



Quasiperiodic predictive filtering for robot-assisted beating heart surgery

Citation

Yuen, Shelten G., Paul M. Novotny, and Robert D. Howe. 2008. Quasiperiodic predictive filtering for robot-assisted beating heart surgery. In Proceedings, IEEE International Conference on Robotics and Automation: May 19-23, 2008, Pasadena, CA, 3875-3880. Piscataway, NJ: IEEE Xplore.

Published Version

doi:10.1109/ROBOT.2008.4543806

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:22085975>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Quasiperiodic Predictive Filtering for Robot-Assisted Beating Heart Surgery

Shelton G. Yuen, Paul M. Novotny, and Robert D. Howe, *Member, IEEE*

Abstract—Beating heart procedures promise significant health benefits to patients but the fast motion of the heart poses a serious challenge to the surgeon. Robotic motion synchronization to heart movements could facilitate these surgeries, although for intracardiac procedures this requires the development of a predictive filter to compensate for the measurement noise and time delay present in 3D ultrasound imaging. In this paper, we present a quasiperiodic cardiac motion model and apply the extended Kalman filter to estimation of its parameters in real-time. We experimentally demonstrate high accuracy robot tracking to heart motion using this filter.

Index Terms—Ultrasound, visual servoing, cardiac surgery

I. INTRODUCTION

Real-time 3D ultrasound (3DUS) imaging has enabled new surgical procedures within the beating heart that are not possible with current endoscopic technology [1]. These procedures eliminate the need for cardiopulmonary bypass and its well documented adverse effects [2] [3]. They also make possible the real-time evaluation of reconstructive operations while the heart beats and 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 (Fig. 2), which recoils rapidly with every heartbeat.

A surgical robot could be of great assistance under these conditions. With real-time image guidance, a robot could synchronize its movements to the beating heart and the valve would appear stationary to a surgeon. Prior work achieved this for external heart procedures using fast, highly accurate measurement systems and exploiting near-periodicity in the heart motion [4] [5] [6]. Intracardiac procedures under 3DUS guidance present new challenges due to lower accuracy and time-delayed measurements. Servoing a robot with these measurements directly would result in large errors that could damage the valve and nearby delicate structures [7].

To overcome these difficulties, we have developed an extended Kalman filter (EKF) that can accurately track and predict the motion of the mitral valve in the presence of noise. The EKF can be used to feed-forward the trajectory of a cardiac target to the robot controller for synchronization. Unlike previous work in this area, we account for imperfect periodicity in heart motion by using a time-varying Fourier

This work is supported by the US National Institutes of Health under grant NIH R01 HL073647-01.

The authors are with the School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, 02138 USA. R.D. Howe is also with the Harvard-MIT Division of Health Sciences & Technology, Cambridge, MA, 02139 USA. email: howe@seas.harvard.edu

series model in the filter. We show that the EKF permits accurate robot synchronization to a mitral valve target and also adds robustness to cases where there is high variability in heart rate. In the following, we first discuss mitral valve annulus motion. We then describe the EKF and several other predictive filtering methods and compare them in simulation. A subsequent 3DUS servoing experiment in a water tank allows us to evaluate the overall performance of the robot system to surgical tasks like those encountered in mitral valve annuloplasty.

II. MITRAL VALVE ANNULUS MOTION

We first highlight several features of mitral valve motion. Previous characterizations of mitral valve motion indicate that it is strongly constrained to translational motion along a single axis [8] with negligible rotational movements [9]. Fig. 1A is representative of several cycles of motion along this axis at 1 Hz, digitized from 3DUS sequences of human mitral valves [8]. Spectral decomposition of this trajectory (Fig. 1B) indicates a dominant motion component at 1 Hz corresponding to the heart rate with several significant harmonics appearing above the noise floor. This suggests that the mitral valve trajectory may be suitably approximated by a limited number of harmonics. Such an approach has been attempted in the past for modeling the motion of the external heart wall [4] [5].

III. PREDICTIVE FILTERS

A. Time-Varying Fourier Series Model with Extended Kalman Filter

A periodic motion model is obtained by an m -order Fourier series with a DC offset

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

where $y(t)$ is the position in ultrasound coordinates, ω is the heart rate, c is the DC offset, and r_i and ϕ_i are respectively the harmonic amplitudes and phases. Accurate modeling of




Fig. 1. Mitral Annulus Motion. A: Trajectory; B: Power Spectral Density.




Fig. 2. Mitral Valve Anatomy. The valve consists of two leaflets surrounded by the mitral valve annulus. When functioning properly, it regulates blood flow from the left atrium to the left ventricle. (Image source: Patrick J. Lynch, medical illustrator; C. Carl Jaffe, MD, cardiologist)

quasiperiodicity requires a more flexible model in which the heart rate and signal morphology can evolve over time. Using the parameterization of [10], the trajectory can be expressed as the following time-varying Fourier series

$$y(t) = c(t) + \sum_{i=1}^m r_i(t) \sin \theta_i(t), \quad (2)$$

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 (1).

Defining the state vector $\mathbf{x}(t) \triangleq [c(t), r_i(t), \omega(t), \theta_i(t)]^T$, $i \in (1, \dots, m)$ and assuming that the first $m+2$ state variables and $\phi_i(t)$ evolve through random walk, the state space model for this system is

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

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

$h(\mathbf{x}(t)) \triangleq y(t)$ from (2), and $\nu(t) \sim \mathcal{N}(0, \sigma_R^2)$ and $\boldsymbol{\mu}(t) \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$ are independent Gaussian noise terms.

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

$$\begin{aligned} \mathbf{P}(t + \Delta t|t) &= \mathbf{F}\mathbf{P}(t|t)\mathbf{F}^T + \mathbf{Q} \\ \mathbf{S} &= \sigma_R^2 + \mathbf{H}\mathbf{P}(t + \Delta t|t)\mathbf{H}^T \\ \mathbf{K} &= \mathbf{P}(t + \Delta t|t)\mathbf{H}^T\mathbf{S}^{-1} \end{aligned}$$

$$\begin{aligned} \hat{\mathbf{x}}(t + \Delta t|t + \Delta t) &= \mathbf{F}\hat{\mathbf{x}}(t|t) + \mathbf{K}(z(t + \Delta t) - h(\mathbf{F}\hat{\mathbf{x}}(t|t))) \\ \mathbf{P}(t + \Delta t|t + \Delta t) &= (\mathbf{I} - \mathbf{K}\mathbf{H})\mathbf{P}(t + \Delta t|t), \end{aligned}$$

where

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

and $\mathbf{P}(\cdot) \triangleq E[\hat{\mathbf{x}}(\cdot)\hat{\mathbf{x}}(\cdot)^T]$ denotes the state estimate covariance, whose initialization is described later in this section. Note that the time dependencies for \mathbf{F} , \mathbf{K} , \mathbf{S} , and \mathbf{H} have been dropped for notational convenience. The EKF as presented here is a slight variant on the one first introduced by Parker and Anderson [10].

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 Fig. 1B indicates that the dominant frequency peak is the fundamental frequency of the signal (i.e., heart rate). We apply an 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 (1) is equivalently reparameterized as

$$y(t) = c_0 + \sum_{i=1}^m [a_i \sin(i\hat{\omega}_0 t) + b_i \cos(i\hat{\omega}_0 t)]. \quad (3)$$

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

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

to which the least squares estimate is obtained as $\hat{\tilde{\mathbf{x}}}_0 = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \tilde{\mathbf{z}}$. $\hat{\mathbf{x}}(T|T)$ 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 $\mathbf{P}(T|T) = \text{diag}[\sigma_R^2/N, \sigma_1^2, \sigma_1^2/2^2, \dots, \sigma_1^2/m^2, \sigma_\omega^2, 0.02, \dots, 0.02]$. The relationship between the amplitude uncertainties is chosen to reflect the decreasing harmonic strength seen in

Fig. 1B, while the phase uncertainties follow those used in [10]. Parameters σ_1^2 and σ_ω^2 are determined through experimentation. Last we assume that the process noise covariance \mathbf{Q} is diagonal with all values set to 10^{-4} except for q_ω , the entry corresponding to $\omega(t)$.

B. Autoregressive Model with Least Squares Estimator

Another model that can be employed for mitral valve trajectory prediction is an n -order autoregressive (AR) model

$$y[k] = \sum_{i=1}^n \alpha_i y[k-i], \quad (4)$$

where $\alpha_i, i \in \{1, \dots, n\}$ are the model coefficients and $y[k]$ is the target position in ultrasound coordinates at time sample k . Rather than explicitly assuming periodicity in the target motion, this model predicates that the k^{th} 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 $z[k] = y[k] + \nu[k]$ to be the noise-corrupted 3DUS measurement at time sample k with $\nu[k] \sim \mathcal{N}(0, \sigma_R^2)$, this estimator is expressed compactly in matrix form as

$$\begin{aligned} \mathbf{Z}[k] &= [z[k-n-1], \dots, z[k-1]] \\ \mathbf{R}[k] &= \mathbf{R}[k-1] + \mathbf{Z}[k]^T \mathbf{Z}[k] \end{aligned} \quad (5)$$

$$\hat{\alpha}[k] = \hat{\alpha}[k-1] + \mathbf{R}[k]^{-1} \mathbf{Z}[k]^T (z[k] - \mathbf{Z}[k] \hat{\alpha}[k-1]), \quad (6)$$

with initial conditions $\mathbf{R}[0] = \mathbf{0}$ and $\hat{\alpha}[0] = 0$. The autoregressive model order was determined using the Akaike Information Criteria on data collected and processed off-line, yielding $n = 30$. Predicted target locations, $\hat{y}[k]$, are obtained through evaluation of (4). Additionally, the target trajectory was interpolated from its inherent measurement rate to the higher control rate of the robot using the Whittaker-Shannon interpolation formula.

C. Autoregressive 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 (6) does not update $\hat{\alpha}$ because $\mathbf{R}[k]^{-1} \mathbf{Z}^T \rightarrow \mathbf{0}$ as $k \rightarrow \infty$. Exponential weighting of previous measurements in the iterative least squares estimator is achieved through simple modification of (5):

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

where $0 < f \leq 1$ is the so-called *fading* factor. Choosing $f = 1$ recovers the estimator of Section III-B 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 paper. As can be imagined, reducing the contribution of previous measurements ($f < 1$) can be desirable if the trajectory evolves through time; though doing so incurs increased estimate error when the trajectory is not time varying.

D. 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 to a simpler method of using the previous cardiac cycle trajectory for prediction of the next. A variation of this method, termed *Last Cycle* here, was used successfully in a beating heart tracking system [6].

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 Fig. 1A was downsampled 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 RMS errors for each predictor averaged across 100 monte-carlo trials are shown in Fig. 3A. The EKF, 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 value lower would give performance below the Last Cycle method. The EKF was run with $m = 8$ harmonics, $N = 280$ initialization points (10 sec), $\sigma_1^2 = 2$, $\sigma_\omega^2 = 0.11$ (roughly twice the frequency resolution of the FFT), and $q_\omega = 10^{-3}$.

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 sec of heart motion at 60 bpm and 10 sec of motion at $(60 + \Delta\text{HR})$ bpm. The second portion of the trajectory was generated by compression/dilation of the trajectory in Fig. 1A to obtain the desired heart rate. Like before, the composite trajectory was downsampled to 28 Hz and corrupted with $\sigma_R = 1.3$ mm noise (the measured uncertainty of the 3DUS annulus estimates [7]) and the last 10 sec were used to evaluate performance. A reasonable range of $-10 \text{ bpm} \leq \Delta\text{HR} \leq 10 \text{ bpm}$ was determined from clinical data (Fig. 4), which is discussed later.

Fig. 3B shows the mean RMS errors for each predictor across 100 monte-carlo trials. The EKF provided better predictions than the other three methods. It was also the only method that yielded sub- σ_R error for the majority of heart rate changes. As expected, the accuracy of the AR filter 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 lessened accuracy when the trajectory did not change significantly. Finally, the Last Cycle method




Fig. 3. RMS prediction error results for parametric simulation studies. A: Error for varying measurement noise; B: Error for step heart rate changes.




Fig. 4. Summary of Heart Rate Statistics on Five Human Subjects

showed performance comparable to the Fading AR filter. For this simulation all filter parameters were chosen the same as in the previous simulation, with the exception of $q_\omega = 5 \times 10^{-3} \text{ (rad/s)}^2$ and $\sigma_\omega^2 = 1 \text{ (rad/s)}^2$ 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 [11] and composite mitral valve trajectories were generated in a manner similar to the previous simulation study. Noise-corrupted measurements were generated as before, with $\sigma_R = 1.3 \text{ mm}$. Summary statistics for each subject are presented in Fig. 4 and an example of the beat-to-beat heart rate for subject number 1 is shown in Fig. 7d.

Results from this study suggest that the EKF is more suited to tracking and prediction in this application than the other three methods (Table I). 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.

TABLE I
SIMULATED PREDICTION ERROR ON CLINICAL DATA

Subject No.	RMS (Max) Prediction Error [mm]			
	EKF	AR	Fading AR	Last Cycle
1	0.83 (4.92)	1.48 (5.13)	1.69 (6.45)	1.99 (9.95)
2	0.88 (5.08)	1.66 (6.19)	1.63 (6.65)	1.87 (9.00)
3	0.88 (5.44)	1.98 (8.27)	2.03 (8.27)	2.24 (11.77)
4	1.33 (6.36)	1.81 (6.70)	2.07 (6.85)	2.72 (13.50)
5	0.88 (5.53)	1.58 (6.17)	1.88 (7.54)	2.13 (9.96)
Mean	0.99 (5.60)	1.76 (6.83)	1.90 (7.33)	2.24 (11.06)




Fig. 5. (A) Tank experiments setup. The ultrasound probe images an instrument and target submerged in a water tank. The instrument is attached to a robot and passes through a port. During the trials, the target (B) moves left to right, simulating the 1D motion of the mitral valve.

IV. EXPERIMENTAL VALIDATION

A. Experimental Setup

The usefulness of the EKF, AR, and Fading AR filters to a 3DUS servoing task was tested in a water tank. Details of the setup can be found in [7]. Briefly, a surgical instrument was attached to the end of a small robot (Phantom 1.5, Sensable Technologies, Woburn, MA), which was commanded to maintain a constant offset from the instrument tip to the target position. A sliding ball joint approximated an instrument passing through a port in the heart wall. The target and instrument were both imaged by a real-time 3DUS probe in a water tank at 28 Hz (Fig. 5).

Data was streamed from the ultrasound machine (SONOS 7500, Philips Medical, Andover, MA) to a PC over an ethernet connection. The data stream was captured on the target PC and passed to a graphics processing unit (8800GTS, nVidia Corp, Santa Clara, CA) where it was automatically segmented using the modified Radon transform segmentation algorithms from [12]. Target position measurements were passed to a third thread for filtering and returned predicted target positions 132 ms in the future (68 ms from imaging and segmentation, 64 ms from robot lag time). Simultaneously, the instrument and target measurements were used to automatically register the robot and imaging coordinate frames [7]. These were then sent to the 1 kHz PD control loop for the robot ($K_P = 0.16 \frac{\text{N}}{\text{mm}}$, $K_D = 0.002 \frac{\text{N}}{\text{mm/s}}$).

The system was tested against both fixed and variable heart rates to assess the performance of the EKF, AR filter, and Fading AR filter. For fixed-rate experiments, the target speed corresponded to a heart rate of 60 bpm. In variable heart rate experiments, the target motion was driven by the clinical RR

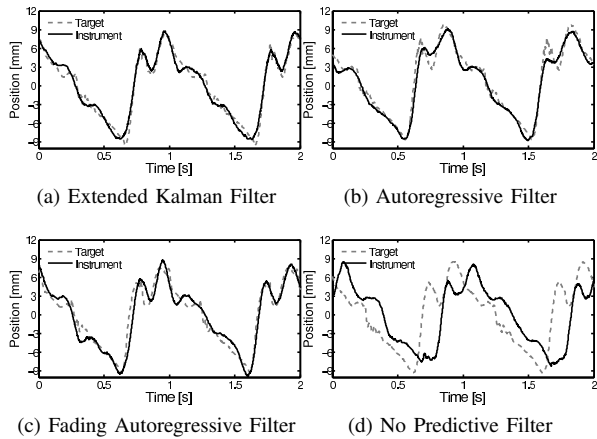


Fig. 6. Results of robot tracking at 60 bpm. Only one dimension of the 3D positions is shown because the movement is 1D.

TABLE II
ROBOT TRACKING ERROR ON A FIXED-RATE TARGET (60 BPM)

Filter Type	RMS Error [mm]	Max Error [mm]
EKF	0.97	3.26
AR	1.15	3.91
Fading AR	1.62	7.60
None	4.57	12.36

interval data from the five subjects described in Section III-D. Each filter used the same parameters as those in the simulation studies of Section III-D. The target was attached to a cam that mimicked mitral valve movements [8] and ground truth was measured with an angular potentiometer. The robot tip position was recorded using the joint encoders of the robot. Both of these measurements were made at 1 kHz and stored to a data file for analysis.

B. Experimental Results

Fixed-rate robot tracking with feed-forward target prediction provided RMS errors under 2 mm, as shown in Fig. 6 and Table II. In agreement with our simulations, the EKF yielded improved results over both the AR and Fading AR filters, with each having RMS errors of 0.97, 1.15, and 1.62 mm, respectively. Likewise, as expected for a fixed-rate target, the AR filter outperformed the Fading AR filter. To illustrate the importance of compensating system delays with a predictive filter, robot tracking performance using raw measurements are also given. Qualitative observation of Fig. 6d indicates that delays in the system are the dominant source of error. Left uncompensated, they cause 4.57 mm of RMS error (Table II).

The EKF significantly outperformed the other filters against a target with heart rate variability (Table III). Indeed, it was the only filter to consistently yield RMS errors under 2 mm. While the introduction of variable heart cycles decreased performance appreciably for all filters, the EKF is noticeably more robust to this error source than the AR and Fading AR filters. This is apparent in Fig. 7, a representative example of robot synchronization for subject number 1.

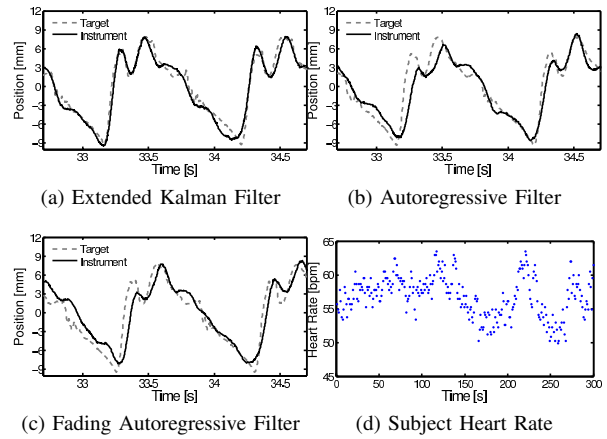


Fig. 7. (a-c) Robot tracking example on a variable heart rate target derived from subject 1. Only one dimension of the 3D positions is shown because the movement is 1D. (d) Heart rate data for subject 1.

TABLE III
ROBOT TRACKING ERROR ON VARIABLE-RATE CLINICAL DATA

Subject No.	RMS (Max) Tracking Error [mm]		
	EKF	AR	Fading AR
1	1.30 (6.98)	2.18 (8.12)	2.21 (8.58)
2	1.50 (7.30)	1.92 (7.47)	1.79 (6.70)
3	1.51 (5.53)	2.22 (9.52)	2.27 (9.68)
4	1.29 (4.56)	1.68 (6.08)	1.87 (8.78)
5	1.53 (6.80)	2.06 (9.05)	2.26 (8.78)
Mean	1.43 (6.23)	2.01 (8.05)	2.08 (8.50)

Averaged across all five patients, the EKF provided $\approx 30\%$ less error than the other two filters.

V. DISCUSSION

Robot-assisted beating intracardiac surgery guided by 3D ultrasound faces two major obstacles: time delay and low imaging accuracy. In this paper, we employed an extended Kalman filter (EKF) that provided high accuracy predictions of heart motion in the presence of both. The EKF model was selected to directly treat quasiperiodicity and this enabled superior tracking performance compared to existing methods with clinically-observed heart rate patterns. Using the EKF in experimental trials, we demonstrated that ultrasound-based visual servoing can be used to accurately synchronize a surgical instrument to a rapidly moving target mimicking the motion of the mitral valve.

Previous work on coronary artery bypass graft (CABG) procedures obtained similar or higher accuracy results than here for 3D synchronization to the external heart wall. Using a high speed camera operating at 500 Hz, Ginhoux *et al.* applied an adaptive filter bank in a generalized predictive controller to obtain robot tracking errors of ≈ 1.5 mm in a porcine heart [5]. In principle, their filtering approach is equivalent to the AR and Fading AR filters described in this work and has their attendant assumption of a fixed heart rate. In similar but independent work, Bebek and Cavusoglu demonstrated <1 mm robot tracking error by incorporating ECG [6] into a variation of the Last Cycle method discussed here. Their system used sonomicrometry sensors sutured to

the surface of a porcine heart, sampling at 257 Hz with $\approx 250\mu\text{m}$ measurement accuracy. These researchers estimate that beating heart CABG requires robot tracking accuracies of ≈ 2 mm in order to safely manipulate blood vessels.

Mitral valve repair shares many of the challenges faced in CABG; however, it assumes additional challenges by being a procedure inside the heart. A major consideration is the need for a real-time imaging technology that can image through blood. We adopt 3D ultrasound because it is currently the only technology that has these capabilities; endoscopic systems are obscured by blood, CT and MRI are not yet able to image at the speeds necessary for beating heart surgery, and attached sensors (i.e., sonomicrometry crystals) are unlikely to be placed reliably inside a moving heart. 3D ultrasound also overcomes the difficulties in spatial perception experienced with traditional 2D ultrasound [13]. For all of its advantages, 3D ultrasound has poor accuracy (≈ 1 mm) and a relatively low sampling frequency (28 Hz). The acquisition and processing of 3D ultrasound volumes also introduces a time delay of 68 ms into the system, during which time the annulus can recoil 15 mm [7].

The EKF developed in this paper enabled us to overcome the apparent shortcomings of 3D ultrasound imaging. The tank trials presented here demonstrated RMS synchronization errors of ≈ 1.5 mm for trajectories derived from clinical heart rate variability data and 0.97 mm for a fixed-rate trajectory. However, additional errors are expected for *in vivo* conditions due to ultrasound probe motion and lowered imaging resolution in blood and through tissue. Future surgical experiments will be conducted open-chest on subjects under artificial ventilation. We expect that the out-of-axis motion due to respiration can be compensated by mounting the robot to a translational stage that is controlled to follow the forced-respiration of the subject. We anticipate that accuracies of ≈ 2 mm will be needed for mitral valve repair, although the requirements should be eased from those in CABG because small, delicate structures are not handled. Future work must determine reasonable tracking accuracies for such procedures. Likewise, the influence of respiratory motion on the EKF and the verisimilitude of our cam-driven heart motion emulator must be addressed.

Should better tracking performance be needed, there are several potential routes for improvement. First, heart rate variability could be reduced through drug treatment or electrical pacing of the heart. Second, following [6], ECG information could be used to obtain direct measurements of heart rate. This could reduce the effect of the nonlinearity in $h(\mathbf{x}(t))$ since, as mentioned before, perfect knowledge of $\omega(t)$ turns this into a linear estimation problem (3). Note that the EKF, a nonlinear filter, would still be useful with these measurements because they are sampled at discrete intervals but $\omega(t)$ is continuous. Finally, robot tracking error may be reduced through the use of a model predictive controller, as done in [5] [6]. A comparison of simulated and experimental results in this paper suggest that robot control accounts for ≈ 0.5 mm of error.

In closing, we note that while the EKF presented here

was chosen for 3D ultrasound-guided mitral valve repair, it is not limited to that application. The EKF proved to be a superior predictor over those presuming a fixed rate (AR and Fading AR filters) or self-similarity between consecutive cycles (Last Cycle method) for every trial in this study. Averaged across the five clinically-derived data sets, the EKF provided $\approx 30\%$ less error than other filters. We reasonably expect that the EKF could be applied to other applications of motion compensation. For instance, the EKF could be applied to external heart procedures where, in addition to increasing robustness to heart rate fluctuations and noise, it could mitigate errors caused by unexpected heart beats. Such events can occur when instruments contact the heart.

ACKNOWLEDGEMENTS

The authors would like to acknowledge Dan Kettler for his contributions to the initial experimental setup and software for this work. Additional thanks to Mahdi Tavakoli, Peter Hammer, and Petr Jordan for many insightful conversations.

REFERENCES

- [1] Y. Suematsu, J. Martinez, B. Wolf, G. Marx, J. Stoll, P. DuPont, R. Howe, J. Friedman, and P. del Nido, "Three-dimensional echo-guided beating heart surgery without cardiopulmonary bypass: atrial septal defect closure in a swine model," *J. Thorac. Cardiovasc. Surg.*, vol. 130, pp. 1348–1357, 2005.
- [2] J. Zeitlhofer, S. Asenbaum, C. Spiss, A. Wimmer, N. Mayr, E. Wolner, and L. Deecke, "Central nervous system function after cardiopulmonary bypass," *European Heart Journal*, vol. 14, pp. 885–890, 1993.
- [3] D. Bellinger, D. Wypij, K. Kuban, L. Rappaport, P. Hickey, G. Wer-novsky, R. Jonas, and J. Newburger, "Developmental and neurological status of children at 4 years of age after heart surgery with hypothermic circulatory arrest or low-flow cardiopulmonary bypass," *Circulation*, vol. 100, pp. 526–532, 1999.
- [4] Y. Nakamura, K. Kishi, and H. Kawakami, "Heartbeat synchronization for robotic cardiac surgery," in *Proc. IEEE International Conference on Robotics and Automation (ICRA'01)*, May 2001, pp. 2014–2019.
- [5] R. Ginhoux, J. Gangloff, M. de Mathelin, L. Soler, M. A. Sanchez, and J. Marescaux, "Active filtering of physiological motion in robotized surgery using predictive control," *IEEE Transactions on Robotics*, vol. 21, no. 1, pp. 27–79, 2006.
- [6] O. Bebek and M. Cavusoglu, "Intelligent control algorithms for robotic assisted beating heart surgery," *IEEE Transactions on Robotics*, vol. 23, no. 3, pp. 468–480, June 2007.
- [7] P. Novotny, "Real-time processing of three dimensional ultrasound for intracardiac surgery," Ph.D. dissertation, Harvard University, 2007.
- [8] D. Kettler, R. Plowes, P. Novotny, N. Vasilyev, P. del Nido, and R. Howe, "An active motion compensation instrument for beating heart mitral valve surgery," in *Proc. IEEE International Conference on Intelligent Robots and Systems*, San Diego USA, Oct. 2007.
- [9] J. Gorman III, K. Gupta, J. Streicher, R. Gorman, B. Jackson, M. Ratcliffe, D. Bogen, and L. E. Jr., "Dynamic three-dimensional imaging of the mitral valve and left ventricle by rapid sonomicrometry array localization," *J. Thorac. Cardiovasc. Surg.*, vol. 112, no. 3, pp. 712–726, 1996.
- [10] P. Parker and B. Anderson, "Frequency tracking of nonsinusoidal periodic signals in noise," *Signal Process.*, vol. 20, pp. 127–152, 1990.
- [11] A. Goldberger, L. Amaral, L. Glass, J. Hausdorff, P. Ivanov, R. Mark, J. Mietus, G. Moody, C.-K. Peng, and H. Stanley, "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," *Circ.*, vol. 101, no. 23, 2000.
- [12] P. Novotny, J. Stoll, N. Vasilyev, P. del Nido, T. Zickler, P. Dupont, and R. Howe, "GPU based real-time instrument tracking with three-dimensional ultrasound," *Medical Image Analysis*, vol. 11, pp. 458–464, 2007.
- [13] J. Cannon, J. Stoll, I. Salgo, H. Knowles, R. Howe, P. Dupont, G. Marx, and P. del Nido, "Real-time three-dimensional ultrasound for guiding surgical tasks," *Computer Aided Surgery*, vol. 8, pp. 82–90, 2003.