/pii/S1053811904003362`.

[40] R 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. *British Journal of Cancer*, 93(8):890–895, September 2005. ISSN 0007-0920. doi: 10.1038/sj.bjc.6602803. URL `http://www.nature.com/bjc/journal/v93/n8/full/6602803a.html`.

[41] Yasutoshi Ishihara, Arturo Calderon, Hidehiro Watanabe, Kazuya Okamoto, Yoshinori Suzuki, Kagayaki Kuroda, and Yutaka Suzuki. A precise and fast temperature mapping using water proton chemical shift. *Magnetic Resonance in Medicine*, 34(6): 814–823, December 1995. ISSN 1522-2594. doi: 10.1002/mrm.1910340606. URL `http://onlinelibrary.wiley.com/doi/10.1002/mrm.1910340606/abstract`.

[42] F. A. Jolesz, A. R. Bleier, P. Jakab, P. W. Ruenzel, K. Huttl, and G. J. Jako. MR imaging of Laser-Tissue interactions. *Radiology*, 168

(1):249–253, July 1988. ISSN 0033-8419, 1527-1315. URL `http://radiology.rsna.org/content/168/1/249`.

[43] Ferenc A. Jolesz. MRI-Guided focused ultrasound surgery. *Annual Review of Medicine*, 60(1):417–430, 2009. doi: 10.1146/annurev.med.60.041707.170303. URL `http://www.annualreviews.org/doi/abs/10.1146/annurev.med.60.041707.170303`.

[44] Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: open source scientific tools for python, 2001. URL `http://www.scipy.org/`.

[45] R E Kalman. A new approach to linear filtering and prediction problems. *Journal Of Basic Engineering*, 82(Series D):35–45, 1960. doi: 10.1109/ICASSP.1982.1171734.

[46] Kim Butts Pauly, Andrew Holbrook, Juan Santos, and Pejman Ghanouni. Focused ultrasound of the liver during free breathing. In *Program & Abstract Book*, page 95, Washington, D.C., USA, October 2010. Focused Ultrasound Surgery Foundation. URL `http://www.fusfoundation.org/media/MRgFUS%20Symposium%202010%20Final%20Program.pdf`.

[47] Donald Ervin Knuth. *The art of computer programming: Sorting and searching*, volume 3. Addison-Wesley, 1998. ISBN 9780201896855.

[48] Andrew V. Knyazev. Toward the optimal preconditioned eigensolver: Locally optimal block preconditioned conjugate gradient method. *SIAM Journal on Scientific Computing*, 23(2):517–541, 2001. doi: 10.1137/S1064827500366124. URL `http://link.aip.org/link/?SCE/23/517/1`.

[49] Douglas Kondziolka, David Mathieu, L. Dade Lunsford, Juan J. Martin, Ricky Madhok, Ajay Niranjan, and John C. Flickinger. Radiosurgery as definitive management of intracranial meningiomas. *Neurosurgery*, 62(1), 2008. ISSN 0148-396X. URL `http://journals.lww.com/neurosurgery/Fulltext/2008/01000/Radiosurgery_As_Definitive_Management_of.4.aspx`.

[50] D. Le Bihan, J. Delannoy, and R L Levin. Temperature mapping with MR imaging of molecular diffusion: Application to hyperthermia. *Radiology*, 171(3):853–857, June 1989. ISSN 0033-8419, 1527-1315. URL `http://radiology.rsna.org/content/171/3/853`.

[51] P. P. Lele. A simple method for production of trackless focal lesions with focused ultrasound: physical factors. *The Journal of Physiology*, 160 (3):494–512, March 1962. ISSN 0022-3751. PMID: 14463953 PMCID: PMC1359560.

[52] F. Marquet, M. Pernot, J. -F Aubry, M. Tanter, G. Montaldo, and M. Fink. In-vivo non-invasive motion tracking and correction in high intensity focused ultrasound therapy. In *28th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 2006. EMBS '06*, pages 688–691. IEEE, September 2006. ISBN 1-4244-0032-5. doi: 10.1109/IEMBS.2006.259963.

[53] Ernst Martin, Daniel Jeanmonod, Anne Morel, Eyal Zadicario, and Beat Werner. High-intensity focused ultrasound for noninvasive functional neurosurgery. *Annals of Neurology*, 66(6):858–861, December 2009. ISSN 1531-8249. doi: 10.1002/ana.21801. URL `http://onlinelibrary.wiley.com/doi/10.1002/ana.21801/abstract;jsessionid=E602F286D84F0E5EF733523E9D7696BE.d04t03?userIsAuthenticated=false&deniedAccessCustomisedMessage=`.

[54] Michael V McConnell, Vaibhav C Khasgiwala, Bernard J Savord, Ming Hui Chen, Michael L Chuang, Warren J Manning, and Robert R Edelman. Prospective adaptive navigator correction for breath-hold MR coronary angiography. *Magnetic Resonance in Medicine*, 37 (1):148–152, January 1997. ISSN 1522-2594. doi: 10.1002/mrm.1910370121. URL `http://onlinelibrary.wiley.com/doi/10.1002/mrm.1910370121/abstract`.

[55] N McDannold. Quantitative MRI-based temperature mapping based on the proton resonant frequency shift: review of validation studies. *International Journal of Hyperthermia: The Official Journal of European Society for Hyperthermic Oncology, North American Hyperthermia Group*, 21(6):533–546, September 2005. ISSN 0265-6736. doi: 10.1080/02656730500096073. URL `http://www.ncbi.nlm.nih.gov/pubmed/16147438`. PMID: 16147438.

[56] N McDannold, N Vykhodtseva, and K Hynynen. Targeted disruption of the blood–brain barrier with focused ultrasound: association with cavitation activity. *Physics in Medicine and Biology*, 51(4):793–807, February 2006. ISSN 0031-9155, 1361-6560. doi: 10.1088/0031-9155/51/4/003. URL `http://iopscience.iop.org/0031-9155/51/4/003`.

[57] Nathan McDannold and Stephan E. Maier. Magnetic resonance acoustic radiation force imaging. *Medical Physics*, 35(8):3748–3758, August

2008. ISSN 0094-2405. doi: 10.1118/1.2956712. PMID: 18777934 PM-CID: PMC2673647.

[58] Nathan McDannold, Greg T. Clement, Peter Black, Ferenc Jolesz, and Kullervo Hynynen. Transcranial magnetic resonance imaging– guided focused ultrasound surgery of brain tumors. *Neurosurgery*, 66(2):323–332, February 2010. ISSN 0148-396X. doi: 10.1227/01.NEU.0000360379.95800.2F. URL `http://journals.lww.com/neurosurgery/Abstract/2010/02000/Transcranial_Magnetic_Resonance_Imaging__Guided.20.aspx`.

[59] Nathan J McDannold, Natalia I Vykhodtseva, and Kullervo Hynynen. Microbubble contrast agent with focused ultrasound to create brain lesions at low power levels: MR imaging and histologic study in rabbits. *Radiology*, 241(1):95–106, October 2006. ISSN 0033-8419, 1527-1315. doi: 10.1148/radiol.2411051170. URL `http://radiology.rsna.org/content/241/1/95`.

[60] Kazushi Numata, Hiroyuki Fukuda, Masao Ohto, Ryu Itou, Akito Nozaki, Masaaki Kondou, Manabu Morimoto, Eii Karasawa, and Katsuaki Tanaka. Evaluation of the therapeutic efficacy of high-intensity focused ultrasound ablation of hepatocellular carcinoma by three-dimensional sonography with a perflubutane-based contrast agent. *European Journal of Radiology*, 75(2):e67–75, August 2010. ISSN 1872-7727. doi: 10.1016/j.ejrad.2009.11.022. URL `http://www.ncbi.nlm.nih.gov/pubmed/20015603`. PMID: 20015603.

[61] Dennis L. Parker, Vernon Smith, Philip Sheldon, Lawrence E. Crooks, and Lauren Fussell. Temperature distribution measurements in two-dimensional NMR imaging. *Medical Physics*, 10(3):321–325, 1983. doi: 10.1118/1.595307. URL `http://link.aip.org/link/?MPH/10/321/1`.

[62] Ari Partanen, Charles Mougenot, and Teuvo Vaara. Feasibility of Agar-Silica phantoms in quality assurance of MRgHIFU. *AIP Conference Proceedings*, 1113(1):296–300, April 2009. doi: 10.1063/1.3131434. URL `http://link.aip.org/link/?APC/1113/296/1`.

[63] Maggy Patrick. Tremors halted with painless, Non-Invasive brain surgery. *ABC World News with Diane Sawyer*, January 2012. URL `http://abcnews.go.com/Health/patients-tremors-stopped-invasive-pain-free-brain-surgery/story?id=15406244`.

[64] Mathieu Pernot, Mickaël Tanter, and Mathias Fink. 3-D real-time motion correction in high-intensity focused ultrasound therapy. *Ultrasound in Medicine & Biology*, 30(9):1239–1249, September 2004. ISSN 0301-5629. doi: 10.1016/j.ultrasmedbio.2004.07.021. URL `http://www.sciencedirect.com/science/article/pii/S0301562904001942`.

[65] Jason A. Polzin, David G. Kruger, David H. Gurr, Jean H. Brittain, and Stephen J. Riederer. Correction for gradient nonlinearity in continuously moving table MR imaging. *Magnetic Resonance in Medicine*, 52 (1):181–187, July 2004. ISSN 1522-2594. doi: 10.1002/mrm.20117. URL `http://onlinelibrary.wiley.com/doi/10.1002/mrm.20117/full`.

[66] Bruno Quesson, Jacco A. de Zwart, and Chrit T. W. Moonen. Magnetic resonance temperature imaging for guidance of thermotherapy. *Journal of Magnetic Resonance Imaging*, 12(4):525–533, October 2000. ISSN 1522-2586. doi: 10.1002/1522-2586(200010)12:4<525::AID-JMRI3> 3.0.CO;2-V. URL `http://onlinelibrary.wiley.com/doi/10.1002/1522-2586(200010)12:4<525::AID-JMRI3>3.0.CO;2-V/abstract`.

[67] Natalya Y. Rapoport, Anne M. Kennedy, Jill E. Shea, Courtney L. Scaife, and Kweon-Ho Nam. Controlled and targeted tumor chemotherapy by ultrasound-activated nanoemulsions/microbubbles. *Journal of Controlled Release*, 138(3):268–276, September 2009. ISSN 0168-3659. doi: 10.1016/j.jconrel.2009.05.026. URL `http://www.sciencedirect.com/science/article/pii/S0168365909003599`.

[68] Viola Rieke, Karl K Vigen, Graham Sommer, Bruce L Daniel, John M Pauly, and Kim Butts. Referenceless PRF shift thermometry. *Magnetic Resonance in Medicine*, 51(6):1223–1231, June 2004. ISSN 1522-2594. doi: 10.1002/mrm.20090. URL `http://onlinelibrary.wiley.com/doi/10.1002/mrm.20090/abstract`.

[69] Mario Ries, Baudouin Denis de Senneville, Sébastien Roujol, Yasmina Berber, Bruno Quesson, and Chrit Moonen. Real-time 3D target tracking in MRI guided focused ultrasound ablations in moving tissues. *Magnetic Resonance in Medicine*, pages 1704–1712, September 2010. ISSN 07403194. doi: 10.1002/mrm.22548. URL `http://onlinelibrary.wiley.com/doi/10.1002/mrm.22548/abstract`.

[70] Torsten Rohlfing, Calvin R. Maurer, Walter G. O'Dell, and Jianhui Zhong. Modeling liver motion and deformation during the respiratory cycle using intensity-based nonrigid registration of gated MR images. *Medical Physics*, 31:427, 2004. ISSN 00942405. doi: 10.1118/

1.1644513. URL `http://online.medphys.org/resource/1/mphya6/v31/i3/p427_s1?isAuthorized=no`.

[71] James C. Ross, Rekha Tranquebar, and Dattesh Shanbhag. Real-Time liver motion compensation for MRgFUS. In Dimitris Metaxas, Leon Axel, Gabor Fichtinger, and Gábor Székely, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2008*, volume 5242, pages 806–813. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008. ISBN 978-3-540-85989-5. URL `http://www.springerlink.com/content/v14lt36810p82707/`.

[72] J.M. Santos, G.A. Wright, and J.M. Pauly. Flexible real-time magnetic resonance imaging framework. In *Engineering in Medicine and Biology Society, 2004. IEMBS '04. 26th Annual International Conference of the IEEE*, volume 1, pages 1048–1051, 2004. doi: 10.1109/IEMBS.2004.1403343.

[73] Stephen A. Sapareto and William C. Dewey. Thermal dose determination in cancer therapy. *International Journal of Radiation Oncology*Biology*Physics*, 10(6):787–800, April 1984. ISSN 0360-3016. doi: 10.1016/0360-3016(84)90379-1. URL `http://www.sciencedirect.com/science/article/pii/0360301684903791`.

[74] W. G Schneider, H. J Bernstein, and J. A Pople. Proton magnetic resonance chemical shift of free (Gaseous) and associated (Liquid) hydride molecules. *The Journal of Chemical Physics*, 28(4):601–607, April 1958. ISSN 00219606. doi: doi:10.1063/1.1744199. URL `http://jcp.aip.org/resource/1/jcpsa6/v28/i4/p601_s1`.

[75] Andrew D. Scott, Jennifer Keegan, and David N. Firmin. Motion in cardiovascular MR imaging. *Radiology*, 250(2):331 –351, February 2009. doi: 10.1148/radiol.2502071998. URL `http://radiology.rsna.org/content/250/2/331.abstract`.

[76] Baudouin Denis de Senneville, Charles Mougenot, and Chrit T.W. Moonen. Real-time adaptive methods for treatment of mobile organs by MRI-controlled high-intensity focused ultrasound. *Magnetic Resonance in Medicine*, 57(2):319–330, 2007. doi: 10.1002/mrm.21124. URL `http://dx.doi.org/10.1002/mrm.21124`.

[77] Simon A Shakespeare, Rachel J Moore, John A Crowe, Penny A Gowland, and Barrie R Hayes-Gill. A method for foetal heart rate monitoring during magnetic resonance imaging using doppler ultrasound. *Physiological Measurement*, 20:363–368, November 1999. ISSN

0967-3334, 1361-6579. doi: 10.1088/0967-3334/20/4/303. URL `http://iopscience.iop.org/0967-3334/20/4/303`.

[78] Claude E. Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10– 21, January 1949. ISSN 0096-8390. doi: 10.1109/JRPROC.1949.232969.

[79] S D Sokka, R King, and K Hynynen. MRI-guided gas bubble enhanced ultrasound heating in in vivo rabbit thigh. *Physics in Medicine and Biology*, 48(2):223–241, January 2003. ISSN 0031-9155. doi: 10.1088/0031-9155/48/2/306. URL `http://iopscience.iop.org/0031-9155/48/2/306`.

[80] Elizabeth A. Stewart, Jaron Rabinovici, Clare M.C. Tempany, Yael Inbar, Leslie Regan, Bobbie Gastout, Gina Hesley, Hyun S. Kim, Suzanne Hengst, and Wladyslaw M. Gedroye. Clinical outcomes of focused ultrasound surgery for the treatment of uterine fibroids. *Fertility and Sterility*, 85(1):22–29, January 2006. ISSN 0015-0282. doi: 10.1016/j.fertnstert.2005.04.072. URL `http://www.sciencedirect.com/science/article/pii/S0015028205036551`.

[81] A. M Tang, D. F Kacher, E. Y Lam, K. K Wong, F. A Jolesz, and E. S Yang. Simultaneous ultrasound and MRI system for breast biopsy: Compatibility assessment and demonstration in a dual modality phantom. *IEEE Transactions on Medical Imaging*, 27(2):247–254, February 2008. ISSN 0278-0062. doi: 10.1109/TMI.2007.911000.

[82] M Tanter, M Pernot, J F Aubry, G Montaldo, F Marquet, and M Fink. Compensating for bone interfaces and respiratory motion in high-intensity focused ultrasound. *International journal of hyperthermia: the official journal of European Society for Hyperthermic Oncology, North American Hyperthermia Group*, 23(2):141–151, March 2007. ISSN 0265-6736. URL `http://www.ncbi.nlm.nih.gov/pubmed/17578338`. PMID: 17578338.

[83] Junichi Tokuda, Gregory S. Fischer, Xenophon Papademetris, Ziv Yaniv, Luis Ibanez, Patrick Cheng, Haiying Liu, Jack Blevins, Jumpei Arata, Alexandra J. Golby, Tina Kapur, Steve Pieper, Everette C. Burdette, Gabor Fichtinger, Clare M. Tempany, and Nobuhiko Hata. OpenIGTLink: an open network protocol for image-guided therapy environment. *The international journal of medical robotics + computer assisted surgery : MRCAS*, 5(4):423–434, December 2009. ISSN 1478-5951. doi: 10.1002/rcs.274. PMID: 19621334 PMCID: 2811069.

[84] Warren S. Torgerson. Multidimensional scaling: I. theory and method. *Psychometrika*, 17(4):401–419, December 1952. ISSN 0033-3123. doi: 10.1007/BF02288916. URL http://www.springerlink.com/content/86u3178hvh181k57/.

[85] G. Vallancien, M. Harouni, B. Veillon, A. Mombet, D. Prapotnich, J.M. Brisset, and J. Bougaran. Focused extracorporeal pyrotherapy: Feasibility study in man. *Journal of Endourology*, 6(2):173–181, April 1992. ISSN 0892-7790, 1557-900X. doi: 10.1089/end.1992.6.173. URL http://online.liebertpub.com/doi/abs/10.1089/end.1992.6.173.

[86] Karl K. Vigen, Bruce L. Daniel, John M. Pauly, and Kim Butts. Triggered, navigated, multi-baseline method for proton resonance frequency temperature mapping with respiratory motion. *Magnetic Resonance in Medicine*, 50(5):1003–1010, 2003. doi: 10.1002/mrm.10608. URL http://dx.doi.org/10.1002/mrm.10608.

[87] F. Viola and W. F Walker. A comparison of the performance of time-delay estimators in medical ultrasound. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 50(4):392–401, April 2003. ISSN 0885-3010. doi: 10.1109/TUFFC.2003.1197962.

[88] Natalia Vykhodtseva, Nathan Mcdannold, Heather Martin, Roderick T Bronson, and Kullervo Hynynen. Apoptosis in ultrasound-produced threshold lesions in the rabbit brain. *Ultrasound in Medicine and Biology*, 27(1):111–117, January 2001. ISSN 0301-5629. URL http://www.umbjournal.org/article/S0301-5629(00)00275-1/abstract.

[89] Patrick D. Wall, William J. Fry, Ralph Stephens, Don Tucker, and Jerome Y. Lettvin. Changes produced in the central nervous system by ultrasound. *Science*, 114(2974):686–687, December 1951. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.114.2974.686. URL http://www.sciencemag.org/content/114/2974/686.

[90] Y Wang, P J Rossman, R C Grimm, S J Riederer, and R L Ehman. Navigator-echo-based real-time respiratory gating and triggering for reduction of respiration effects in three-dimensional coronary MR angiography. *Radiology*, 198(1):55–60, January 1996. ISSN 0033-8419. URL http://www.ncbi.nlm.nih.gov/pubmed/8539406.

[91] J. C Waterton, J. P. R Jenkins, X. P Zhu, H. G Love, I. Isherwood, and D. J Rowlands. Magnetic resonance (MR) cine imaging of the human heart. *British Journal of Radiology*, 58(692):711–716, August

1985. ISSN 0007-1285, 1748-880X. doi: 10.1259/0007-1285-58-692-711. URL http://bjr.birjournals.org/content/58/692/711.

[92] Feng Wu, Zhi-Biao Wang, Wen-Zhi Chen, Hui Zhu, Jin Bai, Jian-Zhong Zou, Ke-Quan Li, Cheng-Bing Jin, Fang-Lin Xie, and Hai-Bing Su. Extracorporeal high intensity focused ultrasound ablation in the treatment of patients with large hepatocellular carcinoma. *Annals of Surgical Oncology*, 11(12):1061–1069, December 2004. ISSN 1068-9265. doi: 10.1245/ASO.2004.02.026. URL http://www.springerlink.com/content/c9ct12f3lc6lxa27/.

[93] E. C. Zeeman. Unknotting combinatorial balls. *The Annals of Mathematics*, 78(3):501–526, November 1963. ISSN 0003-486X. doi: 10.2307/1970538. URL http://www.jstor.org/stable/1970538. ArticleType: research-article / Full publication date: Nov., 1963 / Copyright © 1963 Annals of Mathematics.

Colophon