Distribution Agreement

In presenting this thesis as a partial fulfillment of the requirements for a degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis in whole or in part in all forms of media, now or hereafter now, including display on the World Wide Web. I understand that I may select some access restrictions as part of the online submission of this thesis. I retain all ownership rights to the copyright of the thesis. I also retain the right to use in future works (such as articles or books) all or part of this thesis.

Alec Reinhardt April 13, 2021

Recursive Phase Estimation with Applications to Heartbeat Processing

by

Alec Reinhardt

Gordon Ramsay
Advisor

Department of Mathematics

Gordon Ramsay
Advisor

Bree Ettinger
Committee Member

Alessandro Veneziani
Committee Member

2021

Recursive Phase Estimation with Applications to Heartbeat Processing

by

Alec Reinhardt

Gordon Ramsay
Advisor

An abstract of
a thesis submitted to the Faculty of Emory College of Arts and Sciences
of Emory University in partial fulfillment
of the requirements of the degree of
Bachelor of Science with Honors

Department of Mathematics

2021

Abstract

Recursive Phase Estimation with Applications to Heartbeat Processing

by Alec Reinhardt

In this thesis, we explore the estimation of nonlinear, time-varying phase. We first introduce how to conceptualize and estimate phase using global techniques from Fourier analysis and Functional Data Analysis, before going on to develop recursive formulations using a constrained Extended Kalman Filter approach. We then evaluate the performance of the global and recursive methods on a variety of simulated examples, and discuss respective limitations of each. In general, we find that the recursive methods can outperform traditional estimation techniques in high-noise settings, and may be better suited to tracking local phase variations. Finally, we illustrate how the recursive models may be applied in order to robustly process and analyze heartbeat signals.

Recursive Phase Estimation with Applications to Heartbeat Processing

by

Alec Reinhardt

Gordon Ramsay
Advisor

A thesis submitted to the Faculty of Emory College of Arts and Sciences
of Emory University in partial fulfillment
of the requirements of the degree of
Bachelor of Science with Honors

Department of Mathematics

2021

Contents

List of Figures

List of Tables

1 Introduction

1.1 Background

The automatic processing of heartbeat recordings is a complex and well-studied area within biomedical engineering and statistics. Depending on what kind of information is extracted, processing methods are instrumental for development of different kinds of classification schemes, including the detection of cardiac arrhythmias, the biometric recognition of individuals based on their heartbeat, and the monitoring of different behavioral states [3] [7] [12]. This thesis comes out of a long-term study which aims to use Doppler ultrasound fetal heartbeat signals as a real-time behavioral classifier, with the ultimate goal of assessing risk of autism in utero using these processing pipelines.

In order to go about tackling this problem, there is a more general question which needs to be addressed. That is, how do we come up with a robust and adaptive framework for characterizing heartbeat signals (and other waveforms) in real time? We propose that the notion of *phase* (defined below) can give us a useful, clinically relevant representation of heartbeat signals, since it provides information about naturally-occurring fluctuations in heart rate which have previously been studied as indicators of physiological and behavioral status [6] [12]. In particular, we wish to develop a flexible method to track nonlinear *phase distortions* across intra-beat and inter-beat time intervals.

1.2 Defining Phase

In plain terms, *phase* can be thought of as a measure of how far we are along a waveform with respect to a starting time [14]. More generally, we may define the phase as a smooth, monotonically-increasing function of time, $\phi(t)$, representing the state of an underlying dynamical system, which can uniquely represent all versions of a waveform with a given *amplitude* function, $A(t)$. Furthermore, we note that the derivative of phase, $\phi'(t) = \omega(t)$ gives the *instantaneous frequency* of the waveform, another useful measure of wave dynamics.

Here, it is important to mention that phase and amplitude are mutually-dependent concepts and are only clearly defined when we consider them as parameters of a known *reference function*

or a set of reference functions known as a *basis*. For instance, we may consider the simple family of sine waves given by

$$Asin(\phi(t)) = Asin(\omega t + \phi_0) \tag{1}$$

which, for a given frequency ω, are completely described by some constant amplitude A and phase offset ϕ_0. This representation gives rise to Fourier analysis, in which we decompose signals based on the amplitudes and phases of multiple sinusoidal basis functions, each with a given frequency (see Section 2.1).

Alternatively, we may wish to represent sinusoidal phase and amplitude as nonlinear functions corresponding to one overall waveform (rather than a decomposition of multiple sinusoidal bases). This can be expressed in the following complex exponential form

$$A(t)e^{j\phi(t)} = A(t)[cos(\phi(t)) + jsin(\phi(t))] \tag{2}$$

However, for each of these instances, we claim that the choice of representing phase in terms of sine waves can be considered an arbitrary one. Thus, for cases where the underlying waveform is more complex or specific in structure, we would like to define and estimate phase in a way that is more relevant to particular applications (e.g. processing heartbeat signals). By extension, we may also wish to characterize relative temporal alignment, or *phase variability*, for functions which in general are not periodic in nature.

1.3 Motivation

Various approaches exist for globally estimating the phase of a signal in both the sinusoidal and non-sinusoidal case, of which a select few are discussed in the following section [9]. However, there has been less attention given to recursive approaches, which may offer benefits for tracking the local behavior of phase and for use in real-time processing systems [21].

Therefore, in this thesis, we aim to develop techniques for recursively estimating the phase of an arbitrary function embedded in noise using the Extended Kalman Filter [19]. We evaluate the performance of these methods on simulated examples and show correspondence to global tech-

niques for phase and phase variability estimation. Applications related to analysis of heartbeat signals are considered.

2 Theory

Before introducing our recursive framework, we discuss global methods for phase estimation, both in the sinusoidal and non-sinusoidal case.

2.1 Sinusoidal Phase Estimation

Classically, the phase (or phase components) of a sinusoidal waveform embedded in additive Gaussian noise is estimated using methods from Fourier analysis. We describe here three of the most standard techniques, ordered in terms of how local (i.e. nonlinear) their phase estimates are.

2.1.1 Discrete Fourier Transform

For a discrete-valued signal x_n of length N, its Discrete Fourier Transform (DFT) is defined as [8]

$$\mathcal{F}\{x_n\} = \sum_{n=0}^{N-1} x_n e^{-i2\pi kn/N} = X_k \tag{3}$$

and gives the *frequency-domain* representation of the signal. Namely, the DFT allows for the estimation of constant amplitude and phase parameters of sinusoidal components of the form (1) across different frequencies. The phase estimate for the k-th sinusoidal component is given as the angle of the complex-valued spectral representation, X_k.

Note that the DFT is invertible, meaning these spectral representations may be transformed back into the time-domain representation (i.e. the signal x_n). This inverse is given by

$$\mathcal{F}^{-1}\{X_k\} = \frac{1}{N}\sum_{k=0}^{N-1} X_k e^{i2\pi kn/N} = x_n \tag{4}$$

Because the phase and amplitude estimates obtained from the DFT are constants, this method is only suitable for estimation when we have stationary periodic signals – i.e. the waveform retains

its properties over time.

2.1.2 Short-Time Fourier Transform

As a more suitable approach for phase and amplitude estimation of non-stationary signals, we may use the Short-Time Fourier Transform (STFT). This method is based on taking the DFT over overlapping, windowed segments of the signal, and is given for the discrete time signal x_n by [13]

$$X(m,\omega) = \sum_{n=-\infty}^{\infty} x(n)w(n-m)e^{-j\omega n} \tag{5}$$

where m is a time index, ω is frequency, and w is a chosen *window function*.

This transformation gives a *time-frequency representation* of the signal which allows for time-varying amplitude, frequency, and phase estimation across multiple major components (e.g. harmonics). However, the STFT comes with its own key limitation, namely a tradeoff in temporal resolution and frequency resolution due to the Heisenberg Uncertain Principle [1]. Practically, this means it is not possible to get the true instantaneous phase estimate for arbitrarily local phase distortions.

2.1.3 Hilbert Transform

Due to the limitation of the DFT and STFT for estimating local phase changes, we consider the Hilbert Transform as the primary Fourier-based method for obtaining estimates of the phase and amplitude functions of a sinusoidal waveform of the form (2).

In discrete time, the Hilbert Transform is conveniently expressed via the DFT and inverse DFT [8]. Namely, we have

$$\mathcal{H}(x_n) = \mathcal{F}_D^{-1}\{-jsgn(N/2-k)sgn(k)\mathcal{F}_D\{x_n\}\} \tag{6}$$

From this transform, we can define the complex-valued *analytic signal* z_n as

[1] This tradeoff always exists for the discrete-time STFT, but it can be properly balanced by choosing different window functions

$$z_n = x_n + j\mathcal{H}(x_n) = A_n e^{j\phi_n} \tag{7}$$

and obtain the wrapped instantaneous phase ϕ_n^{wr} using

$$\phi_n^{wr} = arg(z_n) = tan^{-1}(\mathcal{H}(x_n)/x_n) \tag{8}$$

Note that the Hilbert phase estimate is "wrapped" in the sense that the inverse tangent function is only defined modulo pi, i.e. $\phi_n^{wr} \in (-\pi, \pi]$. To obtain a more useful representation, we need to transform, or unwrap, ϕ_n^{wr} such that its tangent is preserved but it is monotonically increasing for increasing n. This is a well-studied problem known as Phase Unwrapping Problem which can be nontrivial when the signal-to-noise ratio (SNR) is high [15].

In addition to the problem of Phase Unwrapping, the Hilbert Transform has the key limitation that it is not easily interpretable for signals which are non-sinusoidal or contain multiple frequency components. An illustration of this is shown in section 4.1.

2.2 General Phase Estimation

Given that Fourier techniques may not always be appropriate for analyzing arbitrary waveforms or temporal deviations from a non-periodic curve, we require a more general technique for phase estimation. In particular, we are interested in modeling an observed signal $x(t)$ as a phase and amplitude-distorted instance of some underlying (non-sinusoidal) *reference function* $\tilde{y}(t)$, i.e.

$$x(t) \sim A(t)\tilde{y}(\phi(t)) \tag{9}$$

In order to address this problem more concretely, we introduce some theory from Functional Data Analysis (FDA). In particular, we will frame this phase estimation problem in relation to the constrained optimization problem known as Curve Registration, which was posed in terms of FDA by Ramsay and Silverman [17] [18]. Note that the optimization technique described below can be viewed as a generalization to other kinds of curve-alignment methods, including landmark-based registration and Dynamic Time Warping [4].

2.2.1 Curve Registration

Suppose we are given a set of time-varying curves, and we wish to characterize how these curves differ from one another, both in terms of *amplitude variation* (i.e. differences along the y-axis) and *phase variation* (i.e. differences along the t-axis).

In general, Curve Registration seeks to find the so-called *warping functions*, $h(t)$, which map a curve's *observed time* onto some common *reference time* in such a way that minimizes phase variation between all the warped curves. For our purposes, we will assume that amplitude variation between phase-aligned curves is linear – that is, amplitude deviations can be described with a scaling and shifting factor. In this case, we can relate each observed curve $x_i(t)$ to a common template curve $\tilde{y}(t)$ with the following model

$$\tilde{y}(t) = a_i x_i(h_i(t)) + b_i + \epsilon_i \tag{10}$$

where $h_i(t)$ is the warping function of the i-th curve, a_i is the scale factor, b_i is the shift factor, and ϵ_i is a noise term.

The constraint for the Curve Registration problem comes from the form of the warping functions. In particular, as is the case with our notion of phase, we expect that these functions are continuous, invertible, and monotonically increasing, since the flow of time is unidirectional and we wish to recover any given curve from its temporally-transformed counterpart. In fact, we can think of these warping functions, or more specifically, their inverses, as a kind of general phase, since they summarize temporal location along the reference curve $y(t)$.

In [17], Ramsay and Li proposed using a class of functions given by

$$h(t) = c_0 + c_1 \int_0^t e^{\int_0^t w(s)ds} d\tau \tag{11}$$

to represent monotonically-increasing warping functions, where $w(t)$ is some unconstrained function representing the *relative curvature* of $h(t)$.

From these expressions, the overall FDA optimization problem can be expressed as a global minimization of the squared error between the reference curve $y(t)$, and each warped curve, after

scaling and shifting has been applied. In addition, we add a *regularization factor* λ to avoid estimates which have an overly large relatively curvature. In the case where all scale and shift factors a_i and b_i are equal, we may perform this optimization using the target function

$$F_\lambda(\tilde{y}, x_i|h_i) = \int ||\tilde{y}(t) - x_i(h_i(t))||^2 dt + \lambda \int w_i^2(t) dt \tag{12}$$

from which the minimum is obtained by changing parameters of $w_i(t)$. A more robust criterion when the values of a_i and b_i may vary for each curve, is based on maximizing the *correlation* between the observed curve and unwarped template, and can be expressed as a minimum eigenvalue criterion as in [18] (p. 139-140).

Limitations

While Curve Registration gives a solid framework for phase variability estimation, it may not always be suitable for practical purposes since it is a global optimization scheme which can be computational expensive. Furthermore, the globally formulated constraint of monotonicity on the warping function $h(t)$ may be unwieldy or too restrictive in some cases. Thus, we would like to shift from this global formulation to a recursive one, where we can impose local constraints in order to get monotonically increasing phase or inverse warping function estimates. Additionally, recursive approaches are advantageous in that they can incorporate information about state dynamics and thus may ultimately result in a more intuitive formulation of FDA.

2.3 Recursive Estimation

In order to develop this recursive framework, we first introduce the general approach for solving recursive estimation problems (Bayesian filtering), followed by a specific implementation of this approach known as the Extended Kalman Filter (EKF) [1] [19].

2.3.1 Bayesian Filter

In general, the problem of recursive estimation of an underlying state (e.g. warping function parameters) given a set of observations can be solved using a Bayesian filter [1]. Formally, the Bayesian filter provides a way for estimating the posterior distribution $p(x_k|Z_{1:k})$ for a random

variable x_k representing some underlying dynamical state and $Z_{1:k}$ representing realizations (i.e. observations) of a random variable z_k up to the current time point.

In order for this filtering problem to be tractable, we make the following assumptions about random variables x_k and z_k:

- The PDFs $p(x_k)$, $p(x_{k+1}|x_k)$, and $p(z_k|x_k)$ are defined for all k, $1 \leq k \leq n$, and $p(x_0)$ is known
- x_k is a discrete-time Markov process,
  i.e. $p(x_k|x_0, x_1, \ldots, x_{k-1}) = p(x_k|x_{k-1})$
- x_k can be expressed as a function of the previous state x_{k-1}, along with current and previous time point and a noise with a known PDF variable through a given *state model*
- z_k can be expressed as a function of x_k and a noise variable with a known PDF through a given *measurement model*
- The noise variables from the state and measurement model and x_0 are mutually independent

Given these assumptions, we employ Bayes' rule to derive the a pair of recursive equations which can be used to find $p(x_k|Z_{1:k})$, namely:

$$\begin{aligned} p(x_{k+1}|Z_{1:k}) &= \int p(x_{k+1}, x_k|Z_{1:k})dx_k \\ &= \int p(x_{k+1}|x_k, Z_{1:k})p(x_k|Z_{1:k})dx_k \\ &= \int p(x_{k+1}|x_k)p(x_k|Z_{1:k})dx_k \end{aligned} \tag{13}$$

$$\begin{aligned} p(x_{k+1}|Z_{1:k+1}) &= p(x_{k+1}|z_{k+1}, Z_{1:k}) \\ &= \frac{p(z_{k+1}|x_{k+1}, Z_{1:k})p(x_{k+1}|Z_{1:k})p(Z_{1:k})}{p(z_{k+1}, Z_{1:k})} \\ &\propto p(z_{k+1}|x_{k+1})p(x_{k+1}|Z_{1:k}) \end{aligned} \tag{14}$$

In order to utilize these equations, we require expressions for the PDFs $p(x_{k+1}|x_k)$ and $p(z_k|x_k)$. Most often, we do this by deriving them from the state and measurement models after introducing further assumptions about state and measurement noise.

2.3.2 Extended Kalman Filter

The Extended Kalman Filter (EKF) is a widely-used special case of a Bayesian Filter where the state and measurement noise PDFs (and thus all subsequent PDFs involved in the estimation of $p(x_k|Z_{1:k})$) are assumed to be Gaussian [19]. Formally, the general state-space model for an EKF can be expressed in terms of a (differentiable) state function f and measurement function h with the following *state-space model*:

$$x_{k+1} = f(x_k, u_k) + w_k \tag{15}$$

$$z_k = h(x_k) + v_k \tag{16}$$

where w_k and v_k are the respective state and measurement noise variables, which have the following statistical properties ($\mathcal{N}$ represents a multivariate normal distribution given by a mean and covariance matrix)

$$w_k \sim \mathcal{N}(0, Q_k) \tag{17}$$

$$v_k \sim \mathcal{N}(0, R_k) \tag{18}$$

$$E(w_k v_k) = 0 \tag{19}$$

Furthermore, we initialize x_k with a starting state estimate x_0 and state covariance matrix P_0 – i.e. $x_0 \sim \mathcal{N}(x_0, P_0)$ – and assume it is uncorrelated with both w_k and v_k.

The EKF estimates the posterior of the state x_k through a prediction and update step which are each based on a Taylor-series linearization of the nonlinear state and measurement functions. The corresponding solutions have the form

Prediction Step

$$\hat{x}_{k|k-1} = f(\hat{x}_{k-1|k-1}, u_k) \tag{20}$$

$$P_{k|k-1} = F_k P_{k-1|k-1} F_k^T + Q_k \tag{21}$$

Update Step

$$K_k = P_{k|k-1} H_k^T (H_k P_{k|k-1} H_k^T + R_k)^{-1} \tag{22}$$

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k(z_k - h(\hat{x}_{k|k-1})) \tag{23}$$

$$P_{k|k} = (I - K_k H_k) P_{k|k-1} \tag{24}$$

Here, $\hat{x}_{k|k-1}$ is the predicted state estimate with covariance $P_{k|k-1}$, $\hat{x}_{k|k}$ is the updated (i.e. posterior) state estimate with covariance $P_{k|k}$, and K_k is the Kalman gain matrix. Because the EKF relies on linearization, we set F_k and H_k to be the Jacobians of the state and measurement function, respectively, evaluated at the most recently estimated state value, i.e.

$$F_k = \left.\frac{\partial f}{\partial \mathbf{x}}\right|_{\hat{x}_{k-1|k-1}} \tag{25}$$

$$H_k = \left.\frac{\partial h}{\partial \mathbf{x}}\right|_{\hat{x}_{k|k-1}} \tag{26}$$

Note that when the state space model is linear, these Jacobians do not depend on state estimates, and we are left with the optimal linear filter, i.e. the Kalman filter [1].

3 Methods

In this section, we propose two different state space models which make use of the EKF to recursively estimate of the instantaneous phase of some known reference curve $\tilde{y}(t)$. In both formulations, we consider modeling an observed signal $x(t)$ as a phase-distorted, amplitude scaled, and vertically shifted noisy instance of the reference curve using the model

$$x(t) \sim a\tilde{y}(\phi(t)) + b \tag{27}$$

These formulations, which we call the Raw Phase Model and Phase Function Model, differ in how they represent the phase estimates, and thus can be useful for different kinds of applications. In particular, the Raw Phase Model includes a simple form of state dynamics based on the instantaneous frequency; this formulation may be altered if we have prior knowledge about the dynamical form of the phase, for instance in cases of polynomial or oscillatory phase signals. On the other hand, the Phase Function Model does not contain state dynamics, but still allows for locally-imposed constraints on an estimated continuous phase function.

Following our formulations, we also show how to impose constraints of monotonicity and, if needed, bounded endpoints, using a projection method. Finally, we discuss practical considerations, including choice of model parameters and methods for utilizing these models to obtain phase estimates when the true reference curve $\tilde{y}(t)$ is not known a priori.

3.1 Model 1: Raw Phase Model

Let us assume for now that we are given the form of a continuous, differentiable, and periodic *template function*, denoted $\tilde{y}(t)$, along with its first derivative, $\tilde{y}'(t)$, and its period, T. We first consider the problem of estimating the raw value of the phase of $\tilde{y}(t)$ at a discrete set of observed time points $t_{1:n}$. In this case, we assume that the observed values, $Z_{1:n}$, represent noisy discretizations of a phase-distorted template curve, whose amplitude may also be scaled and shifted by some constant amounts over the time interval.

3.1.1 State Space Model

Define the state variable at time t_k to be $x_k = [\phi_k, \omega_k, a_k, b_k]^T$, where ϕ_k is the raw phase value of $\tilde{y}$, ω_k is frequency, a_k is an amplitude scaling constant, and b_k is an amplitude shifting constant. We assume that these parameters represent discrete values of underlying functions $\phi(t)$, $\omega(t)$, $a(t) = a$, and $b(t) = b$.

Using the relationship $\phi'(t) = \omega(t)$ for phase function $\phi(t)$ and frequency function $\omega(t)$, we can write the Taylor series expansion of $\phi(t)$ about time t_k as

$$\phi(t) = \phi(t_k) + (t - t_k)\omega(t_k) + \frac{(t - t_k)^2}{2}\omega'(t_k) + \dots \tag{28}$$

which implies for discrete values

$$\phi_{k+1} = \phi_k + \Delta t_k \omega_k + \frac{{\Delta t_k}^2}{2}\omega'_k + \dots \tag{29}$$

Since we are primarily dealing with signals that are nearly-periodic, we make the assumption that the frequency function $\omega(t)$ is approximately constant for time difference much smaller than period T; formally, this is expressed as $|\omega(t + \Delta t) - \omega(t)|/|\omega(t)| \approx 0$ for $\Delta t << T$. Thus, we can treat the higher-order terms of the Taylor expansion as state noise and are left with the following state-space model:

$$x_{k+1} = f(x_k) + w_k = A_k x_k + w_k = \begin{bmatrix} 1 & \Delta t_k & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \phi_k \\ \omega_k \\ a_k \\ b_k \end{bmatrix} + w_k \tag{30}$$

$$z_k = h(x_k) + v_k = a_k \tilde{y}(\phi_k) + b_k + v_k \tag{31}$$

Since the state equation is linear, we use $F_k = A_k$ for the EKF prediction step. The measurement equation, however, is nonlinear, so we need an expression for the Jacobian to use for the EKF update step, namely,

$$H_k = \left.\frac{\partial h}{\partial \mathbf{x}}\right|_{\hat{x}_{k|k-1}} = \begin{bmatrix} \hat{a}_{k|k-1}\tilde{y}'(\hat{\phi}_{k|k-1}) & 0 & \tilde{y}(\hat{\phi}_{k|k-1}) & 1 \end{bmatrix} \tag{32}$$

Monotonicity Constraint

Given that the true phase function $\phi(t)$ is monotonically increasing across time, we have that $\omega(t) = \phi'(t) > 0$. Thus, we can impose the constraint that $\omega_k > 0$ for all k. Note that applying this constraint at each iteration of the EKF is a soft constraint of sorts, because the posterior phase estimates for the current time point will not change (only later phase estimates will be

affected).

In order to avoid estimating frequencies which would be subject to aliasing, we also constrain $\omega(t)$ to be less than the Nyquist frequency $f_s/2$ for sampling rate $f_s = \frac{1}{\Delta t_k}$.

3.2 Model 2: Phase Function Model

Our second model considers the problem of recursively estimating the *parameters of a function* representing the phase of the periodic waveform $\tilde{y}(t)$. This method can also be thought of as a recursive analog to curve registration in FDA, where here we are estimating the *inverse warping functions*, $h^{-1}(t)$, that optimally map template time onto observed time. However, unlike FDA curve registration where the reference curve is defined for one period (finite time) and each instance of the phase distorted reference curve is known a priori, we allow the EKF method to uncover multiple instances of the template waveform within a single phase-distorted signal. To show correspondence with FDA curve registration, we impose endpoint constraints that allow for unwarping when the number of periods spanned by each observed curve is known (e.g. if we have a set of curves each representing one period of the template).

3.2.1 B-Spline Functions

In order to develop this method, we require a basis function representation of the phase $\phi(t)$ which can describe smooth, arbitrarily local distortions. B-Splines provide a natural choice of basis because they are (d-1)-times continuously differentiable for order d, and their locality can be easily adjusted by changing the number and position of control points, or knots. Furthermore, as we will see, B-splines can lend themselves well to imposing local monotonicity constraints, as well as endpoint constraints.

Given a set of $n = J + d + 1$ strictly increasing knots $\mathbf{K} = (K_1, ..., K_n)$, a *B-Spline function* can be defined as a weighted sum of J nonnegative piecewise polynomial functions of degree d. Each B-spline basis, denoted $b_i(t)$ for $i = 1, 2, ..., J$, is nonzero for $t \in [K_i, K_{i+d+1})$ and zero elsewhere. Carl de-Boor showed that these basis functions can be defined recursively for increasing order d as follows, with $b_{j,\delta}(t)$ representing the j-th δ-degree basis [5]:

Cox-de Boor formula

$$b_{j,0}(t) = \begin{cases} 1 & K_j \leq t < K_{j+1} \\ 0 & \text{otherwise} \end{cases} \tag{33}$$

$$b_{j,\delta}(t) = \frac{t - K_j}{K_{j+\delta+1} - K_j} b_{j,\delta-1}(t) + \frac{K_{j+\delta} - t}{K_{j+\delta} - K_{j+1}} b_{j+1,\delta-1}(t) \tag{34}$$

These expressions imply that for a given $t \in [K_j, K_{j+1}]$, the value of the unweighted B-spline function can be obtained from d+1 of the J bases, namely: $b_{j-d}(t)$, $b_{j-d+1}(t)$, ..., $b_j(t)$. We can compute the row vector of these d+1 basis values, denoted $\mathbf{b}_{j,d}(t)$, as in [11]:

$$\mathbf{b}_{j,d}(t) = \mathbf{B}_{j,1}(t)\mathbf{B}_{j,2}(t)...\mathbf{B}_{j,\delta}(t)...\mathbf{B}_{j,d}(t) \tag{35}$$

where each $\mathbf{B}_{j,\delta}(t)$ is a $\delta \times (\delta + 1)$ matrix defined as

$$\mathbf{B}_{j,\delta}(t) = \begin{bmatrix} \frac{K_{j+1}-t}{K_{j+1}-K_{j+1-\delta}} & \frac{t-K_{j+1-\delta}}{K_{j+1}-K_{j+1-\delta}} & 0 & \cdots \\ & \ddots & \ddots & \\ 0 & \cdots & \frac{K_{j+\delta}-t}{K_{j+\delta}-K_j} & \frac{t-K_j}{K_{j+\delta}-K_j} \end{bmatrix} \tag{36}$$

The overall B-spline function $f(t)$ is defined for $t \in [K_{d+1}, K_{J+1})$ as

$$f(t) = \mathbf{b}_{j,d}(t)\mathbf{x}_{j,d} \tag{37}$$

where j indexes the knot interval such that $t \in [K_j, K_{j+1})$ and $\mathbf{x}_{j,d} \in \mathbb{R}^{(d+1)\times 1}$ is the coefficient vector over that knot interval.

Monotonicity Constraint

In order to impose a monotonicity constraint, we will also require the form of the first derivative of the B-spline function. This is given by

$$\frac{\partial}{\partial t} f_d(t) = \begin{cases} \frac{\partial}{\partial t}\mathbf{b}_{d+1,d}(t)\mathbf{x}_{d+1,d} & t \in [K_{d+1}, K_{d+2}) \\ \frac{\partial}{\partial t}\mathbf{b}_{d+2,d}(t)\mathbf{x}_{d+2,d} & t \in [K_{d+2}, K_{d+3}) \\ \dots & \\ \frac{\partial}{\partial t}\mathbf{b}_{J,d}(t)\mathbf{x}_{J,d} & t \in [K_J, K_{J+1}) \end{cases} \tag{38}$$

$$\frac{\partial}{\partial t}\mathbf{b}_{j,d}(t) = d\mathbf{B}_{j,1}(t)\dots\mathbf{B}_{j,d-1}(t)\mathbf{B}'_{j,d}(t) \tag{39}$$

where each $\mathbf{B}_{j,\delta}$ is defined above, and $\mathbf{B}'_{j,d}(t)$ is the element-wise derivative of $\mathbf{B}_{j,d}(t)$, i.e.:

$$\mathbf{B}'_{j,d}(t) = \begin{bmatrix} \frac{-1}{K_{j+1}-K_{j+1-d}} & \frac{1}{K_{j+1}-K_{j+1-d}} & 0 & \cdots \\ 0 & \frac{K_{j+2}-t}{K_{j+2}-K_{j+2-\delta}} & \frac{t-K_{j+2-\delta}}{K_{j+2}-K_{j+2-\delta}} & \cdots \\ & \ddots & \ddots & \\ 0 & \cdots & \frac{-1}{K_{j+d}-K_j} & \frac{1}{K_{j+d}-K_j} \end{bmatrix} \tag{40}$$

In the EKF model (given below), we will consider imposing the monotonicity constraint on the B-spline phase function one of two ways. Either we can constrain the derivative to be positive *at the current timepoint* after each update step of the EKF, or we can impose a global constraint by restricting the B-spline coefficients $\mathbf{x}$ to be increasing values. For verification that increasing coefficient values leads to monotonically increasing B-spline functions of first order, see Appendix.

Endpoint Constraints

We may also wish to constrain the B-spline function so that the values at each endpoint differ by a fixed amount. This is especially useful for comparing the EKF results to the FDA Curve Registration results, where we set the difference in endpoints to be equal to one period of the template waveform. Additionally, these constraints may allow for more accurate estimates of interbeat phase distortions (i.e. microstructure) when the number of observed beats is known (or approximately known). The form for imposing these constraints can be easily obtained from the B-spline evaluation vectors, $\mathbf{b}_{j,d}$, for known ending time points.

3.2.2 State-Space Model

Now that we know the form of the B-spline functions and we have natural ways of constraining them to be monotonically increasing or bounded at the endpoints, we can define our state space model.

Let h_k represent J coefficients of a degree d B-spline function with an increasing knot sequence $\mathbf{K} = (K_1, \ldots, K_{J+d+1})$. For simplicity, we will assume that knots are uniformly spaced and defined such that $K_{d+1} \leq t_1$ and $K_{J+1} > t_n$ for the observed starting and ending time points, t_1 and t_n, respectively.

For greater ease of imposing the increasing-coefficient monotonicity constraint, we define a new state variable x_k to represent consecutive differences of the spline coefficients, namely

$$x_k = L^{-1} h_k \tag{41}$$

for lower triangular matrix $L \in \mathbb{R}^{J \times J}$ with all nonzero elements equal to 1. Thus, we can obtain monotonicity by restricting the last $J-1$ elements of x_k to be greater than 0. Now letting $\tilde{x}_k = [x_k, a_k, b_k]^T$ be the overall state variable with amplitude scale factor a_k and shift factor b_k, we write the state-space model as

$$\tilde{x}_{k+1} = A_k \tilde{x}_k + w_k \tag{42}$$

$$z_k = h(\tilde{x}_k) + v_k = a_k \tilde{y}(C_k L x_k) + b_k + v_k \tag{43}$$

where $A_k = \mathbf{I}$, and $C_k = [\mathbf{0}_{1\times(j-d-1)}, \mathbf{b}_{j,d}(t_k), \mathbf{0}_{1\times(J-j)}]$ gives the value of the J B-spline bases at time $t_k \in [K_j, K_{j+1}) \subset [K_{d+1}, K_{J+1})$.

Lastly, to linearize the measurement function for the update step, we use the following Jacobian:

$$H_k = \left.\frac{\partial h}{\partial \mathbf{x}}\right|_{\hat{\hat{x}}_{k|k-1}} = \begin{bmatrix} \hat{a}_{k|k-1}\tilde{y}'(C_k L\hat{x}_{k|k-1})(C_k^{(1)} + C_k^{(2)} + \ldots + C_k^{(J)}) \\ \vdots \\ \hat{a}_{k|k-1}\tilde{y}'(C_k L\hat{x}_{k|k-1})(C_k^{(J)}) \\ \tilde{y}(C_k L\hat{x}_{k|k-1}) \\ 1 \end{bmatrix}^T \tag{44}$$

3.3 State Constraints

The constraints on monotonicity that we have discussed can each be expressed as linear inequality constraints of the form (see Appendix for particular choices of constraint matrices)

$$Dx_k \leq d \tag{45}$$

We impose these constraints within the EKF using a projection method, as presented in [20], which we then feed into the next prediction step. This projection is found by minimizing

$$(x - \hat{x}_{k|k})^T W (x - \hat{x}_{k|k}) \tag{46}$$

subject to the constraint in (45), where W is a positive definite weighting matrix and $\hat{x}_{k|k}$ is the posterior unconstrained state estimate. For linear *equality* constraints, the solution to this projection is given by

$$x_k^c = x_k^{uc} - W^{-1}D^T(DW^{-1}D^T)^{-1}(Dx_k^{uc} - d) \tag{47}$$

While inequality constraints typically require more steps, often utilizing *active set methods*, we note here that our constraint matrices D either consist of one row, or constrain each state parameter independently. Thus, to impose our constraints, we simply check if the inequality constraint is satisfied, and if not, solve the projection for corresponding equality constraint.

When implementing these constraints within the EKF, we choose $W = P_k^{-1}$ for the current state estimate covariance P_k, which gives a minimum variance filter [20].

3.4 Practical Considerations

Model Parameters

The two models we have laid out depend on a variety of parameters, including the state and measurement noise covariance matrices, Q and R, the initial state estimate covariance matrix P_0, the degree and number of knots of the B-spline function, and the given template curve $\tilde{y}(t)$.

Broadly speaking, the choice of these parameters can adjust one of three things: the locality or nonlinearity of the phase estimates (i.e. how rapidly the phase functions vary), the number of curve features which can be properly aligned, and in some cases, the stability and accuracy of the EKF results. For instance, a higher choice of the state noise of the frequency estimate ω_k in the Model 1 is more suitable when we expect to have more nonlinear phase, whereas a higher choice of the state noise for the amplitude factors may allow for more detrending (but should be carefully chosen to avoid overfitting amplitude and giving poor phase results).

For the heartbeat simulations presented in the next section, we obtain simulated electrocardiogram (ECG) waveforms as numerical solutions of the dynamical model proposed in [16]. Furthermore, we set B-spline knots to be uniformly spaced, with the number of knots chosen to have approximately 2-3 knots within each beat cycle so as to provide some information about inter-beat phase. Exact choices of parameters for each presented simulation are given in the Appendix; these choices were determined by experiment.

Initial Phase Estimates

Another issue with our EKF formulations is that they can be subject to biased estimates of the initial phase offset. One method for getting around this is to use a Kalman Smoother, which essentially runs the EKF in both directions, adjusting constraints and state models accordingly. The resulting estimates for both the forward and backward EKF may then be weighted together according to their respective state estimate covariance matrices P_k [2].

Obtaining Template Curves

Lastly, there is the fundamental problem of obtaining a proper choice of $\tilde{y}(t)$ when we do not have a good a priori representation. One crude method that works when we have a known set of

curves is the Procrustes algorithm as originally proposed by Ramsay and Silverman, which uses the cross-sectional average of those curves as an initial guess of template, registers each curve to that template, and iteratively obtains new templates after each registration step [17].

We can come up with an analogous algorithm using our EKF models when we do not know the set of curves in advance. In particular, we consider the case of ECG or ECG-like signals, where there are dominant peaks occurring within each beat cycle. In these cases, we can first run our EKF with an impulse function as our template [2] to segment individuals beats. Following this, we obtain a more refined estimate of the template curve in a similar manner to the Procrustes algorithm (i.e. using cross-sectional average of each waveform period). Note that since we require $\tilde{y}(t)$ to be differentiable and periodic, we can fit the averaged curve with a smoothing spline function (i.e. B-spline function) and constrain its endpoints to be equal so that it can be made periodic without losing continuity.

4 Numerical Studies

In this section, we present various simulated examples which show the performance of the EKF methods across different kinds phase-distorted signals. These numerical examples are used to show the correspondence between each of the discussed methods for phase estimation (i.e. Fourier methods, global FDA optimization, and the EKFs), and the respective limitations and guiding factors behind each one. Illustrations of how to use the EKF methods for recursive curve registration and iterative ECG template uncovering are also shown.

4.1 Sinusoidal Examples

We simulated three classes of noisy sinusoidal signals corresponding to increasing phase complexity: namely, linear phase, quadratic phase (i.e. chirp signals), and cubic B-spline phase. In each case, we applied three methods of phase estimation – the Hilbert Transform, Model 1 EKF, and Model 2 EKF – and calculated the error between estimates and the known phase function. Results for each class of sinusoidal signal are shown in Figures 1, 2, and 3, respectively.

[2] Alternatively, we may start with a simulated ECG template curve, which we use to iteratively find a a template curve more suited to that particular ECG signal

Class	MSE (Hilbert)	MSE (EKF 1)	MSE (EKF 2)
Linear	2.51×10^{-4}	1.08×10^{-5}	1.01×10^{-6}
Quadratic	4.17×10^{-4}	3.47×10^{-5}	2.71×10^{-5}
Cubic B-Spline	3.51×10^{-4}	6.42×10^{-5}	1.87×10^{-6}

Table 1: Mean squared errors in phase estimates for Hilbert method, EKF Model 1, and EKF Model 2 for three classes of sinusoidal signals.

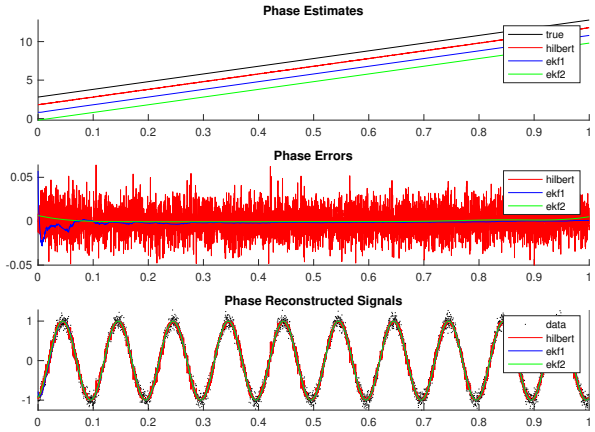

Figure 1: Linear Phase Sine Wave with mean-zero Gaussian noise (sd = .1). True normalized phase function is given by $\phi(t) = 10t - .2$. Phase estimates are shown with offsets of 1 period for ease of visibility.

In all shown cases, the mean squared errors were highest for the Hilbert method and lowest for the Phase Function EKF model (see Table 1). However, as seen in Figure 4, the Hilbert transform performs better than the EKF methods for the linear phase case when the amount of noise is very low; as more noise is injected, the EKF methods become more suitable as choices of models and ultimately perform similar to each other. Additionally for the nonlinear cases, the Hilbert transform method exhibits more artefacts at the edges.

Note that while the Phase Function Model generally performs better than the Raw Phase Model for relatively low levels of noise, this discrepancy is especially present in the nonlinear phase cases. This makes sense given that we formulated Model 1 in a way that is best suited for near-linear phase estimation.

Figure 5 provides an example for why the Hilbert transform is only appropriate in select cases. Here, we have a sine wave with two harmonic components – in this case $y(t) = sin(4\pi t) + sin(2\pi t)$

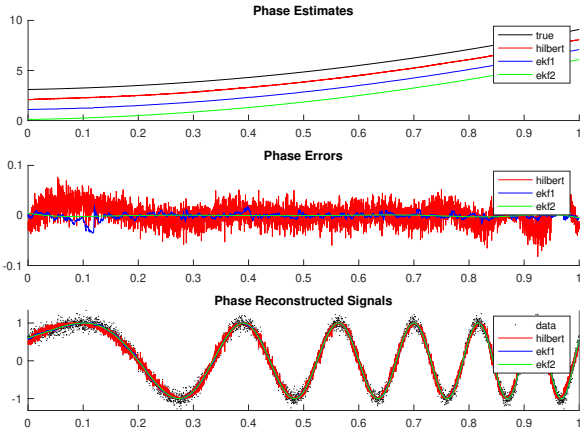

Figure 2: Quadratic phase sine wave with mean-zero Gaussian noise (sd = .1). True normalized phase function is given by $\phi(t) = 5t^2 + t + .1$. Phase estimates are shown with offsets of T=1 for ease of visibility.

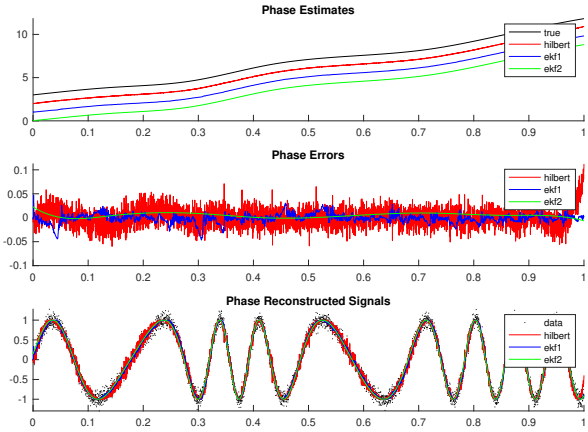

Figure 3: Sine wave with monotonically increasing cubic B-spline with J=10 coefficients in Gaussian noise (sd = .1). Phase estimates are shown with offsets of T=1 for ease of visibility.

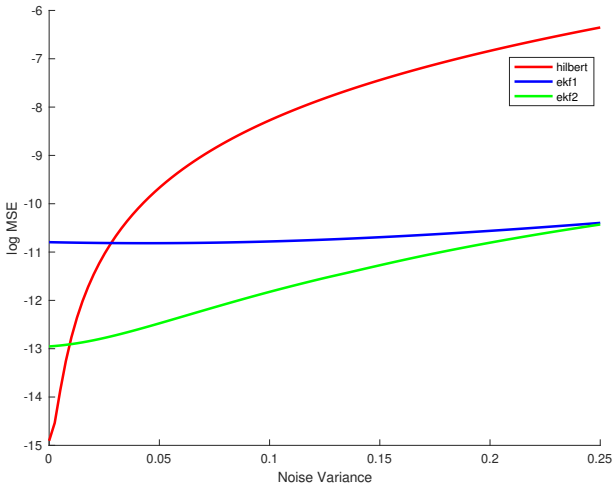

Figure 4: Performance of phase estimation techniques for increasing noise variances in the linear sinusoid case.

– with nonlinear phase given by an increasing spline function. The Hilbert estimate for this signal picks up the double-beat structure (as evident from the phase reconstructed signal), but does not provide an adequate, continuous characterization of the overall phase evolution of the double-beat waveform (see sharp cutoffs between sub-beat components).

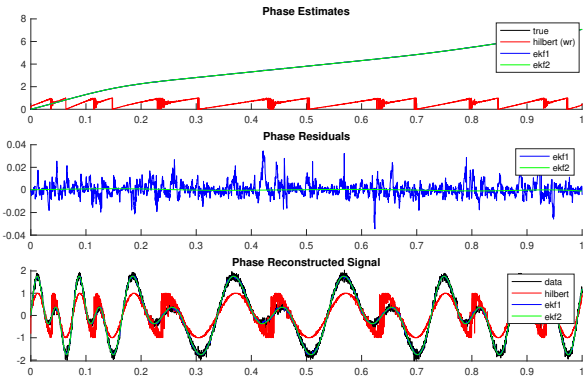

Figure 5: Phase estimates for double beat sinusoid with cubic B-spline phase function. Hilbert estimates are shown in wrapped form for ease of visibility.

4.2 Recursive Curve Registration

Next, we simulated how the Phase Function EKF Model performs relative to global FDA optimization on a multi-component sinusoidal signal. For each method, we constrained the difference in endpoints for the phase function (i.e. inverse warping function) to be equal to $T = 1$, and used cubic B-splines with $J = 10$ coefficients to obtain estimates of the phase function (either directly for the EKF and or using the integral constraint for FDA). For this example, the template curve was given by $\tilde{y}(t) = sin(2\pi t) + sin(4\pi t) + sin(6\pi t)$, and the true inverse warping functions were set as one-parameter exponential functions of the form

$$h^{-1}(t|\beta) = \frac{1}{\beta} log\left[t(e^{\beta} - 1) + 1\right] \tag{48}$$

As seen in Figure 6, curve registration was successfully applied in both the FDA and EKF approaches. We note that, potentially due to the flexibility of the locally-constrained B-spline state functions, the recursive method performed slightly better than the FDA method. The corresponding averages in mean squared errors across all 6 inverse warping function estimates were 1.42×10^{-7} for the EKF method and 1.97×10^{-5} for the FDA method.

To get a sense of how the EKF estimates unfold over time, the inverse warping function estimates are shown after 50% of the data has been processed (Figure 7) and after 100% of data has been processed (Figure 8). We note here that monotonicity was applied locally using the

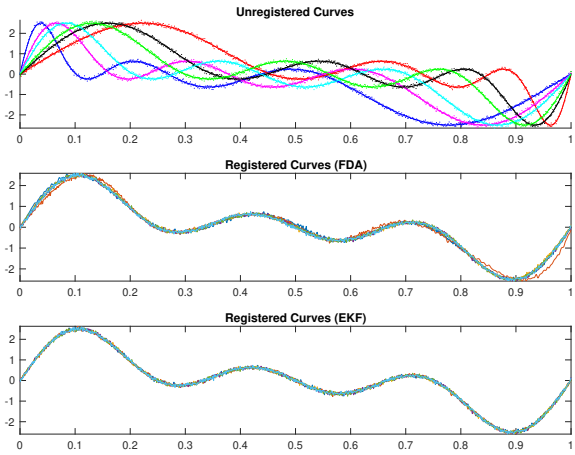

Figure 6: Result of curve registration using the FDA and EKF approaches for a three-component sinusoid template. Endpoint constraints are imposed to prevent registered curves from exceeding 0 to 1 range

derivative formulation, so only the parts of the function up to the current time point are expected to be increasing.

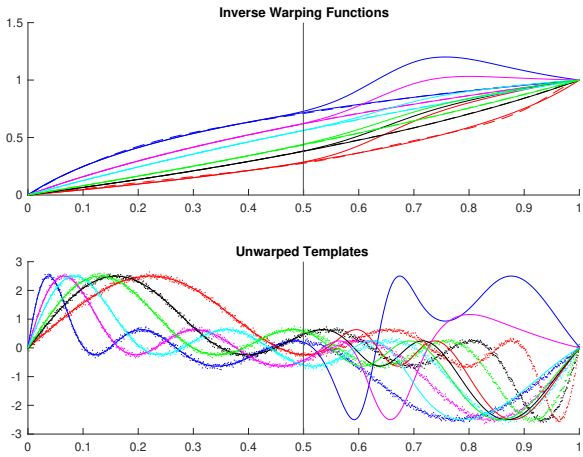

Figure 7: Inverse warping functions and unwarped template curves after 50% of data has been processed by EKF. Solid lines represent EKF estimates, while dashed lines represent FDA estimates of inverse warping functions. Note that monotonicity constraint is local, and does not apply to future curve values, unlike endpoint constraints.

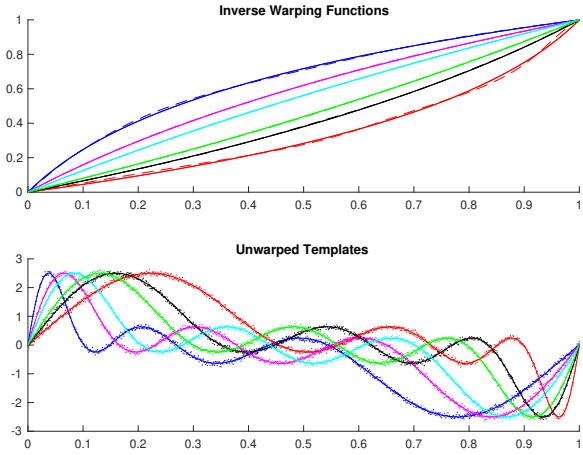

Figure 8: Inverse warping functions and unwarped template curves after 100% of data has been processed by EKF. Solid lines represent EKF estimates, while dashed lines represent FDA estimates of inverse warping functions.

4.3 Simulated ECG Signals

Lastly, we applied the recursive techniques to simulated noisy electrocardiogram (ECG) signals. Broadly speaking, these signals can be described in terms of five major landmarks within each heartbeat cycle, known as the P, Q, R, S, and T wave [16]. The most dominant of these waves is the R wave, which is most frequently used to estimate heart rate (HR) and beat-to-beat variability. However, the relative positioning of these R waves alone does not fully encapsulate the micro-structure phase variability across local time scales. Therefore, we applied our techniques of recursive phase estimation to determine whether or not local, nonlinear phase distortions could be estimated from simulated ECG data. As seen in Figure 9, nonlinear phase distortions were reasonably estimated by both EKF models using a known ECG waveform as a template, with mean squared errors of 1.72×10^{-4} for the Model 1 phase estimates, and 1.07×10^{-5} for the Model 2 phase estimates.

However, since in practice we do not known the particular ECG waveform of a real signal in advance, we demonstrate the results of the proposed refinement method in Figure 10. Here, the initial choice of template function was an impulse-like wave of the form $\tilde{y}(t) = sin(\pi t)^{100}$. Using this initial template, the location of dominant R peaks was identified by running the Phase Function EKF Model on a simulated ECG with low signal-to-noise ratio; the phase estimates for this initial iteration corresponded with a mean squared error of 0.0121 and a mean error of 0.0855. After refining the template curve using a smoothing spline (order 3, 15 coefficients), the resulting phase estimates had a mean squared error of 4.93×10^{-4} and a mean error of .0072, indicating that more accurate results were achieved following refinement. We note that the initial iteration gave systematically biased phase estimates with errors greater than zero for nearly all observed time points. This can be explained by the fact that the initial impulse curve is symmetric about its peaks, whereas the ECG waveform is not.

Comments

It is worth mentioning that the the EKF results can depend heavily on whether or not we implement constraints. This is particularly important when we have uncertainty about amplitude variation, and thus require a method for "guiding" the phase estimates in a way that reduces

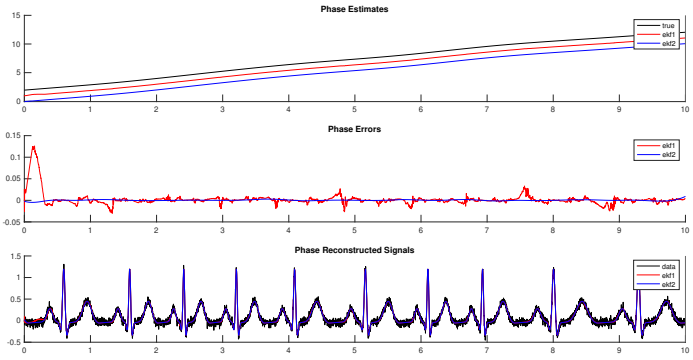

Figure 9: Simulated ECG Curve with phase distortions modeled by an increasing cubic B-spline

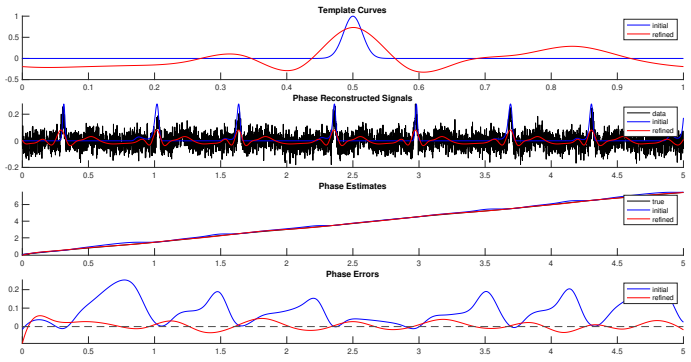

Figure 10: Illustration of iterative EKF process used to get refined estimates of underlying template curve. Initial estimates are given by the blue curves, whereas red curves come from refinement using averaged initially-identified beat intervals.

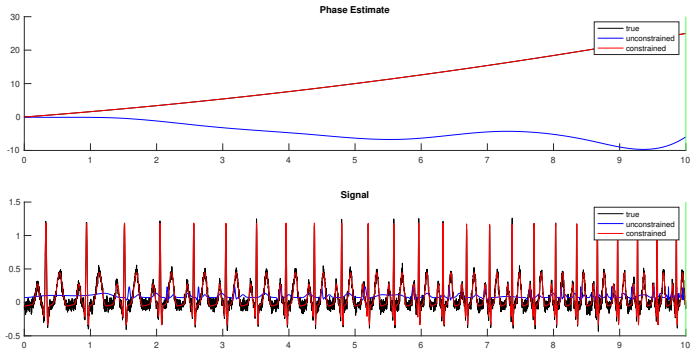

Figure 11: Results of EKF Functional Phase Model with and without derivative constraints when initial amplitude uncertainty is high (P=10^4)

ambiguity in our state. Figure 11 provides an example for how important this guiding process can be in the case where we set the initial covariance of amplitude parameters to be high. We also mention, but do not illustrate here, that constraints may aid in accuracy when the template function is highly nonlinear.

Additionally, it is worth mentioning that the problem of optimizing EKF model parameters for a given data set merits further attention. This is especially important because it may allow for more adaptive modeling techniques, where initial choices of model parameters can be corrected over time.

5 Discussion

5.1 Conclusions and Limitations

Our results using simulated data indicate that the constrained Extended Kalman Filter may be used as a reasonably robust tool for the recursive estimation of the phase of a noisy waveform, particularly in cases when underlying template structure is known or approximately known. Furthermore, we have presented examples for how these approaches can achieve better results than standard methods like the Hilbert transform with signal-to-noise ratio is low. By imposing appropriate constraints, these methods also effectively function as a recursive analog to FDA curve registration, but with potential advantages for modeling local phenomena.

That being said, the proposed models come with several limitations, mostly stemming from the fact that the EKF approximates systems which may be quite nonlinear in nature using a linearization. This can lead us to rely heavily on parameter tuning, which may prevent models from being easily generalizable to broad classes of nonlinear and nonstationary signals. As a more robust alternative, it may be worthwhile to use Machine Learning methods for parameter tuning, or studying how nonlinear filters, like the Gaussian filter, may be used to tackle the problem of recursive general phase estimation.

5.2 Theoretical Considerations

Here, we note a few other conceptual frameworks which were considered for addressing the problem of recursive phase estimation but ultimately not applied in this thesis due to practical limitations.

For one, we know that we can represent a sine or cosine wave as the solution to a second-order linear differential equation of the form

$$\ddot{x} = -k^2 x \tag{49}$$

Therefore, for at least one of the cases presented above (i.e. the linear phase, sinusoidal case), we may come up with a linear state-space model to track the dynamics of the observed signal recursively. This is advantageous because it means we can use an optimal Kalman Filtering approach for state estimation, versus an EKF which in general is not optimal. However, it is not obvious how to utilize this dynamical formulation when the sinusoidal phase is nonlinear (or naturally, when the reference function is non-sinusoidal).

Another kind of estimation scheme which was partially derived but ultimately not applied here involves the use of a dual-state Bayesian filter, where we wish to estimate both the template curve $\tilde{y}(t)$ and its phase recursively, without placing additional assumptions on the starting guess of $\tilde{y}(t)$. Ultimately, this is worth further study as a step towards a completely recursive version of Functional Data Analysis that does not rely on template-approximation heuristics like the Procrustes algorithm.

5.3 Further Applications

Finally, we discuss further applications which may make use of recursive phase estimation techniques. In particular, we are interested in analyzing a variety of heartbeat signals which have high amounts of noise and thus may not be easily processed using traditional methods like peak picking. One specific application which we began to explore involves a cross-validation scheme on different individuals in order to compare phase deviations in heartbeat microstructure. Given that most classification schemes of this nature are based either on morphological differences (i.e. amplitude variation) [7] or interbeat phase variation (i.e. heart rate variability) [6], assessing microstructure phase variation would be a novel application worth further study.

Another interesting application related to fetal heartbeats is the quantification of *phase synchrony* between the fetal and maternal heartbeats [10]. That is, can the phase nonlinearities between the heartbeat waveforms of the mother and fetus be represented as an interconnected dynamical system? For this to work using recursive approaches, methods for recursive separation of fetal heart beat and maternal heart beat signals are required.

Finally, as stated in the Introduction, the ultimate motivation behind this thesis was to develop pipelines for real-time behavioral classification of fetuses and infants at-risk of autism. As next steps for conducting this research, we plan to use the presented processing methods in order to characterize how the fetal heartbeat responds to external social stimuli like the mother's voice, in controlled laboratory settings.

6 References

[1] Barker, A. L., Brown, D. E., and Martin, W. N. "Bayesian estimation and the Kalman filter". In: *Computers and Mathematics with Applications* 30.10 (1995), pp. 55–77.

[2] Barratt, S. and Boyd, S. "Fitting a Kalman Smoother to Data". In: *arXiv* (2019), pp. 1–15.

[3] Behar, J. A. et al. "Noninvasive fetal electrocardiography for the detection of fetal arrhythmias". In: *Prenatal Diagnosis* 39.3 (2019), pp. 178–187.

[4] Berndt, D. and Clifford, J. "Using dynamic time warping to find patterns in time series". In: *AAAI: Workshop on Knowledge Discovery in Databases* (1994), pp. 359–370.

[5] de-Boor, Carl. *A Practical Guide to Splines*. New York: Springer-Verlag, 1978.

[6] Dipietro, J. A. et al. "Fetal Heart Rate and Variability: Stability and Prediction to Developmental Outcomes in Early Childhood". In: *Child Development* 78.6 (2007), pp. 1788–1798.

[7] Fratini, A. et al. "Individual identification via electrocardiogram analysis". In: *BioMedical Engineering OnLine* (2015), pp. 1–23.

[8] Hahn, S. "Hilbert Transforms". In: *The Handbook of Formulas and Tables for Signal Processing* (1999).

[9] Hansen, J. *Selected approaches to estimation of signal phase*. Tech. rep. 2003, pp. 1–15.

[10] Ivanov, P., Ma, Q., and Bartsch, R. "Maternal – fetal heartbeat phase synchronization". In: *PNAS* 106.33 (2009), pp. 13641–13642.

[11] Jauch, J. et al. "Recursive B-spline approximation using the Kalman filter". In: *Engineering Science and Technology, an International Journal* 20.1 (2017), pp. 28–34.

[12] Kisilevsky, B. and Hains, S. "Onset and maturation of fetal heart rate response to the mother's voice over late gestation". In: *Developmental Science* 14.2 (2011), pp. 214–223.

[13] Krawczyk, M. and Gerkmann, T. "STFT phase reconstruction in voiced speech for an improved single-channel speech enhancement". In: *IEEE/ACM Transactions on Audio Speech and Language Processing* 22.12 (2014), pp. 1931–1940.

[14] Kurz, G. and Hanebeck, U. "Heart phase estimation using directional statistics for robotic beating heart surgery". In: *2015 18th International Conference on Information Fusion* (2015), pp. 703–710.

[15] McGowan, R. "Direct Relation Between a Signal Time Series and Its Unwrapped Phase". In: *IEEE Transactions on Acoustics, Speech, and Signal Processing* 30.5 (1982), pp. 719–726.

[16] Mcsharry, P. et al. "Electrocardiogram Signals". In: *IEEE Transactions on Biomedical Engineering* 50.3 (2003), pp. 289–294. ISSN: 0018-9294.

[17] Ramsay, J. and Li, X. "Curve Registration". In: *Royal Statistical Society* 60.2 (2010), pp. 351–363.

[18] Ramsay, J. and Silverman, B. *Functional data analysis*. Springer, 2010.

[19] Ribeiro, M. "Kalman and Extended Kalman Filters : Concept , Derivation and Properties". In: *Institute for Systems and Robotics* February (2004), p. 42.

[20] Simon, D. "Kalman filtering with state constraints: A survey of linear and nonlinear algorithms". In: *IET Control Theory and Applications* 4.8 (2010), pp. 1303–1318.

[21] Slaets, L., Claeskens, G., and Silverman, B. "Warping functional data in R and C via a bayesian multiresolution approach". In: *Journal of Statistical Software* 55.3 (2013).

7 Appendix

Increasing Coefficients

For first-order B-spline functions, the following holds for $2 \leq j \leq J$:

$$\frac{\partial}{\partial t}\mathbf{b}_{j,1}(t)\mathbf{x}_{j,1} = \begin{bmatrix} \frac{-1}{K_{j+1}-K_j} & \frac{1}{K_{j+1}-K_j} \end{bmatrix} \begin{bmatrix} x_j^{(1)} \\ x_j^{(2)} \end{bmatrix} = \frac{1}{\Delta K_j}(x_j^{(2)} - x_j^{(1)}) \tag{50}$$

from which we obtain the following monotonicity constraints on the first degree B-spline function $f_{d=1}(t)$

$$\frac{\partial}{\partial t} f_{d=1}(t) > 0 \implies \begin{cases} (x_2^{(2)} - x_2^{(1)})/\Delta K_2 > 0 \\ (x_3^{(2)} - x_3^{(1)})/\Delta K_3 > 0 \\ \dots \\ (x_J^{(2)} - x_J^{(1)})/\Delta K_J > 0 \end{cases} \tag{51}$$

$$\implies x^{(1)} < x^{(2)} < \ldots < x^{(J)} \tag{52}$$

Thus, monotonicity over the domain $[K_{d+1}, K_{J+1})$ is obtained by restricting the B-spline coefficients to be increasing.

State Constraints

Model 1: Bounded Frequency

$$0 \leq \begin{bmatrix} 0 & 1 & 0 & 0 \end{bmatrix} x_k \leq f_s/2 \tag{53}$$

Model 2 – Increasing Coefficients (1st order):

$$\begin{bmatrix} -1 & 1 & 0 & \dots & 0 & 0 \\ 0 & -1 & 1 & \dots & 0 & 0 \\ \vdots & & & & & \\ 0 & 0 & 0 & \dots & -1 & 1 \end{bmatrix} h_k = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & & & & \\ 0 & 0 & 0 & \dots & 1 \end{bmatrix} x_k \geq 0 \tag{54}$$

Model 2 – Derivative Constraint

$$D = \left[\mathbf{0}_{1\times(j-d-1)}, \tfrac{\partial}{\partial t}\mathbf{b}_{j,d}(t_k), \mathbf{0}_{1\times(J-j)} \right] \tag{55}$$

$$0 \leq lb \leq Dh_k = DLx_k \leq ub \tag{56}$$

Model 2 – Endpoint Constraints (n periods, for period T)

$$C_j = \left[\mathbf{0}_{1\times(j-d-1)}, \mathbf{b}_{j,d}(t_k), \mathbf{0}_{1\times(J-j)} \right] \tag{57}$$

$$D = C_{J+1} - C_{d+1} \tag{58}$$

$$Dh_k = DLx_k = nT \tag{59}$$

Simulation Parameters

The choice of model parameters for each presented simulation are given here. Subscripts indicate which model the parameters correspond to (EKF Model 1 or EKF Model 2). Each example is based on the following parameters:

- $\tilde{y}(t)$ - Chosen template function
- $\phi(t)$ - True phase function
- Q - EKF state noise covariance matrix
- R - EKF measurement noise variance
- L - Length of signal

- $var(\epsilon)$ - Simulated noise variance
- x_0 - Initial state estimate
- P_0 - Initial state covariance matrix
- d - Degree of estimated phase spline function
- J - Number of coefficients in estimated phase spline function

Simulation 1: Linear Phase Sinusoid (Figure 1)

- $\tilde{y}(t) = sin(2\pi t)$
- $\phi(t) = 10t - .2$
- Q_1 = diag(1e-12,. . .,1e-12), Q_2 = diag(1e-12,. . .,1e-12)
- $R_1 = 1$, $R_2 = 1$
- L = 5000
- d = 3
- J = 10
- $var(\epsilon) = .01$
- $x_0 = 0$
- P_0=diag(100,. . .,10,10)

Simulation 2: Quadratic Phase Sinusoid (Figure 2)

- $\tilde{y}(t) = sin(2\pi t)$
- $\phi(t) = 5t^2 + t + .1$

- Q_1 = diag(1e-12,1,1e-12,1e-12), Q_2 = diag(1e-12,. . .,1e-12)
- $R_1 = 1$, $R_2 = 10$
- L = 5000
- d = 3
- J = 10
- $var(\epsilon) = .01$
- $x_0 = 0$
- P_0=diag(100,. . .,10,10)

Simulation 3: B-Spline Sinusoid (Figure 3)

- $\tilde{y}(t) = sin(2\pi t)$
- $\phi(t)$ = Bspline(K, x = [-1 0 1 1.2 3.8 4.5 5 7 9 10], d=3), K uniformly spaced with K(4)=0 and K(J+1)=1
- Q_1 = diag(1e-12,1,1e-12,1e-12), Q_2 = diag(1e-12,. . .,1e-12)
- $R_1 = 1$, $R_2 = 1$
- L = 5000
- d = 3
- J = 10
- $var(\epsilon) = .01$
- $x_0 = 0$
- P_0=diag(100,. . .,10,10)

Simulation 4: Double Beat Sinusoid (Figure 5)

- $\tilde{y}(t) = sin(4\pi t) + sin(2\pi t + \pi/2)$
- $\phi(t)$ = Bspline(x=[-1 0 1 1.5 2 2.5 3 3.5 3.5 4.5], d=3)
- Q_1 = diag(1e-5,1,1e-12,1e-12), Q_2 = diag(1e-12,. . .,1e-12)
- $R_1 = 1$, $R_2 = 1$
- L = 5000
- d = 3
- J = 10
- $var(\epsilon) = .01$
- $x_0 = 0$
- P_0=diag(100,. . .,10,10)

Simulation 5: Curve Registration (Figure 6)

- $\tilde{y}(t) = sin(2\pi t) + sin(4\pi t) + sin(6\pi t)$
- $\phi_i(t) = (1/\beta_i) log\left[t(e^{\beta_i} - 1) + 1\right]$, $\beta_i = -2, -1, -.5, .5, 1, 2$
- Q_2=diag(1e-12,. . .,1e-12)
- R_2=1
- L = 1000
- $var(\epsilon) = .05$
- d = 3
- J = 10

- $x_0 = 0$
- P_0=diag(100,. . .,10,10)

Simulation 6: ECG Simulation 1 (Figure 9)

- $\tilde{y}(t)$ given by [16] with default parameters and amplitude parameter a_i=[.5 -1 30 -5 0.5]
- $\phi_i(t)$ = Bspline(K, x = [-1.5 0 1.2 3.1 4.2 5.9 7 8.3 10.1 11], d=3), K uniformly spaced with K(4)=0 and K(J+1)=10
- Q_1=diag(1e-12,1e-1,1e-12,1e-12), Q_2=diag(1e-12,. . .,1e-12)
- $R_1 = 1$, $R_2 = 1$
- L = 5000
- $var(\epsilon) = .05$
- d = 3
- J = 30
- $x_0 = 0$
- P_0=diag(100,. . .,1,1)
- Derivative constraint: lb=.5, ub=5

Simulation 7: ECG Simulation 2 / Refinement Method (Figure 10)

– For the true template, we simulate an ECG as in [16] with default parameters and amplitude parameters given by $a_i = [.05 - .13 - .1 .05]$.

Initial Run-through

- $\tilde{y}(t) = sin(\pi t)^{500}$

- $\phi(t) = 1.5t + .04sin(\pi t)$
- Q_2=diag(1e-12,. . .,1e-12)
- $R_2 = 30$
- L = 5000
- $var(\epsilon) = .05$
- d = 3
- J = 15
- $x_0 = 0$
- P_0=diag(100,. . .,1,1)
- Derivative constraint: lb=1, ub=5

Second Run-through

- $\tilde{y}(t)$=Bspline(K, x=[-0.0217, -0.0209, -0.0218, -0.0159, -0.0059, 0.0167, -0.0612, 0.1589, -0.0883, 0.0149, -0.0023, 0.0431, -0.0058, -0.0226, -0.0215])
- $\phi(t) = 1.5t + .04sin(\pi t)$
- Q_2=diag(1e-12,. . .,1e-12)
- $R_2 = 1$
- L = 5000
- $var(\epsilon) = .05$
- d = 3
- J = 20
- $x_0 = 0$

- P_0=diag(100,. . .,1e-12,10)
- Derivative constraint: lb=1, ub=5

Simulation 8: ECG Simulation showing effect of constraint (Figure 11)

- $\tilde{y}(t)$ given by [16] with default parameters and amplitude parameter a_i=[.5 -1 30 -5 0.5]
- $\phi(t) = 1.5t + .1t^2$
- Q_2=diag(1e-12,. . .,1e-12)
- $R_2 = 1$
- L = 5000
- $var(\epsilon) = .05$
- d = 3
- J = 10
- $x_0 = 0$
- P_0=diag(100,. . .,1e9,1e9)
- Derivative constraint (for constrained estimate): lb=1