The International Journal of Robotics Research OnlineFirst, published on May 19, 2009 as doi:10.1177/0278364909104065

Shelten G. Yuen
Daniel T. Kettler
Paul M. Novotny
Richard D. Plowes
Robert D. Howe*

School of Engineering and Applied Sciences,
Harvard University,
Cambridge, MA 02138,
USA
howe@seas.harvard.edu

Robotic Motion Compensation for Beating Heart Intracardiac Surgery

Abstract

3D ultrasound imaging has enabled minimally invasive, beating heart intracardiac procedures. However, rapid heart motion poses a serious challenge to the surgeon that is compounded by significant time delays and noise in 3D ultrasound. This paper investigates the concept of using a one-degree-of-freedom motion compensation system to synchronize with tissue motions that may be approximated by 1D motion models. We characterize the motion of the mitral valve annulus and show that it is well approximated by a 1D model. The subsequent development of a motion compensation instrument (MCI) is described, as well as an extended Kalman filter (EKF) that compensates for system delays. The benefits and robustness of motion compensation are tested in user trials under a series of non-ideal tracking conditions. Results indicate that the MCI provides an approximately 50% increase in dexterity and 50% decrease in force when compared with a solid tool, but is sensitive to time delays. We demonstrate that the use of the EKF for delay compensation restores performance, even in situations of high heart rate variability. The resulting system is tested in an in vitro 3D ultrasound-guided servoing task, yielding accurate tracking (1.15 mm root mean square) in the presence of noisy, time-delayed 3D ultrasound measurements.

KEY WORDS—medical robotics, motion compensation, ultrasound.

1. Introduction

Real-time 3D ultrasound (3DUS) imaging has enabled new surgical procedures within the beating heart that are not possible with current endoscopic technology (Cannon et al. 2003; Suematsu et al. 2005). These procedures eliminate the need for cardiopulmonary bypass and its attendant side effects, such as increased stroke risk, decline in cognitive performance, delay of neural development in children, and damage to the major vessels (Zeitlhofer et al. 1993; Bellinger et al. 1999; Murkin et al. 1999). Beating heart procedures also make the real-time evaluation of structural modifications possible while the heart continues to function. The benefits of such procedures are clear; however, the rapid motion of the heart poses serious challenges to the surgeon. This is especially true of procedures involving intracardiac structures such as the mitral valve, which recoils rapidly with every heartbeat. In order to realize the potential of minimally invasive beating heart mitral valve procedures, the difficulty of manipulating tissues undergoing fast motion must be addressed.

One appealing approach to the beating heart repair of intracardiac structures is robotic assistance to the surgeon. In this setting, the surgeon uses a handheld robotic instrument that tracks the fast motion of the heart so that the operation is performed on an effectively stopped heart (i.e. as if the heart were not beating). This is similar to previously proposed motion compensation systems in which the surgeon could interact with the heart on a supported moving platform (Trejos et al. 1999) or teleoperate on a "virtually" stopped heart (Nakamura et al. 2001; Ortmaier et al. 2005; Ginhoux et al. 2006; Bebek and Cavusoglu 2007). This approach should be distinguished from the recent clinical application of commercial sur-

The International Journal of Robotics Research
Vol. 00, No. 00, Xxxxxxxx 2009, pp. 000–000
DOI: 10.1177/0278364909104065

* Also with the Harvard-MIT Division of Health Sciences and Technology, Cambridge, MA 02139, USA.

gical robots in performing minimally invasive mitral valve repair, with a stopped heart and cardiopulmonary bypass (Reichenspurner et al. 2005; Woo et al. 2006). Prior work on beating heart motion compensation has largely focused on the development of robust tracking controllers intended for multi-degree-of-freedom (DOF) robot manipulators and extracardiac procedures. However, implementing a full six-DOF robot for intracardiac applications has a number of challenges, including the development of a manipulator with sufficient mechanical bandwidth, creating a wrist that can operate in the restricted workspace within the beating heart, and ensuring safety for a complex manipulator system. These requirements are far beyond the capabilities of current commercial surgical robots.

Another major consideration for beating intracardiac procedures is guidance inside of the heart, which requires imaging tissue through blood in real-time. Currently, 3DUS is the only technology that has these capabilities: endoscopic systems are obscured by blood, and computed tomography and magnetic resonance imaging are not yet able to image at the speeds necessary for beating heart surgery. 3DUS also eases the difficulties in spatial perception experienced with traditional 2D ultrasound (Cannon et al. 2003). However, for all of its advantages, the acquisition and processing of 3DUS volumes introduces a substantial time delay, during which time the annulus can recoil over the majority of its trajectory (Novotny 2007). Servoing a robot with these measurements directly would result in large errors that could damage the valve and nearby structures.

In this paper we describe a robotic motion compensation system that is compatible with the space limitations inside the heart and also counteracts the delays and noise in 3DUS imaging. The system comprises an actuated one-DOF instrument commanded by a predictive filter (Figure 1). The surgeon can use this system in procedures where the motion of the cardiac tissue is largely in a single direction. As will be shown shortly, this includes the motion of the mitral valve annulus, and thus the system can be used for beating heart mitral annuloplasty: the most frequently performed element of mitral valve repair. The instrument consists of a linear actuator capable of following the major component of the annulus motion. The filter that is used to command the actuator, an extended Kalman filter (EKF), can accurately track and predict annulus motion in the presence of noise by exploiting quasiperioidicity in its motion. It is used to feed-forward the trajectory of a cardiac target to the instrument controller for accurate synchronization despite time delays.

In the following, we first characterize the motion of the mitral annulus to determine the tracking performance requirements of the system. Next, we present the design of the actuated one-DOF motion compensation instrument (MCI). We then describe the EKF and several other predictive filtering methods and compare them in simulation. Two subsequent user studies evaluate the efficacy of the robotic motion compensation system against traditional non-tracking tools in a simulated *in vitro* surgical task. Furthermore, they validate the robustness of the system in situations of high noise, time delay and heart rate variability. Finally, we measure the position tracking accuracy of the system in a series of 3DUS-guided motion synchronization experiments.

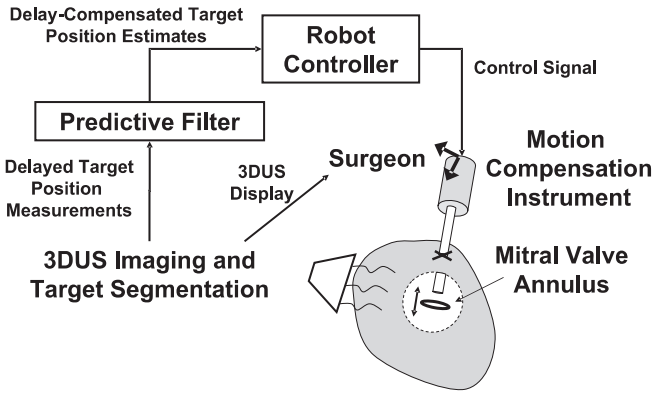

Fig. 1. Motion compensation system. Time-delayed position measurements of the target are provided to a filter which uses prediction to compensate for delay. These are fed-forward to the controller of the one-DOF MCI to synchronize with the target motion. The surgeon holds the instrument and navigates its end-effector to the target using 3DUS imaging.

2. Motion Analysis of the Mitral Valve Annulus

To guide the development of the motion compensation system, the motion of the mitral valve annulus was analyzed using ultrasound image data such as that available in surgery for real-time guidance. A transthoracic 3DUS image-volume sequence of the mitral annulus was acquired[1] at 24 Hz (SONOS 7500, Philips Medical Systems, Andover, MA, USA). This raw 3DUS data was manually segmented to extract mitral annulus trajectory information. For each 3D volume sample, a minimum of 50 data points were selected from the mitral valve annulus and used to locate the annulus centroid. Repeated for each time-stamped data frame, this process generated a record of the annulus centroid position over the cycle of the heartbeat. As this manual segmentation process was extremely labor intensive, only data from two patients were analyzed.

In order to specify hardware design constraints for the motion compensation instrument, position, velocity, and acceleration along the major axis of the valve's motion were extracted

1. 3DUS frame rates are between 24 and 28 Hz and are dependent on machine settings during data acquisition. Rates of 28 Hz are commonly obtained during *in vivo* procedures.

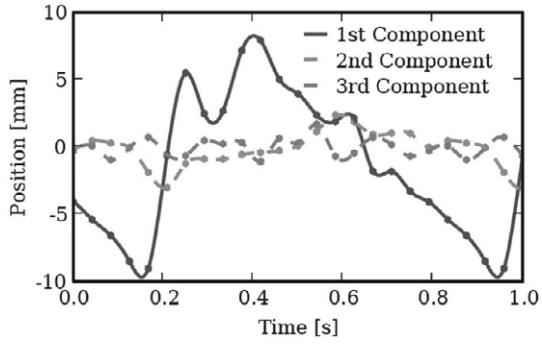

Fig. 2. Orthogonal views of 3DUS tracking data centroid positions plotted in principal component axes. The axes are scaled in millimeters. Note the dominance of motion in the first component.

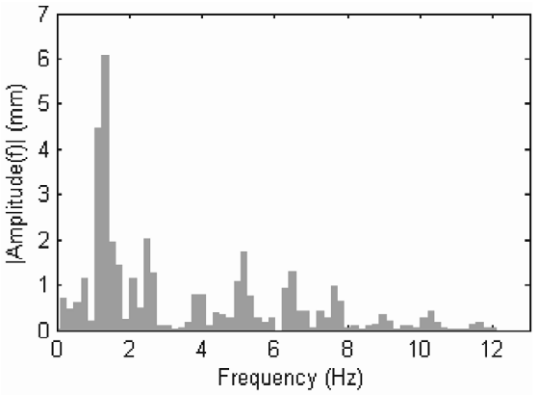

Fig. 3. 3DUS tracking data spectral decomposition. Motion amplitudes decrease quickly with increasing frequency.

from these data. Using singular value decomposition (SVD), a line was fit to each data set. The relative sizes of the singular value associated with each orthonormal basis vector generated by SVD (i.e. 21.75, 5.49, 2.97) suggest that the valve's motion is strongly constrained to a principal axis (Figure 2). Furthermore, working with sheep, Gorman et al. (1996) generated tracking data which indicates that rotational movements around this primary axis are also negligible. Subsequent analysis presented here will therefore focus on the component of motion along this primary axis.

Velocity and acceleration were estimated by considering the relative motion between 3DUS data samples. This method provided mitral valve velocity and acceleration maxima of approximately 210 mm s^{-1} and approximately 3.8 m s^{-2}, respectively, with a total travel of approximately 20 mm at a heart rate of 76 beats per minute (bpm). While only two subjects were analyzed in this fashion, work by Kamigaki and Goldschlager (1972) on the mitral valve leaflets reports similar velocity and amplitude results.

The major frequency components of motion were also reviewed using spectral decomposition (Figure 3). This analysis indicates major motion components at 1.3, 2.6, and 5.2 Hz with further components of decreasing amplitude at higher frequencies. This is consistent with the findings of Nakamura et al. (2001) which show dominant frequency components of 1.5 and 3.0 Hz in the motion of porcine epicardium. Ginhoux et al. (2006) found the same major frequency components in the motion of porcine epicardium. This paper also noted higher-frequency transients that it deemed significant and concluded that a 25 Hz sampling rate would be insufficient to track the motion of the epicardium with high precision (Ginhoux et al. 2006). Bebek and Cavusoglu (2007) found similar frequency components, but concluded that the motion could be adequately characterized using lower sampling frequencies, i.e. 26 Hz.

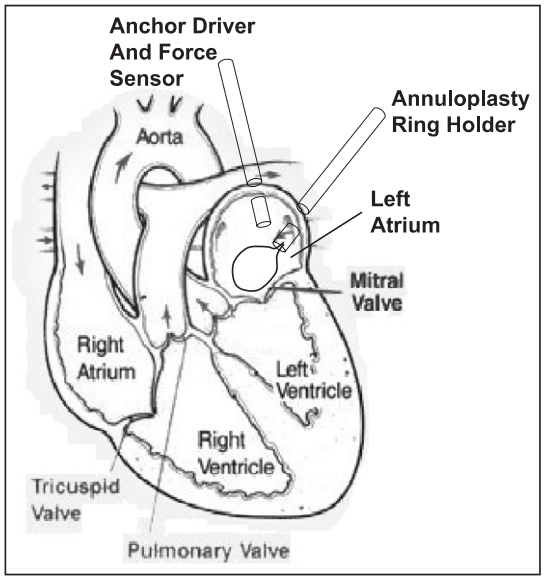

Fig. 4. Diagram of proposed surgical procedure. The anchor driver approaches along the valve's axis, which corresponds to the valve's major component of motion.

3. Motion Compensation Instrument Design

The MCI developed here is an actuated handheld or robot-mounted minimally invasive tool to aid surgeons in working on the moving mitral valve. Rather than attempting to correct for motion components in all three dimensions, the objective was to develop a device that could compensate for the major component of motion and allow surgeons and passive tissue compliance to counter the slow or relatively minor motions along the remaining two axes. The surgical procedure that we are developing for this purpose is a modification of a minimally invasive beating heart procedure for the repair of

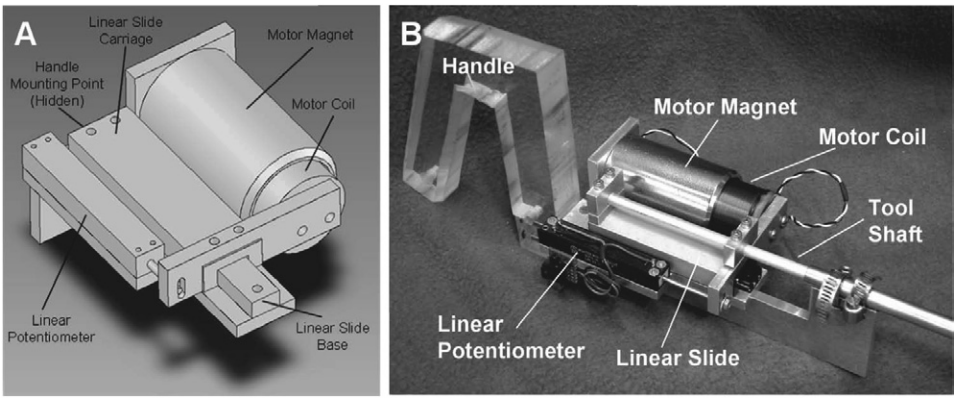

Fig. 5. MCI design and prototype. A, MCI actuation package. The MCI may be mounted by its base (at the upper left of the image) to a robot or handle. The surgical anchor driver is mounted on the linear slide carriage. B, MCI hardware prototype.

atrial–septal defects (Vasilyev et al. 2006). In the new procedure, the instrument will be inserted through the left atrial appendage and controlled by a pursestring suture. A custom annuloplastry ring will be inserted through an adjacent incision and positioned over the mitral annulus (Figure 4). The actuated instrument will be tipped with a modified 14 gauge needle (OD = 0.083 in) that will deploy anchors to "staple" the annuloplasty ring to the annulus (Wagner et al. 2006). The surgeon will then maneuver the ring and anchor driver to the appropriate locations over the annulus and repeatedly fire the anchor driver to attach the ring and reshape the annulus.

The selection of the mechanical mechanism to follow the linear motion component of the mitral valve annulus was guided by the clinical 3DUS trajectory analysis from Section 2. The high velocity and acceleration requirements lead to a linear motor-based design which benefits from low friction and low moving mass (Figure 5). This design format also produces a surgical tool similar in design and function to typical endoscopic tools, supported by a port and maneuvered by hand. Consequently, the MCI will be intuitive to operate for trained endoscopic surgeons.

The MCI uses a voice coil motor (NCC20-18-020-1X, H2W Technologies, Inc., Valencia, CA, USA) and a high linearity potentiometer (P3 America, San Diego, CA, USA) for position sensing. These components are mounted on a linear ball-bearing stage (BX4-3, Tusk Direct, Inc., Bethel, CT, USA). The MCI prototype has a 5.4 cm range of motion and is powered by a linear power amplifier (BTA-28V-6A, Precision MicroDynamics, Victoria, BC, Canada). PID servo control is implemented in a 1 kHz servo loop on a personal computer under Windows XP. The MCI is interfaced to this computer through a data acquisition card (PCI-60-40E, National Instruments, Corp., Austin, TX, USA).

The resulting system has the characteristics required to track the mitral valve. The MCI has demonstrated velocities and accelerations as high as approximately 290 mm s^{-1} and approximately 17.5 m s^{-2}, respectively. Controller gains were hand tuned to achieve good stiffness and response. To avoid potentially dangerous overshoot and instability, the system is overdamped. The tool has a static stiffness of 0.23 N mm^{-1} and a friction force less than 0.009 N. The system's frequency response is similarly adequate for the tracking task (Figure 6). The system has a −3 dB point of 20.6 Hz and roll-off rate of 40 dB per decade. The potentiometer on the MCI is capable of measuring position with a root mean square error (RMSE) of less than 0.01 mm. The system is capable of maintaining stationary at a commanded position with a RMSE of 0.009 mm.

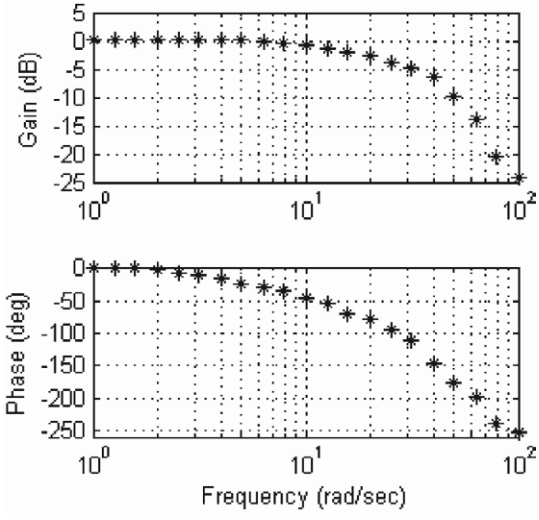

Fig. 6. Frequency response of the MCI. The system is overdamped and has a −3 dB point of 20.6 Hz.

The tracking abilities of the MCI were demonstrated by commanding the system to follow the motion of a mitral valve at 60 bpm (Figure 7). Mitral motion was determined from

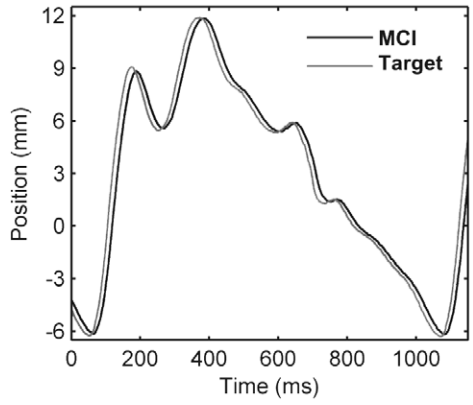

Fig. 7. The MCI tracking a simulated mitral valve trajectory. The MCI motion profile closely mimics that of the target with an approximately 14 ms delay.

the 3DUS data above (Figure 2). The MCI reliably replicated the motion profile of the valve with an effective delay of 14 ms.

4. Time Delay Compensation

The time delay that is intrinsic to 3DUS makes direct servoing of the MCI from this signal potentially dangerous. A previous characterization of the acquisition, transmission, and computation times for 3DUS estimated delays of approximately 70 ms, a sufficient amount of time for the mitral valve annulus to traverse the majority of its path (Novotny 2007). For the purposes of illustration, an example of the error for MCI tracking a mitral valve target with a 75 ms measurement delay is shown in Figure 8. The rapid recoil of the valve associated with ventricular relaxation (100–200 ms in Figure 8) results in a large (approximately 14 mm) tracking error that would cause the instrument to pierce and damage the tissue target.

4.1. Predictive Filters

To avoid this outcome, we exploit the near periodicity of the mitral valve trajectory to predict its path and hence compensate for time delay. However, such predictions must be made in the presence of measurement noise and a potentially variable heart rate. Here we describe and evaluate several predictive filtering methods that can be employed for delay compensation in this setting: an autoregressive (AR) filter, a fading memory AR filter, and an EKF with a quasiperiodic motion model. The AR filter has previously been applied by Nakamura et al. (2001) in a spectral analysis of heart motion for motion compensation in coronary artery bypass graft (CABG) procedures. In principle, this method is equivalent to the adaptive harmonic filter bank used by Ginhoux et al. (2006) for CABG and has its attendant assumption of a fixed heart rate. The fading memory AR filter overcomes this limitation despite using the same model by exponentially discounting the measurements supplied to the filter, thereby allowing it to adjust to more recent information. This approach has been used for motion synchronization in CABG by Franke et al. (2007). In contrast, the EKF permits variations to heart rate by directly accounting for it in a time-varying Fourier series model. A similar model was employed by Riviere et al. (1998) in the weighted Fourier linear combiner (WFLC) estimator for CABG; however, unlike the EKF, this method does not explicitly model noise. Ortmaier et al. (2005) has evaluated other non-linear prediction techniques for CABG such as artificial neural networks and an estimator based on Takens theorem.

4.1.1. AR Model with Least Squares Estimator

Fixed-rate mitral valve motion can be modeled as an n-order AR model

$$y[k] = \sum_{i=1}^{n} \alpha_i y[k-i], \tag{1}$$

where α_i, $i \in \{1, \ldots, n\}$, are the model coefficients and $y[k]$ is the target position at time sample k. Note that rather than explicitly assuming periodicity in the target motion, this model predicates that the kth position can be expressed as a linear combination of the previous n positions.

In order to predict the target position, the model coefficients and order must be estimated. The first can be achieved in real-time using the recursive covariance method estimator. Denoting $z[k] = y[k]+\nu[k]$ as the noise-corrupted position measurement at time sample k with $\nu[k] \sim \mathcal{N}\left(0, \sigma_R^2\right)$, this estimator is expressed compactly in matrix form as

$$\begin{aligned} \boldsymbol{Z}[k] &= [z[k-n-1], \ldots, z[k-1]], \\ \boldsymbol{R}[k] &= \boldsymbol{R}[k-1] + \boldsymbol{Z}[k]^{\mathrm{T}}\boldsymbol{Z}[k], \qquad (2) \\ \hat{\boldsymbol{\alpha}}[k] &= \hat{\boldsymbol{\alpha}}[k-1] + \boldsymbol{R}[k]^{-1}\boldsymbol{Z}[k]^{\mathrm{T}} \\ &\times \left(z[k] - \boldsymbol{Z}\hat{\boldsymbol{\alpha}}[k-1]\right), \qquad (3) \end{aligned}$$

with initial conditions $\boldsymbol{R}[0] = \boldsymbol{0}$ and $\hat{\boldsymbol{\alpha}}[0] = 0$. An appropriate AR model order was determined using the Akaike information criteria on the mitral valve annulus trajectory in Figure 2, yielding $n = 30$. Predicted target locations, $\hat{y}[k]$, can be obtained through evaluation of (1). In addition, the target trajectory can be interpolated from its inherent measurement rate (i.e. 28 Hz, a typical 3DUS frame rate) to the higher control rate of the robot using the Whittaker–Shannon interpolation formula.

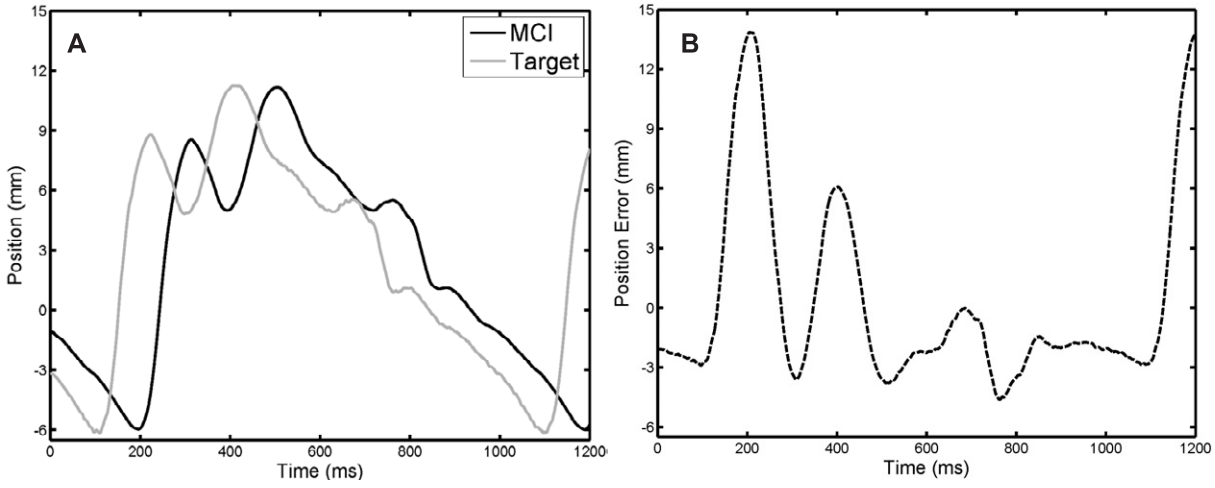

Fig. 8. MCI tracking of a mitral valve target with 75 ms measurement delay. A, MCI and target positions. B, Tracking error. Note that the additional 14 ms response time of the MCI yields an effective delay of 89 ms. Maximum and RMS tracking errors are 14.49 and 4.60 mm, respectively.

4.1.2. AR Model with Fading Memory Estimator

Imperfect periodicity can cause the AR model coefficients to change over time. In this situation, it can be useful to preferentially weight recent measurements over those in the past: otherwise the filter becomes progressively less responsive to new data and (3) does not update $\hat{\boldsymbol{\alpha}}$ because $\boldsymbol{R}[k]^{-1}\boldsymbol{Z}^{\mathrm{T}} \rightarrow \boldsymbol{0}$ as $k \rightarrow \infty$. Exponential weighting of previous measurements in the iterative least squares estimator is achieved through a simple modification of (2):

$$\boldsymbol{R}[k] = f\boldsymbol{R}[k-1] + \boldsymbol{Z}[k]^{\mathrm{T}}\boldsymbol{Z}[k],$$

where $0 < f \leq 1$ is the so-called *fading* factor. Choosing $f = 1$ recovers the estimator of Section 4.1.1 while choosing $f \rightarrow 0$ increases the speed by which previous measurements are discounted. To distinguish between the two estimators, we term the former the *AR filter* and the latter the *fading AR filter* for the remainder of this section. As can be imagined, reducing the contribution of previous measurements ($f < 1$) can be desirable if the trajectory evolves through time, although doing so incurs increased estimation error when the trajectory is not time varying.

4.1.3. Time-varying Fourier Series Model with EKF

The spectral analysis of mitral valve annulus motion from Section 2 suggests that its motion may be approximated by a limited number of harmonics. Consider a perfectly periodic motion model obtained by an m-order Fourier series with a constant offset

$$y(t) = c + \sum_{i=1}^{m} r_i \sin(i\omega t + \phi_i), \tag{4}$$

where $y(t)$ is the position in ultrasound coordinates, ω is the heart rate, c is the constant offset, and r_i and ϕ_i are the harmonic amplitudes and phases, respectively. Accurate modeling of quasiperiodicity requires a more flexible model in which the heart rate and signal morphology can evolve over time. Using the parametrization of Parker and Anderson (1990), the trajectory can be expressed as the following time-varying Fourier series

$$y(t) = c(t) + \sum_{i=1}^{m} r_i(t) \sin\theta_i(t), \tag{5}$$

where $\theta_i(t) = i\int_0^t \omega(\tau)\,d\tau + \phi_i(t)$ and all other parameters are the time-varying equivalents to those in (4).

Defining the state vector $\boldsymbol{x}(t) \triangleq [c(t), r_i(t), \omega(t), \theta_i(t)]^{\mathrm{T}}$, $i \in \{1, \ldots, m\}$ and assuming that $c(t)$, $r_i(t)$, $\omega(t)$, and $\phi_i(t)$ evolve through a random walk, the state space model for this system is

$$\begin{aligned} \boldsymbol{x}(t+\Delta t) &= \boldsymbol{F}(\Delta t)\boldsymbol{x}(t) + \boldsymbol{\mu}(t) \\ z(t) &= h(\boldsymbol{x}(t)) + \nu(t), \end{aligned}$$

where

$$\boldsymbol{F}(\Delta t) = \begin{bmatrix} \boldsymbol{I}_{m+1} & & & & \boldsymbol{0} \\ & 1 & & & \\ & \Delta t & 1 & & \\ \boldsymbol{0} & 2\Delta t & 0 & 1 & \\ & \vdots & & & \ddots \\ & m\Delta t & & & 1 \end{bmatrix},$$

$h(\boldsymbol{x}(t)) \triangleq y(t)$ from (5), $v(t) \sim \mathcal{N}\left(0, \sigma_R^2\right)$ is zero mean Gaussian measurement noise, and $\boldsymbol{\mu}(t) \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{Q})$ is the random step of the states assumed to be drawn from a zero mean multivariate normal distribution with covariance matrix $\boldsymbol{Q}$.

Prediction with this model requires estimation of the $2m+2$ parameters in $\boldsymbol{x}(t)$; a non-linear estimation problem owing to the measurement function, $h(\boldsymbol{x}(t))$. We employ the EKF, a non-linear filtering method that approximates the Kalman filter through linearization about the current state estimate $\hat{\boldsymbol{x}}(t \mid t)$. The EKF can be computed in real-time using the recursion

$$\begin{aligned}
\boldsymbol{P}(t+\Delta t \mid t) &= \boldsymbol{F}\boldsymbol{P}(t \mid t)\boldsymbol{F}^{\mathrm{T}} + \boldsymbol{Q} \\
S &= \sigma_R^2 + \boldsymbol{H}\boldsymbol{P}(t+\Delta t \mid t)\boldsymbol{H}^{\mathrm{T}} \\
\boldsymbol{K} &= \boldsymbol{P}(t+\Delta t \mid t)\boldsymbol{H}^{\mathrm{T}} S^{-1} \\
\hat{\boldsymbol{x}}(t+\Delta t \mid t+\Delta t) &= \boldsymbol{F}\hat{\boldsymbol{x}}(t \mid t) + \boldsymbol{K}(z(t+\Delta t) - h(\boldsymbol{F}\hat{\boldsymbol{x}}(t \mid t))) \\
\boldsymbol{P}(t+\Delta t \mid t+\Delta t) &= (\boldsymbol{I} - \boldsymbol{K}\boldsymbol{H})\boldsymbol{P}(t+\Delta t \mid t),
\end{aligned}$$

where

$$\begin{aligned}
\boldsymbol{H}^{\mathrm{T}} &\triangleq \left(\frac{\partial h}{\partial \boldsymbol{x}}\right)^{\mathrm{T}}\Bigg|_{\hat{\boldsymbol{x}}(t+\Delta t|t)=\boldsymbol{F}\hat{\boldsymbol{x}}(t|t)} \\
&= \begin{bmatrix}
1 \\
\sin\hat{\theta}_1(t+\Delta t \mid t) \\
\vdots \\
\sin\hat{\theta}_m(t+\Delta t \mid t) \\
0 \\
\hat{r}_1(t+\Delta t \mid t)\cos\hat{\theta}_1(t+\Delta t \mid t) \\
\vdots \\
\hat{r}_m(t+\Delta t \mid t)\cos\hat{\theta}_m(t+\Delta t \mid t)
\end{bmatrix},
\end{aligned}$$

and $\boldsymbol{P}(\cdot) \triangleq E\left[\hat{\boldsymbol{x}}(\cdot)\hat{\boldsymbol{x}}(\cdot)^{\mathrm{T}}\right]$ denotes the state estimate covariance, whose initialization is described later in this section. Note that the time dependencies for $\boldsymbol{F}$, $\boldsymbol{K}$, S, and $\boldsymbol{H}$ have been dropped for notational convenience. The EKF as presented here is a slight variant on that first introduced by Parker and Anderson (1990).

To initialize this filter, we first assume that $y(t)$ has constant fundamental frequency yielding N noisy measurements over the interval $t \in [0, N\Delta t]$. Observation of Figure 3 indicates that the dominant frequency peak is the fundamental frequency of the signal (i.e. heart rate). We apply a fast Fourier transform (FFT) to the data sequence with a Hamming window to obtain the power spectrum. The maximum frequency peak within a reasonable human heart rate range (0.5 to 2.5 Hz) is used to initialize the estimate of $\hat{\omega}_0 \triangleq \omega(t = N\Delta t)$.

Assuming $\hat{\omega}_0$ to be correct, the problem can be rewritten as a linear estimation problem to obtain the harmonic amplitudes and coefficients. Equation (4) is equivalently reparametrized as

$$y(t) = c_0 + \sum_{i=1}^{m}\left[a_i \sin(i\hat{\omega}_0 t) + b_i \cos(i\hat{\omega}_0 t)\right]. \qquad (6)$$

Collecting the measurements $z(t)$, $0 \le t \le N\Delta t$ we have $\tilde{\boldsymbol{z}} = \boldsymbol{A}\hat{\boldsymbol{x}}_0 + \boldsymbol{\nu}$, where $\tilde{\boldsymbol{z}} = [z(0), z(\Delta t), \ldots, z(N\Delta t)]^{\mathrm{T}}$, $\hat{\boldsymbol{x}}_0 = [c_0, a_1, \ldots, a_m, b_1, \ldots, b_m]^{\mathrm{T}}$, $\boldsymbol{\nu}$ is a vector of measurement noise, and

$$\boldsymbol{A}^{\mathrm{T}} = \begin{bmatrix}
1 & 1 & \ldots & 1 \\
\sin(\hat{\omega}_0 0) & \sin(\hat{\omega}_0 \Delta t) & \ldots & \sin(\hat{\omega}_0 N\Delta t) \\
\vdots & \vdots & & \vdots \\
\sin(m\hat{\omega}_0 0) & \sin(m\hat{\omega}_0 \Delta t) & \ldots & \sin(m\hat{\omega}_0 N\Delta t) \\
\cos(\hat{\omega}_0 0) & \cos(\hat{\omega}_0 \Delta t) & \ldots & \cos(\hat{\omega}_0 N\Delta t) \\
\vdots & \vdots & & \vdots \\
\cos(m\hat{\omega}_0 0) & \cos(m\hat{\omega}_0 \Delta t) & \ldots & \cos(m\hat{\omega}_0 N\Delta t)
\end{bmatrix},$$

to which the least squares estimate is obtained as $\hat{\tilde{\boldsymbol{x}}}_0 = (\boldsymbol{A}^{\mathrm{T}}\boldsymbol{A})^{-1}\boldsymbol{A}^{\mathrm{T}}\tilde{\boldsymbol{z}}$. $\hat{\boldsymbol{x}}(T|T)$ is then initialized with the values $\hat{c}(T) = \hat{c}_0$, $\hat{\omega}(T) = \hat{\omega}_0$, $\hat{r}_i(T) = (\hat{a}_i^2 + \hat{b}_i^2)^{-1/2}$, and $\hat{\theta}_i(T) = \arctan(\hat{b}_i, \hat{a}_i)$. The state estimate covariance is set to $\boldsymbol{P}(T|T) = \mathrm{diag}[\sigma_R^2/N, \sigma_1^2, \sigma_1^2/2^2, \ldots, \sigma_1^2/m^2, \sigma_\omega^2, 0.02\ \mathrm{rad}^2, \ldots, 0.02\ \mathrm{rad}^2]$. The relationship between the amplitude uncertainties is chosen to reflect the decreasing harmonic strength seen in Figure 3, while the phase uncertainties follow those used in (Parker and Anderson 1990). Parameters σ_1^2 and σ_ω^2 are determined through experimentation. Last we assume that the process noise covariance $\boldsymbol{Q}$ is diagonal with all values set to 10^{-4} except for q_ω, the entry corresponding to $\omega(t)$.

4.2. Simulation Studies

Three simulation studies were conducted to evaluate the capabilities of the EKF, AR filter, and fading AR filter to the primary sources of random error in the system: measurement noise and heart rate variability. For illustrative purposes, the filters were also compared with the WFLC estimator (Riviere et al. 1998) and a simpler method of using the previous cardiac

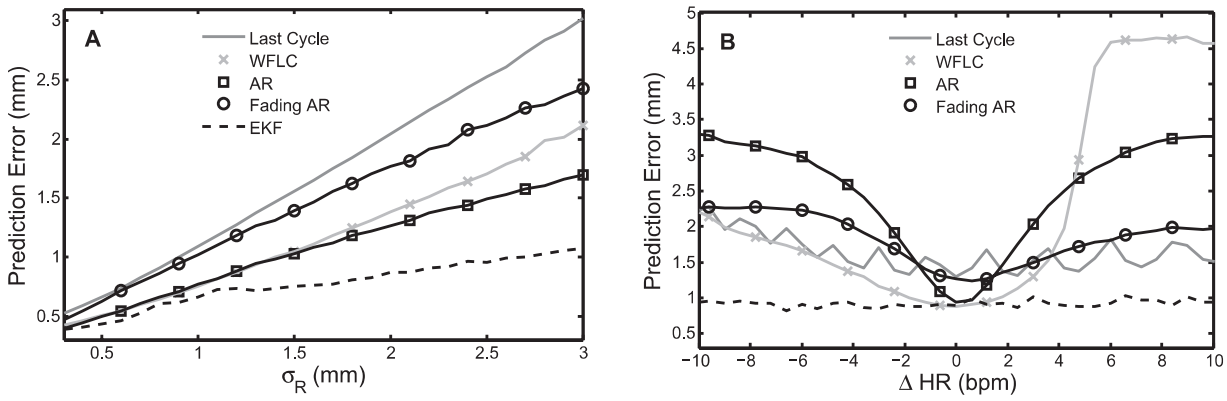

Fig. 9. RMS prediction error results for parametric simulation studies. A, Error for varying measurement noise; B, error for step heart rate changes.

cycle trajectory for the prediction of the next. A more sophisticated version of the latter method, called the *last cycle* here, was used successfully in a beating heart tracking system (Bebek and Cavusoglu 2007).

In the first simulation, we subjected the predictors to varying levels of measurement noise on a fixed-rate trajectory (60 bpm). The mitral annulus trajectory of Figure 2 was reinterpolated to 28 Hz and corrupted by additive zero-mean Gaussian noise with standard deviation $0.3 \leq \sigma_R \leq 3$ mm. Each predictor was then given 30 sec of data to initialize and performance was judged for the following 10 sec on 1-sample ahead predictions.

The RMSEs for each predictor averaged across 100 Monte Carlo trials are shown in Figure 9A. The EKF, WFLC, AR, and fading AR filtering methods clearly give higher-accuracy predictions than the inherent uncertainty of the measurements, with the EKF doing the best. As expected, the last cycle method had error statistics equal to σ_R since it attempts no smoothing. It should be noted that the fading AR filter was tuned with $f = 0.985$ in order to achieve errors that are approximately equal to σ_R. This setting represents the lowest reasonable value for the fading AR filter since a lower value would give performance below the last cycle method. The EKF was run with $m = 8$ harmonics, $N = 280$ initialization points (10 seconds), $\sigma_1^2 = 2$ mm^2, $\sigma_\omega^2 = 0.11$ (rad/s)^2 (roughly twice the frequency resolution of the FFT), and $q_\omega = 10^{-3}$ (rad/s)^2. The WFLC was initialized in the same manner as the EKF, run with $m = 8$ harmonics, and experimentally set with its adaptive gain parameters $\mu_0 = 7 \times 10^{-6}$ and $\mu_1 = 0.03$ for best performance in this and subsequent simulations.

In a second parametric simulation study, we gauged the tolerance of each predictor to a sudden change in heart rate. A trajectory was assembled by piecing together 30 seconds of heart motion at 60 bpm and 10 seconds of motion at (60 + ΔHR) bpm. The second portion of the trajectory was generated by compression/dilation of the target trajectory in Figure 2 to obtain the desired heart rate. Like before, the composite trajectory was reinterpolated to 28 Hz and corrupted with $\sigma_R = 1.30$ mm noise; the measured uncertainty of the 3DUS annulus estimates (Novotny 2007). The last 10 seconds were used to evaluate performance. A reasonable range of -10 bpm $\leq \Delta$HR ≤ 10 bpm was determined from clinical RR data (Figure 10), which is discussed later.

Figure 9B shows the mean RMSEs for each predictor across 100 Monte Carlo trials. The EKF provided better predictions than the other four methods. It was also the only method that yielded sub-σ_R error for the majority of heart rate changes. The WFLC had similar accuracies to the EKF at small ΔHR but showed slow convergence to the new heart rate for ΔHR > 4 bpm. As expected, the accuracy of the AR filter also approached that of the EKF for small ΔHR and quickly degraded as ΔHR increased. Exponential weighting of the measurements allowed the filter to adjust to changes in the trajectory, as demonstrated by the fading AR filter's superior performance over the AR filter for large ΔHR. However, this adaptability decreased accuracy when the trajectory did not change significantly. Finally, the last cycle method showed performance comparable to the fading AR filter. For this simulation all filter parameters were chosen to be the same as in the previous simulation, with the exception of $q_\omega = 5 \times 10^{-3}$ (rad/s)^2 and $\sigma_\omega^2 = 1$ (rad/s)^2 for the EKF.

Finally, to investigate the performance of each predictor to the more realistic case of a continuously changing heart rate, we modulated the period of the annulus trajectory with clinically obtained cardiac cycle records. Annotated ECG records for five human subjects were selected from the MIT–BIH Normal Sinus Rhythm Database (Goldberger et al. 2000) and composite mitral valve trajectories were generated in a manner similar to the previous simulation study. Noise-corrupted measurements were generated as before, with $\sigma_R = 1.30$ mm.

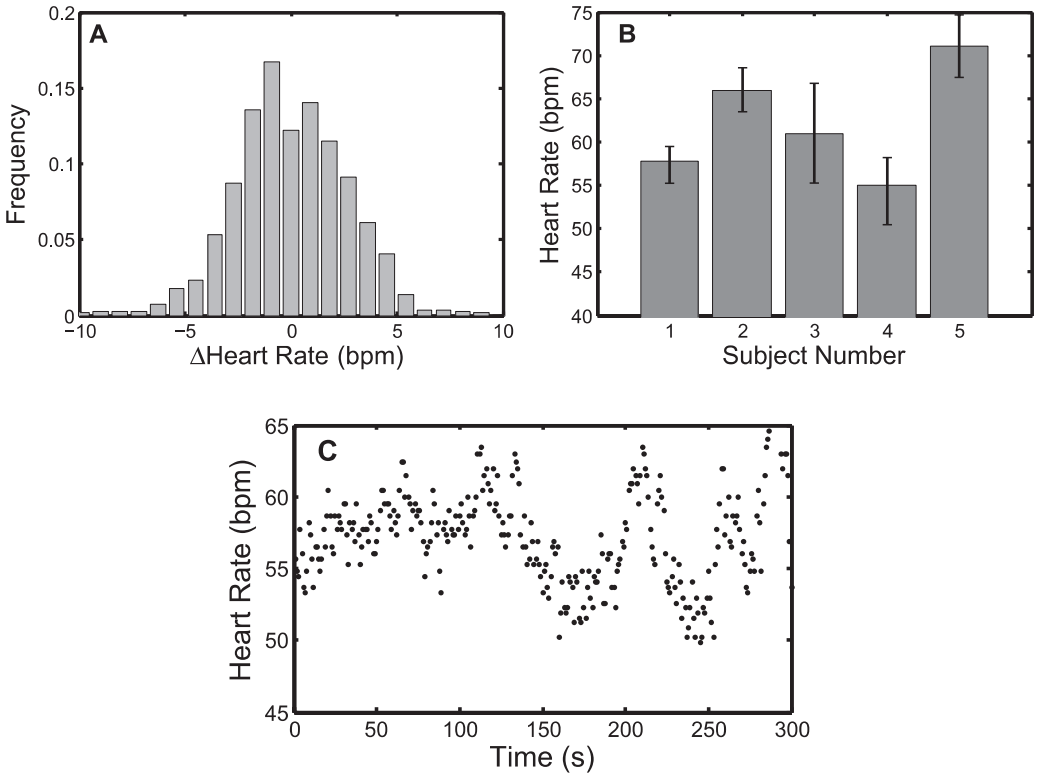

Fig. 10. Heart rate statistics on five human subjects. A, Distribution of beat-to-beat variation; B, subject-specific mean ± one standard deviation; C, heart rate data for subject 1.

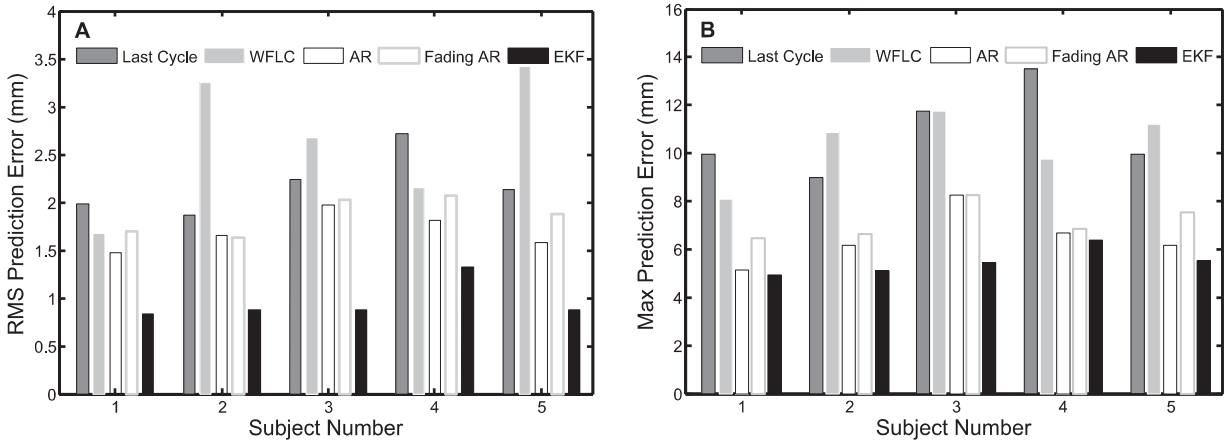

Fig. 11. Comparison of prediction error results on clinical data using last cycle, WFLC, AR filter, fading AR filter, and EKF: A, RMSE; B, maximum absolute error.

Summary statistics for each subject are presented in Figure 10B and an example of the beat-to-beat heart rate for subject number 1 is shown in Figure 10C.

Results from this study indicate that the EKF is more suited to tracking and prediction in this application than the other four methods because it adjusts to rapid changes in heart rate through explicit modeling of quasiperiodicity (Figure 11). Interestingly, the AR filter showed moderately better performance than the fading AR filter. The reason for this is that the AR filter locked on to an "average" trajectory for each

subject while the fading AR filter continuously readjusted to more recent noisy data. Ultimately, deviations from the "average" motion were less than the measurement noise. The last cycle method performed worse for similar reasons: persistent variations in heart rate and measurement noise degraded the accuracy of the previous cycle as a predictor for the next. The slow convergence of the WFLC to changing heart rates caused it to have severely degraded performance.

5. Performance Evaluation in a Surgical Task

In order to quantify the amount of assistance that motion compensation provides to operators working on a moving target, we conducted two studies of user performance with the MCI in an *in vitro* setting. These studies additionally provide insight on how sensitive performance is to potential shortcomings of a 3DUS-guided system. Specifically, user study 1 determined the extent to which user performance is dependent on time delay and random positional error. These error types are of interest because both positional precision errors and computational delays develop while generating 3DUS tracking data in real-time. User study 2 investigated user performance with EKF delay compensation on targets with both fixed and variable heart rates. Subjects performed a drawing task on a moving target using the MCI under different tracking conditions. A total of 18 test subjects (14 male and 4 female, aged 22 to 36; 8 subjects for user study 1 and 10 subjects for user study 2) voluntarily participated following informed consent under a protocol approved by the University Institutional Review Board.

5.1. Experimental Setup

The tests were run on a setup which emulates the intended surgical environment. To simulate the moving mitral valve, a target platform was mounted on a cam-driven device that replicates the 1D motion of the mitral annulus centroid as measured from the 3DUS tracking data. During trials, a paper target was affixed to this platform to record the subject's drawing. A 0.5 cm hard foam rubber pad between the target paper and target platform provided a small measure of compliance. In combination with the pen used in the trials, the pad had a stiffness of 4.5 N mm^{-1}. For the purposes of this experiment, the cam was used to simulate a heart rate of 60 bpm. Opposite the target platform, the MCI was mounted in a gimbal allowing both angular motion and translation towards and away from the target (Figure 12). A rod was mounted on the MCI with a ballpoint pen affixed to one end and a force sensor incorporated along its length. The force sensor had a stiffness of 10 N mm^{-1}. In place of the 3DUS-based tracking and controls algorithms that would be used in surgery, target position was directly measured at 1 kHz by a contact arm with a potentiometer attached to the target platform. This sensing method provided the robust

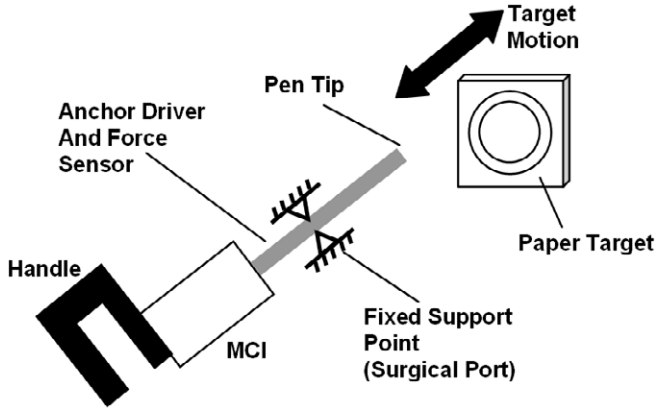

Fig. 12. User trial setup. The bold arrow indicates the cyclical motion of the simulated valve target. The tool shaft has free rotational and sliding motion at the fixed support point.

tracking data necessary to evaluate the efficacy of MCI mitral valve tracking and the performance of predictive filtering algorithms.

5.2. User Task

Subjects were instructed to draw a circle on the moving target platform. The circle must be drawn between two concentric target circles with 2.29 and 2.92 cm diameters. Subjects started at the top of the circle and proceeded in the clockwise direction. If the pen bounced off the target surface or outside the target circles, the subject was instructed to continue drawing from the clockwise-most mark they made between the target circles. Subjects could only draw around the circle once. They could not go back to draw in gaps they originally missed. To prevent subjects from spending an inordinate amount of time on the task, a 25-second time limit was set. All subjects finished every trial before the time expired.

Both dexterity and force metrics were used to evaluate subject performance. In all trials, the quality of each circle drawn was characterized by digitally scanning the target and computing the drawn line's "angular surround" value. This value indicates what percentage of 360 discrete 1° arc segments the user's line covered between the two concentric target circles. This metric reflects positioning accuracy both in the plane of the target and in tracking the motion of the target perpendicular to this plane. If the user-drawn line strayed from the concentric target circles, it did not contribute to the angular surround metric and resulted in a lower score. Similarly, if the user and MCI were unable to track the motion of the target, the pen tended to bounce off of the target surface, producing widely spaced marks and a low angular surround score.

The axial force applied by the subject to the target was also recorded for 4 of the 8 test subjects in user study 1 and 9 of

the 10 subjects in user study 2. In all 18 cases, subjects were informed of both evaluation metrics. They were instructed that their foremost objective was to draw continuous circles conforming to the angular surround metric and only secondly to use the minimum amount of force necessary.

This task was selected to emulate the motion requirements of placing a surgical anchor. In order to apply the surgical anchors developed for this procedure, the tip of the anchor driver, consisting of 14 gauge hypodermic tubing, must be accurately located and pressed against the target surface with a force of at least 1.5 N (Wagner et al. 2006). This contact must be maintained for several seconds as the surgeon inserts the anchor, tests whether it is deployed properly, and then releases the anchor. This process requires a combination of accuracy and prolonged contact with the surface. At the same time, forces must be minimized so as not to cause damage to the valve.

5.3. User Study 1: Tracking with Time Delay and Positional Error

Subjects of this study completed the task under eight different tracking conditions. In the "solid" condition, the motion of the MCI was rigidly locked in order to simulate a traditional, solid endoscopic tool. For the "baseline" MCI tracking condition, the current position of the target (RMSE = 0.015 mm) was sent to the MCI as a position command. This baseline condition resulted in a 14 ms delay. The remaining six tracking conditions were divided into two groups of three conditions corresponding to differing levels of the considered error.

Random positional error was simulated by the superposition of a time-varying error value with the cam position command used in the baseline tracking state. A new positional error was calculated at 8 Hz. These errors were uniformly distributed random values ($[-1, 1]$), multiplied by an amplitude factor of 0.35, 0.70, or 1.05 mm.

Delay error was implemented by recording the cam tracking position and holding it for a specified period before sending this position to the MCI as a motion command. For this set of trials, the three levels of added delay used were 25, 35, and 45 ms. Including the MCI lag time of 14 ms, the effective delay settings were 39, 49, and 59 ms. This range of times was chosen as representative of the imaging and transmission delays associated with real-time 3DUS-guided procedures (Novotny et al. 2007).

5.4. User Study 2: Delay-compensated Tracking with Heart Rate Variation and Measurement Noise

To test the EKF under conditions similar to those seen in 3DUS-guided procedures, the 1 kHz measurements of target position were downsampled to 28 Hz and corrupted by additive Gaussian noise $\sim\mathcal{N}\left(0, 1.30^2 \text{ mm}\right)$(Novotny 2007). The target was commanded to beat at 60 bpm or with a variable rate that had additive Gaussian beat-to-beat fluctuations $\sim \mathcal{N}\left(0, \sigma_{\text{HR}}^2\right)$. Finally, a time delay of 39, 59, or 89 ms was injected into the measurements to simulate the delays encountered with 3DUS. Note that the 89 ms delay exceeds the delays used in user study 1 to also account for the additional computational delays from instrument and target segmentation in 3DUS (Novotny 2007).

Overall, the subjects of this study performed the task under 10 tracking conditions, all with varying amounts of time delay. Three tracking conditions evaluated user performance with EKF delay compensation against 39, 59, and 89 ms time-delayed, noise-corrupted measurements on a 60 bpm target. For comparison, three tracking conditions were conducted on the same set of delays but without delay compensation and without measurement noise.

The remaining four tracking conditions evaluated user performance with EKF delay compensation on a variable heart rate target. Measurements were delayed by 89 ms and noise-corrupted. The four levels of beat-to-beat variation were chosen as $\sigma_{\text{HR}} = \{1.0, 2.0, 3.0, 4.0\}$ bpm. These values span those observed in the clinical RR data from Section 4.2. An illustrative example of this tracking condition with $\sigma_{\text{HR}} = 4.0$ bpm and the subsequent behavior of the EKF and the MCI is shown in Figure 13.

5.5. Testing Protocol

Each subject test consisted of a practice period followed by the trials corresponding to the tracking conditions of their study. Practice was intended to familiarize the test subject with the MCI and the evaluation task in order to bring subjects to a uniform level of ability and to limit learning effects during trials. Practice was divided into three one-minute segments during which the subject was free to experiment with using the MCI to draw on a target paper. During the first minute of training, the target was stationary and the tool was set in the solid condition. The second minute of training involved a moving target and a solid tool. In the third and final minute, the target was moving and the MCI was in the baseline tracking condition. Following the completion of training, each test subject ran through the trials corresponding to the tracking conditions of their study. The order in which these conditions were administered was varied between trials using a balanced Latin square to minimize the effects of between-trial carry-over and learning on collected data.

The means of collected angular surround error metric were compared for statistically significant differences using the SPSS statistical analysis software package (Version 14.0, SPSS Inc., Chicago, IL, USA). These comparisons were made using t-tests and analysis of variance (ANOVA) with a least significant difference (LSD) *post hoc* test. In all cases, significance corresponds to $p < 0.05$.

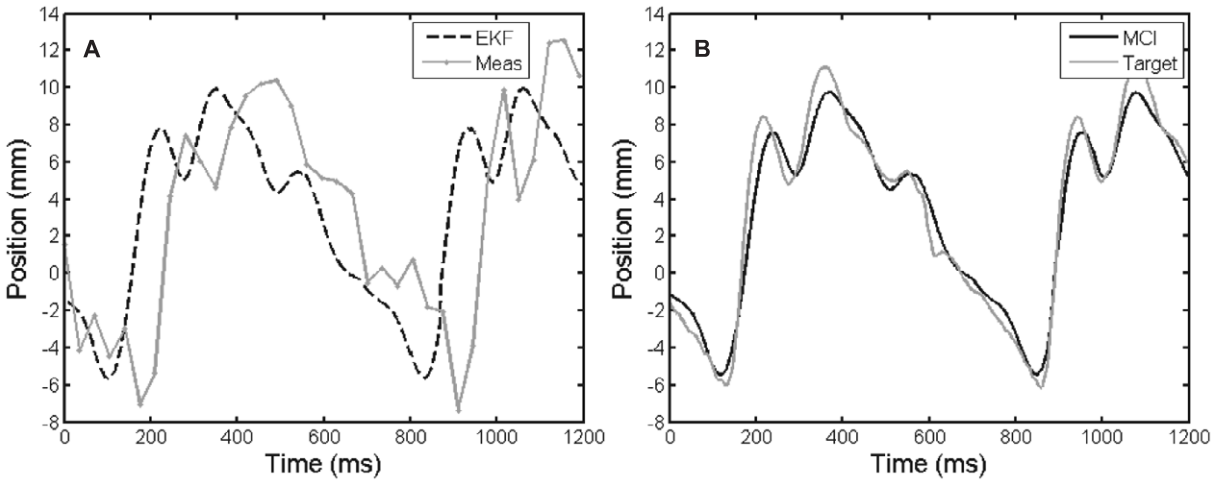

Fig. 13. Variable heart rate tracking condition ($\sigma_{\mathrm{HR}} = 4.0$). A, Noise-corrupted, time-delayed measurements and the resulting EKF predictions; B, MCI position and the true position of the target.

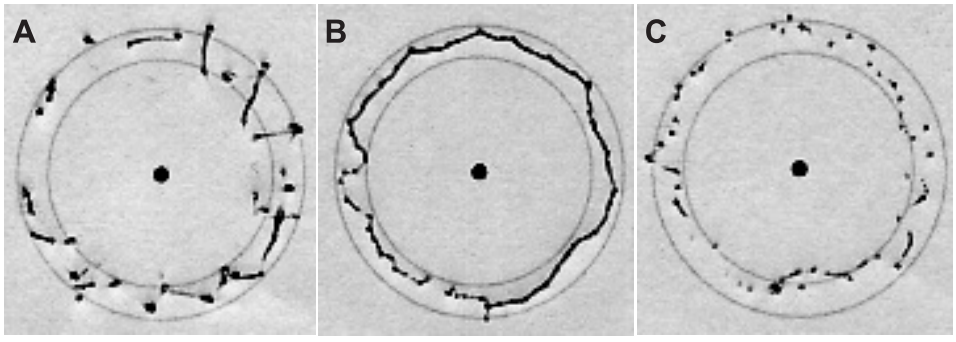

Fig. 14. Targets with median angular surround scores under solid (A), baseline (B), and 59 ms time-delayed (C) tracking conditions. The baseline circle is the most continuous and round. While the solid target has more continuous lines than the time-delayed target, it also has heavy dots and dimpling indicative of high forces.

5.6. User Study 1 Results

The results of this study indicate that the MCI can provide significant assistance while operating on a moving target. Figure 14 shows typical target results. Mean angular surround metric scores (Figure 15) for the baseline tracking condition ($81.9 \pm 4.5\%$ (mean $\pm$ standard error)) were over 50% greater than for the solid condition ($53.8 \pm 5.0\%$), with clear statistical significance ($t(14) = 0.1987$, $p = 0.0009$). Similarly, axial force data indicates that subjects applied less than 50% as much force and spent less time at elevated forces under the baseline MCI tracking condition. Aggregating the data from all four subjects (Figure 16), the 90th percentile for force samples using the solid tool (17.5 N) is roughly double the 90th percentile for the baseline tracking condition (8.5 N).

A comparison of mean angular surround metric scores related to the amount of delay error (Figure 17) demonstrates decreases in performance ranging from 67% to 33% with increasing delay ($f(3, 28) = 16.005$, $p < 0.001$). Statistically significant differences were indicated between the baseline condition and tracking with delays of 39 ms ($61.2 \pm 1.5\%$), 49 ms ($56.0 \pm 3.7\%$), and 59 ms ($48.8 \pm 3.8\%$). A significant difference also exists between the means of the 39 and 59 ms tracking conditions ($p = 0.02$). Trend analysis indicates that the data is well fit by a linear model ($p < 0.001$).

An analysis of mean angular surround metric scores related to positional error did not demonstrate significant differences under ANOVA analysis ($f(3, 28) = 0.638$, $p = 0.597$). As seen in Figure 18, the mean score under the baseline condition differed very little from those with error amplitude factors of

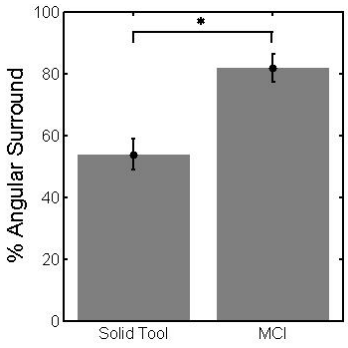

Fig. 15. Performance under baseline MCI and solid tool tracking conditions. Error bars indicate standard error. "*" marks statistical significance.

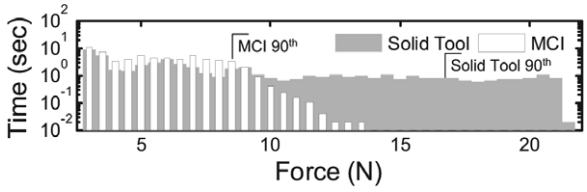

Fig. 16. Histogram of applied force samples under baseline MCI and solid tool tracking conditions. Smaller forces are consistently applied under the baseline condition. 90th percentiles are marked. The zero force bin is not shown.

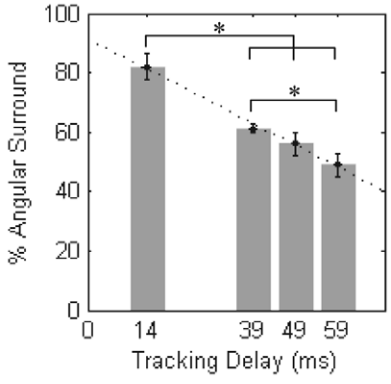

Fig. 17. Performance under delay tracking conditions. Error bars indicate standard error. Asterisks indicate statistical significance. The dotted line shows the fitted linear model ($R^2 = 0.9937$).

0.35 mm (81.7 ± 3.3%), 0.70 mm (78.0 ± 5.5%), and 1.05 mm (74.2 ± 4.6%).

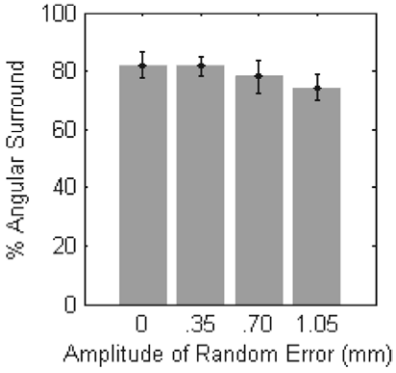

Fig. 18. Performance under positional error tracking conditions. Error bars indicate standard error.

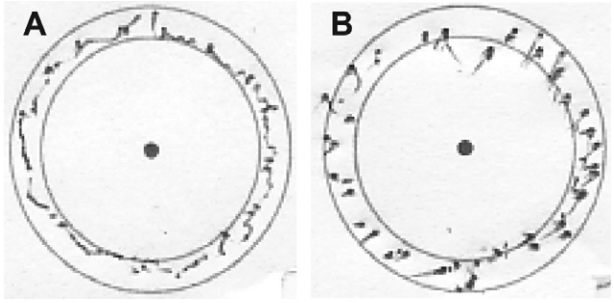

Fig. 19. Targets with median angular surround scores under 89 ms EKF delay-compensated (A) and time-delayed (B) tracking conditions. The delay-compensated example is more continuous and round than the time-delayed example. The latter is fragmented and has heavy dots and dimpling that are indicative of high forces.

5.7. User Study 2 Results

Results from this study demonstrate that the EKF is an effective approach for compensating time delay. Figure 19 shows typical target results. Mean angular surround scores for EKF delay-compensated tracking (Figure 20) were similar to the baseline tracking condition from user study 1 for all three conditions of 39 ms (83.1 ± 3.7%), 59 ms (85.2 ± 2.6%), and 89 ms (84.4 ± 3.8%). Likewise, delay-compensated tracking showed performance increases over delayed tracking ranging from 13% to 24%. The mean angular scores for delayed tracking were 70.4 ± 4.8%, 61.1 ± 4.6%, and 61.0 ± 5.2% for 39, 59, and 89 ms delays, respectively. Statistically significant differences at $p < 0.05$ were observed between the mean scores of each delay-compensated tracking condition to each delayed tracking condition. Delay-compensated tracking also yielded smaller axial forces than those observed for delayed tracking (Figure 21).

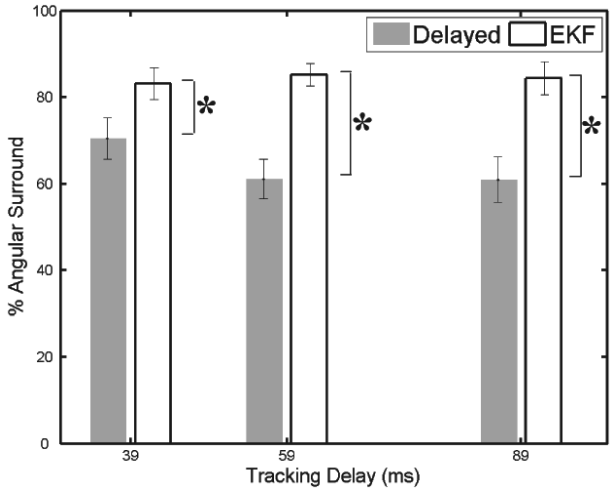

Fig. 20. Performance with EKF delay compensation and without. Error bars indicate standard error. Asterisks indicate statistical significance.

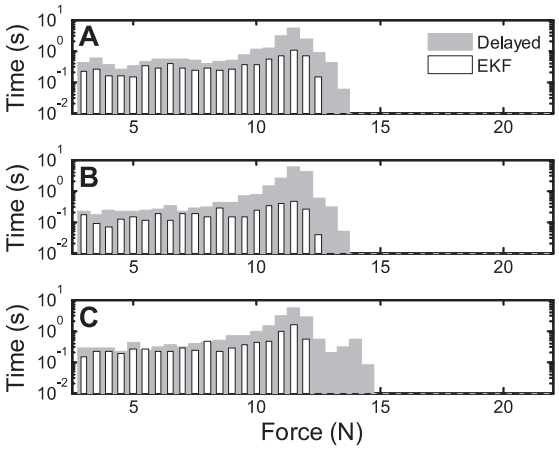

Fig. 21. Force application with and without delay compensation for delays of (A) 39 ms, (B) 59 ms, and (C) 89 ms. Smaller forces are consistently applied under the delay-compensated tracking conditions. The zero force bin is not shown.

An analysis of mean angular surround metric scores related to heart rate variability did not demonstrate significant differences under ANOVA analysis ($f(3, 36) = 0.705$, $p = 0.555$). As seen in Figure 22, the performance against a fixed-rate target was comparable to that against a variable rate target with σ_{HR} equal to 1.0 bpm ($85.9 \pm 2.6\%$), 2.0 bpm ($84.1 \pm 2.9\%$), 3.0 bpm ($80.5 \pm 3.6\%$), and 4.0 bpm ($81.5 \pm 2.5\%$).

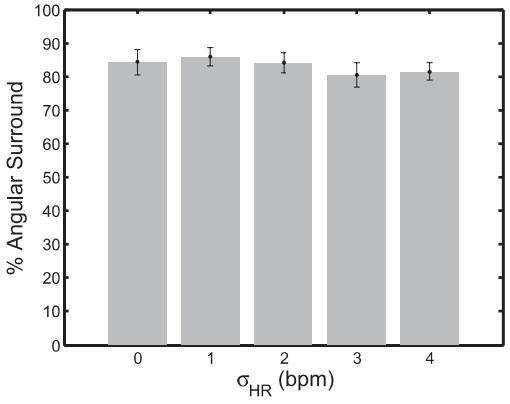

Fig. 22. Performance under variable heart rate tracking conditions. Error bars indicate standard error.

6. System Accuracy Under 3D Ultrasound Guidance

Water tank experiments were conducted to measure the motion synchronization accuracy of the system under 3DUS guidance. To do these, a real-time 3DUS target segmentation algorithm was first incorporated into the system to provide position measurements to the EKF. The target was set to be an $\boldsymbol{X}$-shaped fiducial that can be easily mounted to an annuloplasty ring (Figure 23B). This fiducial was specifically chosen because detecting two intersecting lines is suited for an existing real-time 3DUS segmentation algorithm based on the modified Radon transform (Novotny 2007). This algorithm is known to provide target position measurements with 1.30 mm RMS accuracy under *in vitro* conditions. Because higher noise may be present *in vivo*, we also tested this system in two other noise conditions in which large Gaussian terms $\sim \mathcal{N}(0, \sigma)$, $\sigma = \{2.0, 3.0\}$ mm were added to the segmented target positions. This yielded three noise conditions with overall RMS accuracies of $\sigma_R = \{1.30, 3.30, 4.30\}$ mm. We anticipate *in vivo* noise levels to be less than 3.5 mm RMS.

6.1. Experimental Setup

The target and instrument were imaged by a real-time 3DUS probe in a water tank at 28 Hz (Figure 23A). Data was streamed from the ultrasound machine (SONOS 7500, Philips Medical, Andover, MA) to a computer over an ethernet connection. The stream was captured by the computer and passed to a graphics processing unit (8800GTS, nVidia Corp, Santa Clara, CA) where the volumes were automatically segmented using the modified Radon transform segmentation algorithm (Novotny 2007) to obtain target position measurements. The measurements were corrupted with additive Gaussian noise (depending on the noise condition), then

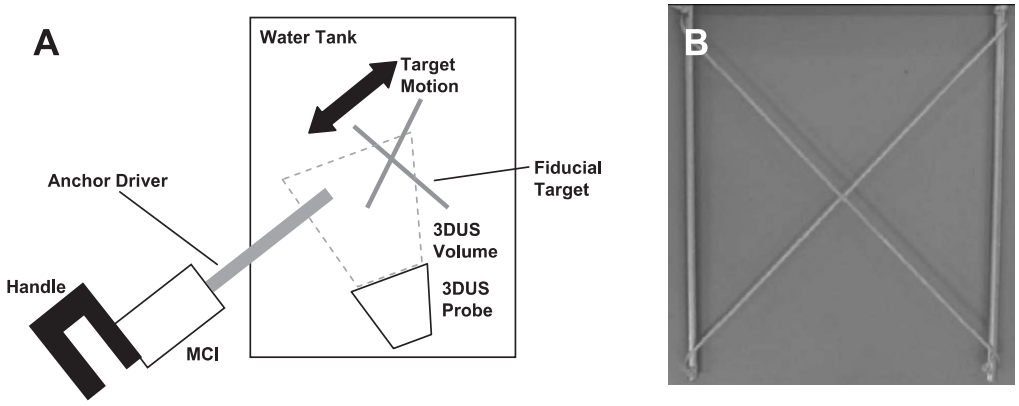

Fig. 23. (A) Setup for motion synchronization experiments and (B) X-shaped fiducial target.

passed to a third thread for EKF processing. This returned predicted target positions 82 ms in the future (68 ms from imaging and segmentation, 14 ms from instrument lag time) and were sent to the 1 kHz PID control loop for the robot.

Five motion synchronization trials were conducted per noise condition. The target simulated mitral valve motion at 60 bpm with the cam-driven mechanism described in Section 5.1. MCI position and true target position were measured at 1 kHz with potentiometers and stored for offline processing. Target measurements and EKF predictions were acquired at the 3DUS sampling rate (28 Hz) and stored as well. Accuracy evaluation was performed on the 20 s of tracking data following 20 s of EKF initialization.

6.2. Results

Results from this experiment demonstrate that the motion compensation system is robust to the noise levels present in 3DUS imaging. Typical data and results for each noise condition are shown in Figure 24 and overall performance is summarized in Figure 25. Mean RMS accuracies for EKF predictions were 1.01 ± 0.02, 1.08 ± 0.03, and 1.22 ± 0.05 mm for $\sigma_R = \{1.30, 3.30, 4.30\}$ mm (respectively) and showed statistically significant reductions in measurement error ($p < 0.0001$ for each case). Mean MCI synchronization errors were 1.15 ± 0.04, 1.23 ± 0.06, and 1.28 ± 0.10 mm RMS for the three noise conditions. The mean RMSE added to the system from the MCI following the prescribed EKF trajectory was 0.14 ± 0.01 mm and this difference was statistically significant ($p = 0.002$). The RMS noise accuracy for the $\boldsymbol{X}$-tracking algorithm observed across all 15 trials was 1.30 ± 0.01 mm, in agreement with previous reports (Novotny 2007).

7. Discussion

User trials demonstrate that the motion compensation system described here is an effective aid for surgical tasks on beating mitral valve targets. Motion synchronization allowed users to operate with both increased dexterity and reduced forces. In user study 1, the angular surround scores indicative of dexterity increased by approximately 50% between the solid tool and baseline tracking conditions while the 90th percentile for force decreased by a similar ratio. These performance gains were maintained when increasing levels of positional error were inserted into the system, but disappeared when uncompensated time delays were added to the tracking data.

The strong dependence of user performance on delay emphasizes that precise timing is essential for successful motion compensation. Left uncompensated, tracking delays produced serious errors during the rapid recoil of the target associated with the relaxation of the left ventricle. During this motion the valve moves towards the operator, nearly covering its entire 18 mm range of motion. Owing to the delay, the MCI continued to servo towards the oncoming target and abruptly collided with it. Qualitatively, while subjects did not always notice added positional error, they universally correctly identified and expressed dismay over tracking delay.

The introduction of a predictive EKF into the system effectively removed the performance decrease associated with tracking delays. The angular surround scores for delay-compensated tracking conditions exceeded their respective delayed tracking conditions by 13% to 24%, depending on the degree of delay. The forces applied in the delay-compensated conditions also decreased relative to the delayed conditions, although to a lesser extent than seen between baseline tracking and the solid tool. The delay-compensated scores of user study 2 marginally exceeded the baseline scores of user study 1 by an amount that is consistent with the removal of the MCI's 14 ms lag time (as predicted by the linear model from user

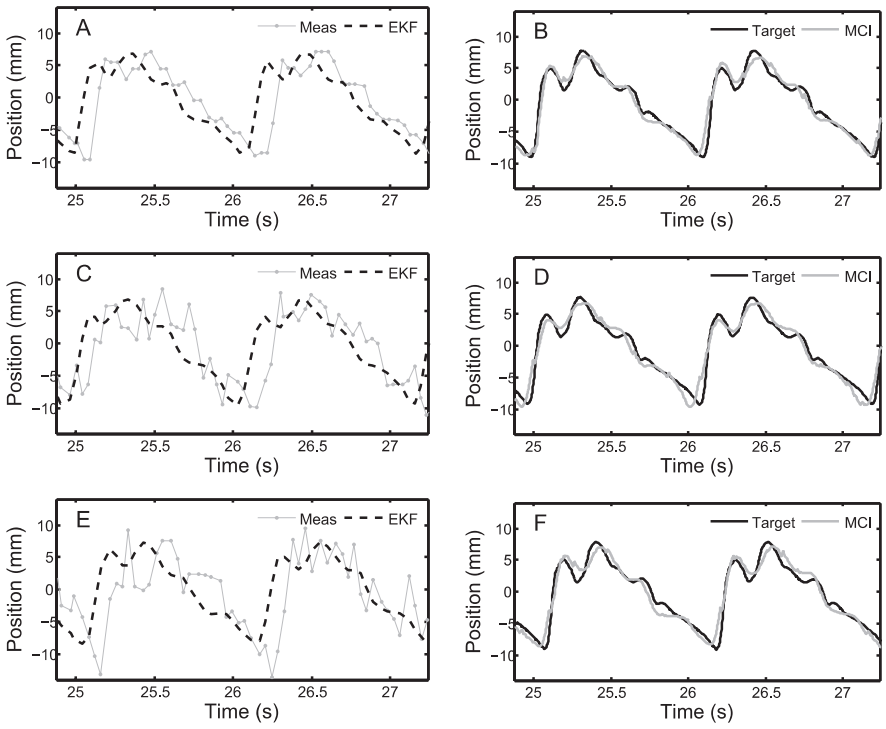

Fig. 24. Examples of data, filtering, and motion tracking results. Left column shows measurements from 3DUS target segmentation and EKF predicted target trajectories for each noise condition: (A) $\sigma_R = 1.30$ mm; (C) $\sigma_R = 3.30$ mm; and (E) $\sigma_R = 4.30$ mm. Corresponding MCI and true target trajectories for each noise condition are shown in the right column: (B) $\sigma_R = 1.30$ mm; (D) $\sigma_R = 3.30$ mm; (F) $\sigma_R = 4.30$ mm.

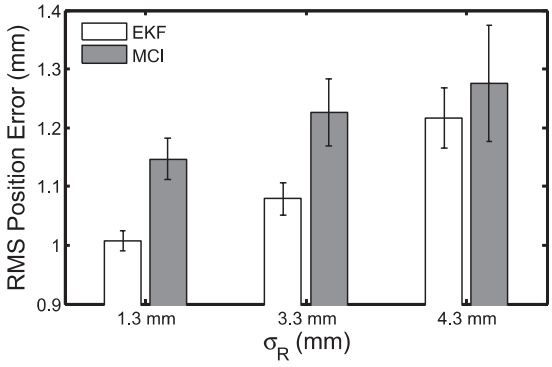

Fig. 25. EKF prediction performance and MCI tracking performance over three increasing noise conditions. Error bars indicate standard error.

study 1). However, caution must be taken when making a direct comparison of the angular surround scores between user studies 1 and 2 because of differing amounts of in-trial training (8 trials in user study 1 and 10 trials in user study 2). It is reasonable to conclude that EKF delay compensation restored the MCI tracking performance to at least baseline tracking conditions in cases of delay. Moreover, the performance increases observed for delay-compensated tracking showed no degradation across the range of time delays expected for 3DUS-guided procedures.

Surprisingly, the levels of random positional error explored in the user studies had a negligible effect on the benefits of MCI tracking. No statistically significant differences were found between user performance under all three positional error conditions and the baseline condition. This result suggests that if the timing with the valve's sudden recoil is maintained, the value of the MCI's tracking is robust in the face of positional errors as high as ± 1 mm. This may be attributed, in part, to compliance in the system: in the MCI's mechanism and in the foam pad on the target platform, which introduce compliance similar in nature to that in the heart. An *in vivo* characterization of MCI performance is necessary to determine whether compliance acts in the same way to assist operation on the mitral valve annulus.

Heart rate variations in the target motion also provided no significant degradation in user performance over the clinically obtained beat-to-beat statistics (Figure 10). This is not surprising given that the EKF estimates did not degrade on these levels of heart rate variation. In some respects, tracking a variable rate target was another approach to injecting

positional error into the system since the filter remained in phase with the target (see, for instance, Figure 13). We anticipate that extremely large beat-to-beat variations will cause the EKF to break synchronization with the target as the linearization of the measurement function, $h(\boldsymbol{x}(t))$, becomes inappropriate. Should more aggressive heart rate variations be encountered in future *in vivo* experiments, there are several potential routes for improvement. First, heart rate variability could be reduced through drug treatment and/or electrical pacing of the heart. Second, following Bebek and Cavusoglu (2007), ECG information could be used to obtain direct measurements of heart rate. This could reduce the effect of the non-linearity in $h(\boldsymbol{x}(t))$ since, as mentioned before, perfect knowledge of $\omega(t)$ turns this into a linear estimation problem (6). Note that the EKF, a non-linear filtering approach, would still be useful with these measurements because they are sampled at discrete intervals but $\omega(t)$ is continuous.

The efficacy of this motion compensation system is based on the assumption that the motion which must be tracked is modeled well by a 1D approximation. Previous research on augmenting surgical procedures with robotic tracking have focused on coronary artery bypass grafting, requiring a 3D model of surface motion (Nakamura et al. 2001; Ortmaier et al. 2005; Ginhoux et al. 2006; Bebek and Cavusoglu 2007). The characterization of the mitral valve annulus motion above shows that the fast motion of the mitral valve is primarily along one axis; however, we also observe approximately 2 mm of off-axis motion (Figure 2). Our preliminary tests suggest that the passive compliance of the mitral valve annulus and the surgeon's hand will assist with these minor off-axis deviations. In addition, *in vivo* conditions will present slow out-of-axis motions owing to the respiration of the subject. Fortunately, surgeons already overcome this in current heart procedures, either by manually compensating for the slow motion or momentarily stopping the controlled ventilation to the patient.

In this work, particular emphasis has been placed on predictive filtering to mitigate the dominant sources of tracking error in the system: 3DUS delay and noise. Using this approach with the MCI in a 3DUS-guided servoing task, we achieved synchronization accuracies of less than 1.3 mm RMS in the presence of large measurement noise (1.3–4.3 mm RMS) and 82 ms of system delay. This is a significant improvement over the approximately 4.6 mm RMSEs that would be incurred in a time-delayed but otherwise noiseless tracking system. Although errors in the current system are low, there is potential for further lowering error through a more sophisticated controller. Repetitive control methods are well suited for quasi-periodic servoing tasks (Hara et al. 1988; Chew and Tomizuka 1990; Kempf et al. 1993; Horowitz 1993) and model predictive control has shown promising results in external beating heart tracking (Ginhoux et al. 2006; Bebek and Cavusoglu 2007). These techniques may reduce the 0.14 mm tracking error that is attributable to the PID controller used in the current system.

Acknowledgements

This work is supported by the US National Institutes of Health under grant NIH R01 HL073647-01. The authors would like to thank the test subjects who graciously donated their time to this study, and Nikolay Vasilyev, Pedro J. del Nido, Mahdi Tavakoli, Peter Hammer, and Petr Jordan for many insightful medical and technical conversations.

References

Bebek, O. and Cavusoglu, M. (2007). Intelligent control algorithms for robotic assisted beating heart surgery. *IEEE Transactions on Robotics*, **23**(3): 468–480.

Bellinger, D., Wypij, D., Kuban, K., Rappaport, L., Hickey, P., Wernovsky, G., Jonas, R. and Newburger, J. (1999). Developmental and neurological status of children at 4 years of age after heart surgery with hypothermic circulatory arrest or low-flow cardiopulmonary bypass. *Circulation*, **100**: 526–532.

Cannon, J., Stoll, J., Salgo, I., Knowles, H., Howe, R., Dupont, P., Marx, G. and del Nido, P. (2003). Real-time three-dimensional ultrasound for guiding surgical tasks. *Computer Aided Surgery*, **8**(2): 82–90.

Chew, K. K. and Tomizuka, M. (1990). Digital control of repetitive errors in disk drive systems. *IEEE Control Systems Magazine*, **10**(1): 16–20.

Franke, T., Bebek, O., and Cavusoglu, C. (2007). Improved prediction of heart motion using an adaptive filter for robot assisted beating heart surgery. *Proceedings of IEEE Intelligent Robots and Systems (IROS '07)*, San Diego, CA, pp. 509–515.

Ginhoux, R., Gangloff, J., de Mathelin, M., Soler, L., Sanchez, M. A. and Marescaux, J. (2006). Active filtering of physiological motion in robotized surgery using predictive control. *IEEE Transactions on Robotics*, **21**(1): 27–79.

Goldberger, A., Amaral, L., Glass, L., Hausdorff, J., Ivanov, P., Mark, R., Mietus, J., Moody, G., Peng, C.-K. and Stanley, H. (2000). PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. *Circulation*, **101**(23): e215–e220.

Gorman, J., III, et al. (1996). Dynamic three-dimensional imaging of the mitral valve and left ventricle by rapid sonomicrometry array localization. *Journal of Thoracic and Cardiovascular Surgery*, **112**(3): 712–726.

Hara, S., Yamamoto, Y., Omata, T. and Nakano, M. (1988). Repetitive control system: a new type servo system for periodic exogenous signals. *IEEE Transactions on Automatic Control*, **33**(7): 659–668.

Horowitz, R. (1993). Learning control of robot manipulators. *Journal of Dynamic Systems, Measurement, and Control*, **115**(2B): 402–411.

Kamigaki, M. and Goldschlager, N. (1972). Echocardiographic analysis of mitral valve motion in atrial septal defect. *American Journal of Cardiology*, **30**: 343–348.

Kempf, C., Messner, W. C., Tomizuka, M. and Horowitz, R. (1993). Comparison of four discrete-time repetitive control algorithms. *IEEE Control Systems Magazine*, **13**(6): 48–54.

Murkin, J., Boyd, W., Ganapathy, S., Adams, S. and Peterson, R. (1999). Beating heart surgery: why expect less central nervous system morbidity? *Annals of Thoracic Surgery*, **68**: 1498–1501.

Nakamura, Y., Kishi, K. and Kawakami, H. (2001). Heartbeat synchronization for robotic cardiac surgery. *Proceedings of IEEE International Conference on Robotics and Automation (ICRA'01)*, pp. 2014–2019.

Novotny, P. (2007). Real-time processing of three dimensional ultrasound for intracardiac surgery. *PhD Thesis*, Harvard University.

Novotny, P., Stoll, J., Vasilyev, N., del Nido, P., Zickler, T., Dupont, P. and Howe, R. (2007). GPU based real-time instrument tracking with three-dimensional ultrasound. *Medical Image Analysis*, **11**: 458–464.

Ortmaier, T., Groger, M., Boehm, D., Falk, V. and Hirzinger, G. (2005). Motion estimation in beating heart surgery. *IEEE Transactions on Biomedical Engineering*, **52**(10): 1729–1740.

Parker, P. and Anderson, B. (1990). Frequency tracking of non-sinusoidal periodic signals in noise. *Signal Processing*, **20**: 127–152.

Reichenspurner, H., Detter, C., Deuse, T., Boehm, D., Treede, H. and Reichart, B. (2005). Video and robotic-assisted minimally invasive mitral valve surgery: a comparison of the port-access and transthoracic clamp techniques. *Annals of Thoracic Surgery*, **79**(2): 485–490.

Riviere, C., Rader, R. and Thakor, N. (1998). Adaptive cancelling of physiological tremor for improved precision in microsurgery. *IEEE Transactions on Biomedical Engineering*, **45**(7): 839–846.

Suematsu, Y., Martinez, J., Wolf, B., Marx, G., Stoll, J., DuPont, P., Howe, R., Triedman, J. and del Nido, P. (2005). Three-dimensional echo-guided beating heart surgery without cardiopulmonary bypass: atrial septal defect closure in a swine model. *Journal of Thoracic and Cardiovascular Surgery*, **130**: 1348–1357.

Trejos, A., Salcudean, S., Sassani, F. and Lichtenstein, S. (1999). On the feasibility of a moving support for surgery on the beating heart. *Medical Image Computing and Computer-Assisted Intervention (MICCAI'99)*, pp. 1088–1097.

Vasilyev, N., Martinez, J., Freudenthal, F., Suematsu, Y., Marx, G. and del Nido, P. (2006). Three-dimensional echo and videocardioscopy-guided atrial septal defect closure. *Annals of Thoracic Surgery*, **82**(4): 1322–1326.

Wagner, C., Perrin, D., Vasilyev, N., del Nido, P. and Howe, R. (2006). Force feedback in a three-dimensional ultrasound-guided surgical task. *Proceedings of 14th Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems*. IEEE Computer Society Press, Washington, DC, pp. 43–48.

Woo, Y., Rodriguez, E., Atluri, P., and Chitwood, W., Jr (2006). Minimally invasive, robotic, and off-pump mitral valve surgery. *Seminars in Thoracic and Cardiovascular Surgery*, **18**(2): 139–147.

Zeitlhofer, J., Asenbaum, S., Spiss, C., Wimmer, A., Mayr, N., Wolner, E. and Deecke, L. (1993). Central nervous system function after cardiopulmonary bypass. *European Heart Journal*, **14**: 885–890.