5-7-2019

Parameter Estimation of a Cardiac Model Using the Local Ensemble Transform Kalman Filter

Nathan Holt
nxh7119@rit.edu

Parameter Estimation of a Cardiac Model Using the Local Ensemble Transform Kalman Filter

by

Nathan Holt

A Thesis Submitted in Partial Fulfillment of the Requirements
for the Degree of Master of Science in Applied and Computational Mathematics
School of Mathematical Sciences, College of Science

Rochester Institute of Technology
Rochester, NY
May 7, 2019

Committee Approval:

R·I·T
College of
Science

Dr. Elizabeth Cherry Date

School of Mathematical Sciences

Thesis Advisor

Dr. Matthew Hoffman Date

School of Mathematical Sciences

Committee Member

Dr. Laura Muñoz Date

School of Mathematical Sciences

Committee Member

Dr. Matthew Hoffman Date

School of Mathematical Sciences

Graduate Program Director

Abstract

Cardiac arrhythmias are irregularities in the electrical activity in the heart; the electrical impulses in the heart become chaotic or disorganized, which can cause a possibly lethal problem to the contraction of the heart. In order to understand the dynamics of arrhythmias and to be able to predict and treat them, numerical models have been developed to capture the dynamics of the electrical impulses in the heart. In a clinical setting, optical mapping technologies — using cameras and voltage-sensitive dyes to capture the electrical impulses propagating across the heart — have been used to capture the dynamics of the electrical activity along the surface of the heart with high spatial and temporal resolution. Despite the high resolution provided by the optical mapping technologies along the surface of the heart, the techniques are unable to capture measurements of the voltage in the interior of the heart. Kalman Filters attempt to solve this problem by combining experimental data — that which is obtained by direct measurement, such as by optical mapping methods — with a mathematical model. This has been shown using synthetic data to be an effective method of reconstructing the dynamics of certain cardiac arrhythmias in tissue. It is desirable to be able to obtain the values of the parameters that guide the dynamical behavior of the cardiac arrhythmias in a given mathematical model. Knowledge of the values of the model parameters can be used to retroactively explain why dynamical effects occurred or to predict future behavior of the electrical impulses propagating throughout the cardiac tissue. In this thesis, we utilize a state-augmentation method of estimating model parameters using a nonlinear extension of the general Kalman Filter. We use a three-variable model of the cardiac action potential in conjunction with the Local Ensemble Transform Kalman Filter (LETKF) in order to estimate the state of the electrical impulses traveling along cardiac tissue. We show the viability of the state-augmentation methods of parameter estimation with the LETKF and determine three criteria that can be used to explain the effectiveness of the parameter estimation algorithm. We first establish the results by estimating a single parameter, and then expand our results by showing the same criteria hold when estimating multiple model parameters simultaneously. The results provide evidence that this method of parameter estimation is useful for cardiac models — both by a good estimation of the state and the predictable estimation of the model parameters — and suggest additional avenues of research for the preliminary work presented in this thesis.

Acknowledgments

I would first like to thank my thesis adviser Dr. Elizabeth Cherry for her continual guidance in my research and her fastidious feedback on all aspects of this document. I would also like to thank my thesis committee for having helped guide me along through this project and having taught courses which have directly impacted the work I was able to accomplish in this document. I would also like to thank my friends and family members for supporting my journey through school, and for their willingness to proofread this document for me. Lastly, I would like to thank the National Science Foundation for funding this project through NSF grants CMMI-1234235 and CNS-1446312.

List of Figures

Contents

I. Introduction

In the United States, cardiovascular disease is the leading cause of death, claiming the lives of over 801,000 people in 2017, accounting for almost one third of all deaths in the country [5]. Heart disease has also been steadily increasing in prevalence over the past decade, causing it to be an ever-increasing threat. Despite advancements in the care and prevention of heart disease, 23.6 million people in the world are expected to die annually from heart disease by 2030 [5]. Cardiac events are often difficult to predict due to their often asymptomatic nature; the person often shows no symptoms until suffering the cardiac event itself, which can prove fatal. These cardiovascular complications can occur when the electrical signals propagating throughout the heart become unorganized or chaotic, resulting in irregular heart contractions, or arrhythmias.

Cardiac arrhythmias are irregular beatings of the heart, such as tachycardia (an abnormally fast heart rate) or bradycardia (an abnormally slow heart rate). Arrhythmias are often the precursors to more serious cardiac complications, such as heart failure, stroke, or cardiac arrest [2].

There has been strong evidence to suggest that potentially lethal arrhythmias could be caused by deformations or instabilities in the cardiac action potential [6, 17], which is a physical occurrence in a cell in which the membrane potential spikes in a process called depolarization, followed by a return to a resting potential in a process called repolarization. This entire process of depolarization and repolarization happens on the time scale of a few hundred milliseconds. These cardiac cells are excitable media, and as such they have a threshold for excitation and an associated refractory period to allow time between excitations in a single cell.

Due to the biological knowledge that the electrical signals propagating throughout the heart act as the coordinator of the heart, telling it to beat together and evenly, it is important to understand exactly how the electrical impulses travel throughout the heart. In a laboratory setting, an electrode mesh can be placed about the surface of the heart in order to take recordings of the electrical potential at each electrode. However, this approach poses two problems: the spatial and temporal resolution limitations constrain how good of a "snapshot" we get of the surface of the heart at a given time, and there are no measurements of the electrical potential in the interior of the heart.

The process of using voltage-sensitive dyes to measure the membrane potential of the heart is the more modern approach to making experimental data about the voltage across the surface of the heart. This has been shown to provide a good approximation to the true voltage using optical

mapping technology, but it does not provide a complete snapshot, as data is subject to noise [12]. Voltage-sensitive dyes have mostly been limited to the surface, and close to the surface of the heart, leaving much of the interior of the heart unobservable. As such, we do not have a precise method to take exact measurements of the voltage in the interior of the heart in the same way that we do about the surface of the heart. This lack of data makes prediction of cardiac events even more difficult, as the surface voltage may look healthy, but the interior voltage may be becoming chaotic or otherwise disorganized, which can lead to fatal cardiac events. There have been experiments that have shown the interior and exterior surfaces of the heart to have very dissimilar dynamics, which indicates the potential complex intramural dynamics occurring in the heart and solidifies the needs for more information about the interior of the heart in order to gain a complete picture of the dynamics occurring in the heart [10].

The lack of complete observational information about the system leads to a problem of attempting to estimate the entire state of the heart given sparse and noisy data (due to the constraints of the method of experimental observation, whether it is an electrode mesh or a more modern method such as optical mapping [3]). The problem of state estimation has in general been solved using a Kalman Filter, with well-known examples coming from the fields of weather forecasting and signal processing. In the context of cardiac modeling, we will be using a specific nonlinear extension of the Kalman Filter, the Local Ensemble Transform Kalman Filter (LETKF) [15], to solve the state estimation problem. This has shown to be a good method of performing the state estimation in the cardiac model we will be using, the Fenton-Karma model [14, 20].

In this thesis, we extend the usage of the LETKF by simultaneously estimating the state of the system and estimating a set of parameters of the Fenton-Karma model. The goal of the simultaneous estimation of both the state and the parameters is to reconstruct the state of the system given only sparse, noisy data and to recover the values of the parameters associated with the true state of the system. We will estimate a single model that is spatially and temporally homogeneous (unchanging across the time of estimation and having a uniform value throughout the entire system). From the estimation of a single parameter, we will analyze patterns in the estimation of the parameter to determine the factors that influence the success of the algorithm, thereby allowing us to predict whether a parameter is likely to be estimated well and if there are any methods to improve the chance of successful estimation. We will then estimate multiple parameters simultaneously in conjunction with the state of the system and show that the same rules govern the success of the algorithm when estimating multiple parameters as they do when estimating a single parameter at a time. This allows us to simplify the analysis in general and look

at the case of the estimation of a single parameter, and to have reasonable assurance that the same results apply when estimating multiple parameters simultaneously.

The remainder of this thesis is organized as follows. In Section II we provide an overview of the mathematical model we are using in conjunction with the LETKF and to generate our synthetic data, including information about the numerical implementation of the model. We also give an introduction to Kalman Filters and the LETKF specifically, and provide technical details about the choices made for tuning parameters and initialization methods of the Local Ensemble Transform Kalman Filter. We also provide a table of common variables used throughout the thesis for reference. In Section III we show the results of the parameter and state estimation with subsections to group related results together. Here we illustrate the three factors we have identified as important for the success of the parameter estimation and we demonstrate that these results also hold in the case of the simultaneous estimation of multiple parameters. Finally, in Section IV we conclude our thesis with a discussion of the insights gained into the parameter estimation problem using our methods and a discussion about future work that could be done to continue this research.

II. Methods

This section is primarily broken down into two parts: the Fenton-Karma model and the Local Ensemble Transform Kalman Filter (LETKF). In the first half of this section, we discuss the history and definition for the Fenton-Karma model, providing information about the motivation for its application in our project, and providing information about the numerical implementation of said model. In the second half of this section, we discuss the LETKF and discuss how we incorporate the Fenton-Karma model in this method of data assimilation to solve the state estimation problem in cardiac tissue.

II.1 Overview of the Fenton-Karma Cardiac Model

II.1.1 History and Definition

The development of the Fenton-Karma model in 1998 was propelled by the desire to model the complex electrical dynamics that occur along the ventricular cell membrane in as simple a way as possible. The goal of the model was to be able to quantitatively reproduce the restitution curves produced from other complex cardiac models and experimental data by using the least amount of ionic channel complexity as possible [9].

The model was inspired in part by the earlier Fitzhugh-Nagumo model. This model is a simplified "caricature" of a general excitable media that can be given the interpretation of a cardiac cell, and can capture the general behavior of the membrane potential. As a tissue-level representation of the cardiac dynamics, it is described by the equations

$$\begin{aligned} \partial_t u &= \nabla \cdot (D \nabla u) + 3u - u^3 - v, \\ \partial_t v &= \varepsilon(u - \delta), \end{aligned} \tag{1}$$

where u is a dimensionless measure of the membrane potential, D is a diffusion coefficient, v is representative of the membrane conductance, ε is a small parameter that measures the abruptness of excitation, and δ is a parameter that measures the excitability of the membrane [9]. The Fitzhugh-Nagumo model continues to be used to study the basic wave behavior of excitable systems, especially based upon its dependence of the simplified parameter set of $\{\varepsilon, \delta\}$. This model has also been used to establish a "baseline" of generic excitable media models and can

be used as a benchmark against other more complicated models to study the effects of more complicated parameter sets being introduced into the mathematical models[26].

However, the Fitzhugh-Nagumo model has been shown to be unable to capture the restitution properties and exact shape of ventricular action potentials [8]. As a result of this shortcoming, subsequent Fitzhugh-Nagumo-inspired models have added complexity in order to be able to capture more realistic dynamics which have been observed experimentally [18]. Other models, such as the 1977 Beeler-Reuter model, applied new-found knowledge of the calcium channel in ventricular cells to make a ventricular cell-specific model in the spirit of the Hodgkin-Huxley model [13], in contrast to a generic excitable media model, as was the case in the Fitzhugh-Nagumo model [4].

The Fenton-Karma model [9] attempts to capture more domain-specific knowledge of cardiac cells, such as the ion channels that contribute to the transmission of voltage flowing through the cell membrane. In doing this, the Fenton-Karma model makes the following simplification:

$$I_{ion} = I_{fi}(V; v) + I_{so}(V) + I_{si}(V; w), \tag{2}$$

where I_{ion} is the total current flowing through the cell membrane, I_{fi} is the fast inward current responsible for depolarization of the membrane, I_{so} is a slow outward current responsible for the re-polarization of the cell membrane, and I_{si} is a slow inward current which helps to balance I_{so} during the plateau phase of the action potential [9]. The currents I_{fi}, I_{so}, and I_{si} correspond to the sodium, potassium, and calcium currents, respectively. This is an oversimplification, as other studies have shown the electrical dynamics across the membrane are much more complex than can be captured by three individual currents [22]. Nevertheless, the minimal ionic channel complexity present in the Fenton-Karma model has still proven able to capture many complex cardiac dynamics, including the ability to generate alternans and recreate restitution curves observed experimentally.

The purpose of the Fenton-Karma model is to only keep the minimal ionic complexity in order to capture the restitution properties of the cardiac cells. As such, this model is unable to represent specific measurements of the trans-membrane currents and cannot claim to have a complete snapshot of the underlying biological mechanisms of the currents. However, this model is able to capture the restitution properties of the ventricular cells quantitatively. As such, the parameters guiding the Fenton-Karma model can be found by curve-fitting of the experimentally-observed restitution curves, in contrast to measurements of the currents across specific ion channels.

With this simplification of ionic current, the Fenton-Karma model can be described by the equations

$$\begin{aligned}
\partial_t u &= \nabla \cdot (D\nabla u) - J_{fi}(u;v) - J_{so}(u) - J_{si}(u;w), \\
\partial_t v &= \frac{(1-\mathrm{H}(u-u_c))(1-v)}{\tau_v^-(u)} - \frac{\mathrm{H}(u-u_c)v}{\tau_v^+}, \\
\partial_t w &= \frac{\mathrm{H}(u_c-u)(1-w)}{\tau_w^-} - \frac{\mathrm{H}(u-u_c)w}{\tau_w^+},
\end{aligned} \tag{3}$$

with the three currents given by

$$\begin{aligned}
J_{fi}(u;v) &= -\frac{v}{\tau_d}\,\mathrm{H}(u-u_c)(1-u)(u-u_c), \\
J_{so}(u) &= \frac{u}{\tau_o}(1-\mathrm{H}(u-u_c)) + \frac{1}{\tau_r}\,\mathrm{H}(u-u_c), \\
J_{si}(u;w) &= -\frac{w}{2\tau_{si}}(1+\tanh\left[k(u-u_c^{si})\right]).
\end{aligned} \tag{4}$$

In these equations, $\mathrm{H}(x)$ denotes the standard Heaviside step function defined by

$$\mathrm{H}(x) = \begin{cases} 0 & x < 0 \\ 1 & x \geq 0 \end{cases}. \tag{5}$$

The variable u represents a dimensionless membrane potential, and both variables v and w are gate variables responsible for the inactivation and reactivation of I_{fi} and I_{si} respectively. The various other parameters of the model consist of the time constants $\tau_d, \tau_o, \tau_v^+, \tau_v^-, \tau_w^+, \tau_w^-$; the threshold potentials u_c, u_c^{si}; and u_v, and the constant k. As previously mentioned, as the Fenton-Karma model aims to reproduce quantitative features of ventricular action potentials, the parameters are chosen to fit specific restitution curves in order to give realistic output.

For our purposes, we have selected parameter values that fit the restitution curves of the original Beeler-Reuter model with standard parameter values [4]. These selected parameter values can be found in the table in the original Fenton-Karma model paper [9] and are provided for reference in Table 2.

II.1.2 Numerical Implementation

While a main motivating factor of employing the methods discussed in this paper relies upon the problem being three-dimensional, and thus having unobservable information in the interior

of the tissue, we will implement these methods in a one-dimensional setting. This is done due to the complexity of the problem drastically increasing in higher dimensions and because of the usefulness of establishing results and interpretations in a one-dimensional setting before moving to a higher-dimensional setting.

We will be using a standard Forward Euler discretization of the Fenton-Karma model to propagate the solution forward in time. As such, we will provide an initial condition of the initial profile of the variables u, v, and w in the one-dimensional setting. We will use a ring geometry, in which we have a cable with periodic boundary conditions, in order to allow the generated action potential to run indefinitely without intervention.

We use a spatial resolution of 0.25 mm and a temporal resolution of 0.05 ms.

II.2 Overview of Kalman Filters

The goal of a Kalman Filter is to produce an improved estimate of a system's state by combining observational data with a numerical model's predictions. Of course, external factors such as noisy observational data or error in the model can limit the success of the forecasting ability. A Kalman Filter seeks to combine a mathematical model and observational data together in a way that is better than either separately, taking into account possible model error or noisy observational input [16].

For our purposes, we use an extension of the general Kalman Filter, the Local Ensemble Transform Kalman Filter [15], in order to improve our estimate of the voltage across the membrane in our one-dimensional ring geometry. We will use the Fenton-Karma model as the numerical prediction model in the Kalman Filter. For the present study, we will create our own synthetic data (with some added Gaussian noise) in order to recreate what realistic observational data would be like, which will allow us to assess the accuracy of our approach due to having access to the true state of our system for all time points.

II.3 The Local Ensemble Transform Kalman Filter

II.3.1 Mathematical Formulation

We will first begin by describing the mathematical formulation of the Local Transform Kalman Filter (LETKF). Then we will discuss the process in a physical, heuristic sense, in terms of the

estimation of the membrane potential u in the Fenton-Karma model.

The LETKF [15] is a particular type of ensemble Kalman Filter which is a nonlinear extension to the general Kalman Filter. Ensemble Kalman Filters such as the LETKF attempt to minimize error variance by using a small number of model states to characterize the covariance of the prior forecast, the estimate of the state of the system including information about observational data points from the previous timestep. This collection of model states is referred to as the ensemble. The number of ensemble members is typically small (which is a big factor in why ensemble Kalman Filters are used in large, complex problems), usually between 10 and 100 [14]. We will denote the number of ensemble states as k.

The process of the LETKF is largely split into two separate halves: the background and the analysis. The background can be thought of as the prior forecast of the system derived purely from the numerical prediction model (in statistics, this would be called the "prior"). The analysis combines observations with the background state estimate (in statistics, this would be called the "posterior").

Let $\{x^{b(i)} : i = 1, 2, \cdots, k\}$ be the background ensemble (the set of all k background ensemble members). Each background member $x^{b(i)}$ is a model forecast from the previous timestep and contains information from previous observations (if this is not the first time step with observation points) and information from the model's propagation forward in time. We assume that the background has an unbiased Gaussian error with mean 0 and variance $(0.05)^2$, and we establish our initial guess of the background state to be the average of the background ensemble,

$$x^b = \frac{1}{k} \sum_{i=1}^{k} x^{b(i)}, \tag{6}$$

and the background covariance is the sample covariance of the ensemble,

$$P^b = \frac{1}{k-1} \sum_{i=1}^{k} \left(x^{b(i)} - x^b \right) \left(x^{b(i)} - x^b \right)^T. \tag{7}$$

With this information, we now have an estimate of the background at the current timestep, x^b, and have successfully parameterized the covariance of the ensemble, P^b. It is at this point that we transition from the background to the analysis step. In doing this, we introduce the H operator which is an operator from the model space to the observation space. If we label the true state of the system at the current time step as x^t, we have

$$y^o = H(x^t) + \varepsilon, \tag{8}$$

where y^o are the observations for the current timestep, and ε is the unbiased Gaussian noise we assume our methods of observation induce. The values for the Gaussian noise of the observations is chosen from the normal distribution with mean 0 and variance $(0.05)^2$, $\mathcal{N}(0, 0.05^2)$, in order to approximate the amount of noise induced by using optical mapping techniques. It is worth noting that H could be a nonlinear map, although in this paper, we will use a linear map for H.

We are using the LETKF, which means that we specify a localization "radius" and that the analysis step for a given point in the model space only includes observational information from points within that localization radius. This can decrease the computational cost in large models and also helps reduce erroneous or unwanted correlations of data across large distances [14]. Numerically, we achieve this in a simplistic method by giving observations within the specified radius a weight of 1 and a weight of 0 to points outside this radius.

All Kalman Filters, like the LETKF, will then combine these observation points with the background estimate in such a way that a specific cost function is minimized. From this analysis step, we can then characterize the analysis covariance in much the same way that we characterized it for the background covariance. The analysis ensemble is then constructed by adding the perturbations predicted by the estimated analysis covariance to this analysis estimate. Then, the new analysis ensemble is used as the initial state in order to propagate the model forward in time. More details of this process are available in detail in previous papers [15, 7].

From here, the model is run forward in time. The LETKF is an iterative process, and this resulting analysis ensemble will become the background ensemble in the next iteration step. As such, the background ensemble at this new timestep will encode previous information about the dynamics of the model and information about previous observations.

Outside of mathematical formalism, it helps to put this in practical, tangible terms that relate directly to the application domain of cardiac electrophysiology. We initialize the model by a set of potential model states — which we will discuss in further detail in section II.3.2 — , and we are interested in assimilating observations of the u variable, a measure of the membrane potential along the tissue. We are interested in seeing how accurately we can estimate the membrane potential with a non-perfect model, and with noisy observational data. The observational data points that we use are synthetic, generated by taking the true state of the system (which we obtain by running the model) and adding Gaussian noise to specific points we want to include in the

analysis step of assimilation.

From this, we have both the "truth" and the observational data points (as the observational data points are created by sampling the truth and adding Gaussian noise). The truth will be used to benchmark against the estimate of the system over time, in order to gauge the accuracy of the estimation of the state and the parameters over the course of assimilation. Ideally, even if the initial estimate is poor, the LETKF should be able to reduce error between the truth and the estimate over time by incorporating observational data which we provide. The goal is that even though the observational points are noisy, they should — hopefully — be able to provide enough information over many iteration cycles to guide the estimate back to the true state of the system.

A concrete example of the usefulness of this algorithm could be imagined in the three-dimensional setting, in which we have a heart and we use optical mapping techniques to make measurements of the voltage along the surface of the heart. We can make measurements of the membrane potential at a specified spatial resolution at fixed time intervals, with some observational error induced from non-perfect sensors. We want to be able to "guess" some possible initial states of the heart to form the background ensemble and then use the LETKF to incorporate the observational data. If the LETKF works properly, this means that the ensemble members should converge over time to a faithful representation of the true system state.

By using a model to generate "truth" data and synthetic observations, we are able to benchmark how well the LETKF algorithm performs over time. In a laboratory setting, the "true" state of the system is not known and, as such, it is rather benchmarked by how well it converges to something that qualitatively and (to some extent) quantitatively captures the future dynamics of the system. In this sense, the test of accuracy is how well the newfound system estimate can capture the future behavior of the true system in the future. As we are not working in a laboratory setting and have access to the true state of the system, we can instead test accuracy directly by checking convergence of the state estimation to the true state of the system. Since we generate the truth and observational data, we will be able to measure accuracy at each timestep and to make conclusions that can be extrapolated to real-world information, where the "true" state of the system is not known at the outset.

This thesis is based upon the seminal work of data assimilation on cardiac electrical wave dynamics, and additional explanation of the LETKF process with respect to cardiac electrodynamics can be found there [14, 20].

For reference purposes, Table 1 provides a table of common variables used throughout the

thesis.

Variable	Meaning
$x^{b(i)}$	The i^{th} member of the background ensemble
x^b	The average of the background ensemble
x^t	The true state of the system at time t
k	The number of ensemble members
P^b	The sample covariance of the background ensemble, the background covariance
X^b	The perturbations of the background ensemble members
y^o	The observational data points of the state
H	An operator from the model space to the observation space
m	Dimension of the state space
$\mathbb{P}^i$	The collection of parameter states of interest of the i^{th} ensemble member

Table 1: Common variables used throughout the thesis.

II.3.2 Initialization of the Ensemble

How the ensemble is initialized can play a vital role in the success of the algorithm. In the previous section, we stated that we take several "guesses" of what the system may look like - be it from propagation of a model or information about the application area - and use those as the initial ensemble members.

In our case, we choose to run the Fenton-Karma model for a "spinup" of 1000ms. We then run the model for an additional 100 iterations - with each iteration corresponding to 0.05ms - and then sample our k initial ensemble members from these last 100 iterations. This idea behind this method of initialization was first proposed in the seminal LETKF paper by Hunt et al [15].

Running the model for 1000 ms will allow any transient effects such as sharp wavefronts to settle down into long-term behavior expected with periodic boundary conditions. For our initial guess of parameters to generate the ensemble (and the truth data) we use the parameters to match the restitution curves of the Beeler-Reuter model [4]. We then randomly sample the next 100 iterations to obtain the k ensemble members ($k = 20$). Sampling the iterations randomly has been shown to be a reasonable choice to characterize the initial ensemble and covariance in the weather forecasting community [15]. This method has also been suggested as a suitable method when no prior observations about the system are available, which is exactly our case, as we are initializing

the ensemble with no observational information available yet.

The creation of the initial ensemble by allowing for a "spinup" and then a random sampling has been shown to capture the covariance of the initial ensemble sufficiently so as to prevent ensemble collapse: when the members of the ensemble become too alike, potentially causing estimation problems and skewing the covariance [21]. While ensemble collapse can still happen after many iterations of the LETKF, as the ensemble members may gradually converge to a single shape, this method of selecting initial ensemble members provides enough spread to initially characterize the system without the possibility of early ensemble collapse.

II.3.3 Ensemble Undersampling

The LETKF performs the analysis step (the combination of the forecast with the observational data for the current iteration) in ensemble space. That is, the analysis step finds an optimal solution (minimizes a cost function) that is a linear combination of the ensemble members. This usage of the ensemble members to estimate the covariance and update the state estimate is a significant computational savings when the number of ensemble members is less than the dimensionality of the state space, as it reduces the problem from an $m \times m$ problem into a $k \times k$ problem, where m is the dimensionality of the model space and k is the number of ensemble members. In our case, this is a reduction by two orders of magnitude (a reduction from 10^3 to 10^1).

However, in cases where the ensemble space is much smaller than the state space — which is most often the case of ensemble-based Kalman Filters — problems can arise from undersampling. The LETKF attempts to estimate the covariance of the state by using the sample covariance obtained from the ensemble members; it also attempts to form the analysis as a linear combination of the ensemble members. This leads to the problem of undersampling of the ensemble variance and filter divergence, an underestimation of the forecast error that causes the algorithm to ignore the observations, making it difficult to regain the spread of the ensemble later [25].

The typical solution to this problem is to artificially inflate the uncertainty. To do this, both multiplicative and additive inflation methods are applied. Multiplicative inflation applies a scaling factor to the $k \times k$ covariance matrix P^b, artificially increasing the uncertainty in the forecast and helping to prevent filter divergence. Additive inflation adds a vector of perturbations to each ensemble member. These vectors correspond to the differences in model states which are 5ms apart during the spinup portion of the initialization algorithm, and are scaled by a positive scaling factor, α. Additive inflation serves to change the span of the ensemble members; since the analysis

step is performed in ensemble space and additive inflation can increase the span of the ensemble members, this can open up new possibilities for the estimation of the state.

For the purposes of this document, we use a value of $\alpha = 0.05$ for additive inflation, building upon previous results [20]. The magnitude of multiplicative inflation varies and can be found in the text near figures.

III. Parameter Estimation

III.1 Introduction

In many scenarios, it is desirable to be able to extract information about the parameters that guide a system. This information can be used to gain insight into the future behavior of the system or to explain retroactively why particular effects occurred. We will attempt to extract the value of a given parameter in the Fenton-Karma model by using the LETKF algorithm and our "truth" data.

We use a ring size of 135 mm, which corresponds to 540 grid points in our discretized model (with a spatial resolution of 0.25 mm). Table 2 provides a list of the parameter values used to generate the truth state of the system from the Fenton-Karma model. These parameter values were chosen to match the Beeler-Reuter parameter set [9]. This selection of parameter values produces alternans — an alternation between long and short action potential durations — and exhibits a stable action potential which will propagate around the one-dimensional ring without being annihilated.

Parameter	Value
τ_v^+	3.33
τ_{v1}^-	15.6
τ_{v2}^-	1250
τ_w^+	870
τ_w^-	41
τ_d	0.25
τ_o	12.5
τ_r	33.33
τ_{si}	29
k	10
u_c^{si}	0.85
u_c	0.13
u_v	0.04

Table 2: Parameter values used to generate the true state of the system, x^t. All time constants are in ms.

III.2 Methods

To implement parameter estimation in the LETKF, we will use state-augmentation methods proposed previously [24, 1]. To estimate p parameter states using k ensemble members, we append the parameter states to the background ensemble states, producing an augmented set of ensemble states. If we let $\mathbb{P}$ be the collection of parameter states we wish to estimate, the augmented ensemble takes on the following form:

$$\{(x^{b(i)}, \mathbb{P}^{b(i)}), i = 1, 2, \cdots, k\}, \tag{9}$$

where $x^{b(i)}$ is the i^{th} member of the background ensemble and $\mathbb{P}^{b(i)}$ is the set of parameter states (values) being estimated of the background ensemble member i. In this case, we have that the dimensionality of $x^{b(i)}$ is m and the dimensionality of $\mathbb{P}^{b(i)}$ is p, making $(x^{b(i)}, p^{b(i)}) \in \mathbb{R}^{m+p}$.

In the method of parameter estimation we are using, the original background ensemble is run through the LETKF algorithm as was previously described to produce the corresponding analysis ensemble. Then we update the parameters in a separate step, obtained from running the non-localized LETKF, the ETKF, on the augmented background ensemble produced previously. Using localization has been shown to provide little increase in the effectiveness of reducing parameter estimation error, so it is common to use the non-localized ETKF in order to reduce the amount of computation necessary [24, 19]. As such, the parameter estimation step is done without localization, but the rest of the forecasting of the state and propagation of the ensemble members forward in time is done using localization, as this can improve the successfulness of the estimation of the state of the system [15, 11].

After the estimated values of the parameters are updated after being run through the ETKF, they are inserted into the model for propagation forward in time. From this, the background ensemble in this next iteration is generated and encodes information about the selection of parameter values. The covariance matrix now includes information about the successfulness of the estimation of the parameters, as it was used to propagate the model forward in time, and the results can be directly compared with the true state of the system, x^t. As such, incorrect perturbations (assuming they are large enough to be non-trivial errors) will be reflected in the covariance, and the algorithm will attempt to vary the parameter more, in a search for the true value of the parameter.

The purpose of running the LETKF on the original background ensemble, and then the ETKF on the augmented background ensemble to update the parameter state, is to make available the

background covariance obtained from propagating forward the model in time, and also because localization can increase the effectiveness of forecasting. This allows the parameter update step to have more information on which to base the updated parameter values. This background covariance acts as an indicator of the "closeness" of the estimated parameter value to the true parameter value. If the state estimation is diverging from the truth, the estimation uncertainty will be reflected in the covariance matrix, and the algorithm will increase the amount it varies the parameter (and the analysis ensemble members as a whole) to compensate for uncertainty in its estimation of the state.

III.3 Estimation of One Spatially and Temporally Homogeneous Parameter

First we attempt to estimate a single parameter that is spatially and temporally homogeneous. Due to this, we have that $\mathbb{P}^{b(i)} \in \mathbb{R}$, since it is a single value for the parameter for each ensemble member. We simply append this value to our background ensemble members, $x^{b(i)}$, and run it through the process as previously described.

III.3.1 Effect of Parameter Dynamics

In this section we explore why this method of parameter estimation works better for some parameter choices than others, with our primary motivation being the system dynamics associated with the parameter's value. The next section will address the same underlying question by looking at the relative magnitudes of the parameter being estimated. The hypothesis is that the more sensitive the dynamics of the model are to changes in the parameter being estimated, the more likely the algorithm is to estimate the parameter successfully.

First we consider the estimation of τ_d, a parameter that appears in the fast inward current of the Fenton-Karma model. In Figure 1, we illustrate the dynamical differences caused by varying the parameter by 0.05 ms^{-1}.

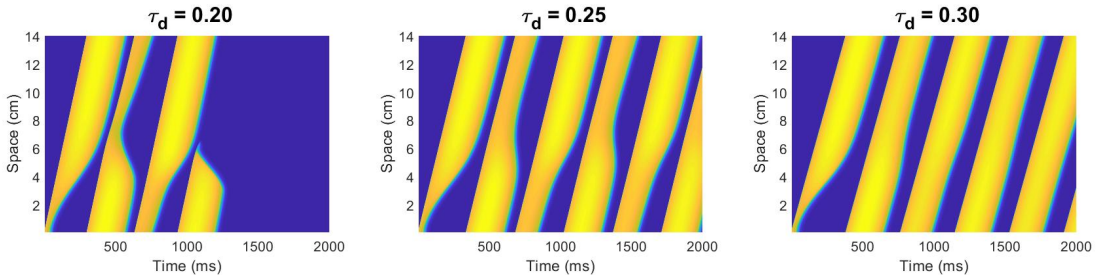

Figure 1: Illustration of the dynamical differences obtained by varying the value of τ_d by 0.05 ms^{-1}.

It can be seen from Figure 1 that changing the value of τ_d by 0.05 ms (a relative change of 20%) results in significant differences in the dynamics of the action potential propagating around the ring; including changes in the speed and wavelength alternation magnitude. With a reduction of 0.05 ms, the wave actually dies off before the simulation of 2000ms finishes. This is due to the fast inward current being increased when τ_d is decreased, causing the wave speed to increase, and stay excited for longer, causing the wave to collide with itself. This collision of the wave front with the wave back causes destructive interference, and causes the action potential to be annihilated. In the case of an increase of 0.05, the action potential stabilizes, removing any alternans in the wave shape. This can be seen in Figure 1, where the action potential no longer has any oscillations in wavelength when $\tau_d = 0.30$. This increase in the value of τ_d stabilizes the action potential by creating system dynamics in which the wave speed is slow enough and the wave dies off quickly enough to prevent any collisions of the wave front and the wave back.

For the purposes of this section, we will be using a multiplicative inflation value of 1.7 for the simulation and leave the discussion of the impact of the choice of multiplicative inflation to section III.3.3. The value of 1.7 for multiplicative inflation is an effective choice for the majority of cases, based on our numerical experiments. We also use 180 observational points, spaced uniformly throughout the ring, which is equivalent to using one of every three grid points observable for usage with the LETKF algorithm.

Figure 2 shows the results of the estimation of τ_d when the initial "guess" of the parameter to be estimated is varied. The initial value of τ_d is labeled on the left of each row as $\tau_d(0)$. The leftmost plot in each row shows the estimation of the parameter over the course of assimilation. The three other plots in each row show the root mean squared error (RMSE) of the state estimation in each of the three state variables, u, v, and w. The gray line is provided as a reference point, and represents the amount of observational error at each assimilation step. The value is held around 0.05 due to our setup of using 0.05 Gaussian observational error in our synthetic data. It is worth noting that

for our purposes, we only take observational data about the value of u, and we do not observe v or w for usage in the LETKF algorithm.

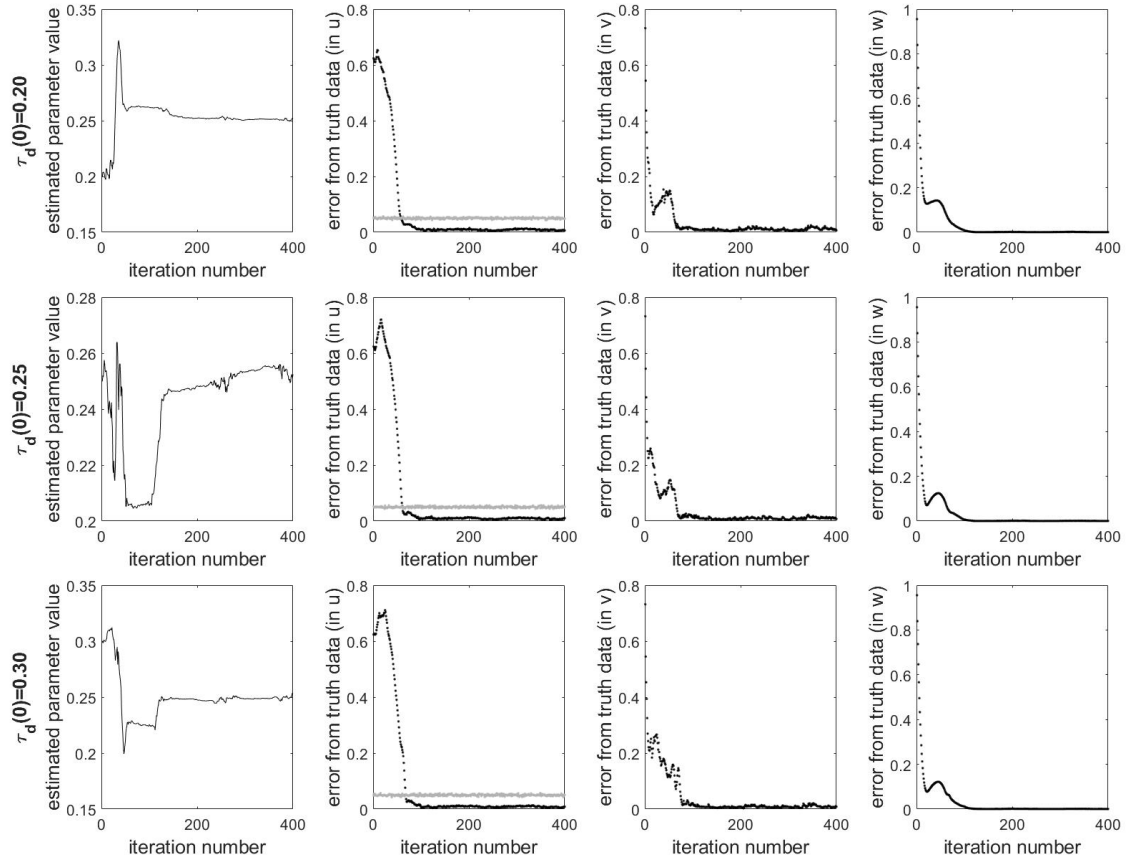

Figure 2: Illustration of the success of the estimation of τ_d regardless of initial guess. Each row corresponds to the estimation of τ_d given a different initial guess of $\tau_d(0)$. The leftmost plot in each row shows the estimation of τ_d, and the right three plots show the RMSE of the state estimation of the 3 model variables (u, v, w), with a gray observation line present in the u figure for reference. Each iteration corresponds to 5 ms.

Figure 2 shows the successful estimation of τ_d regardless of whether the initial estimate of τ_d was over- or underestimated. The error plots provided in all three cases also show the convergence of the estimation of the system state to the truth, as is desired of any Kalman Filter. As part of the way that this method of parameter estimation works, the estimated value of τ_d is put into the model in the analysis step, so it is used to create the background in this next iteration by propagating the model forward in time.

Next we consider another variable of the Fenton-Karma model, τ_{v2}^{-}. Figure 3 illustrates the

dynamical differences obtained by varying the parameter by 250 ms (a relative change of 20%).

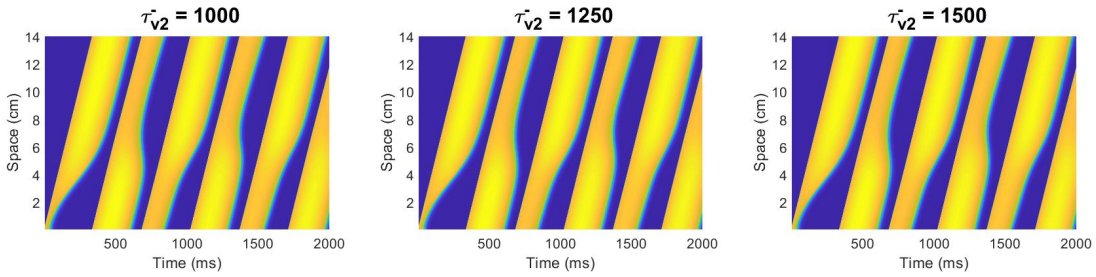

Figure 3: Illustration of the dynamical differences obtained by varying the value of τ_{v2}^{-} by 250 ms.

The dynamical differences in the change of this parameter are much more subtle. Despite using the same relative change in the parameter as for τ_d, the model is much less sensitive to this parameter, and as a result the resulting waves look identical to the unaided eye. Figure 4 illustrates the results of parameter estimation of τ_{v2}^{-} using the same method as in Figure 2.

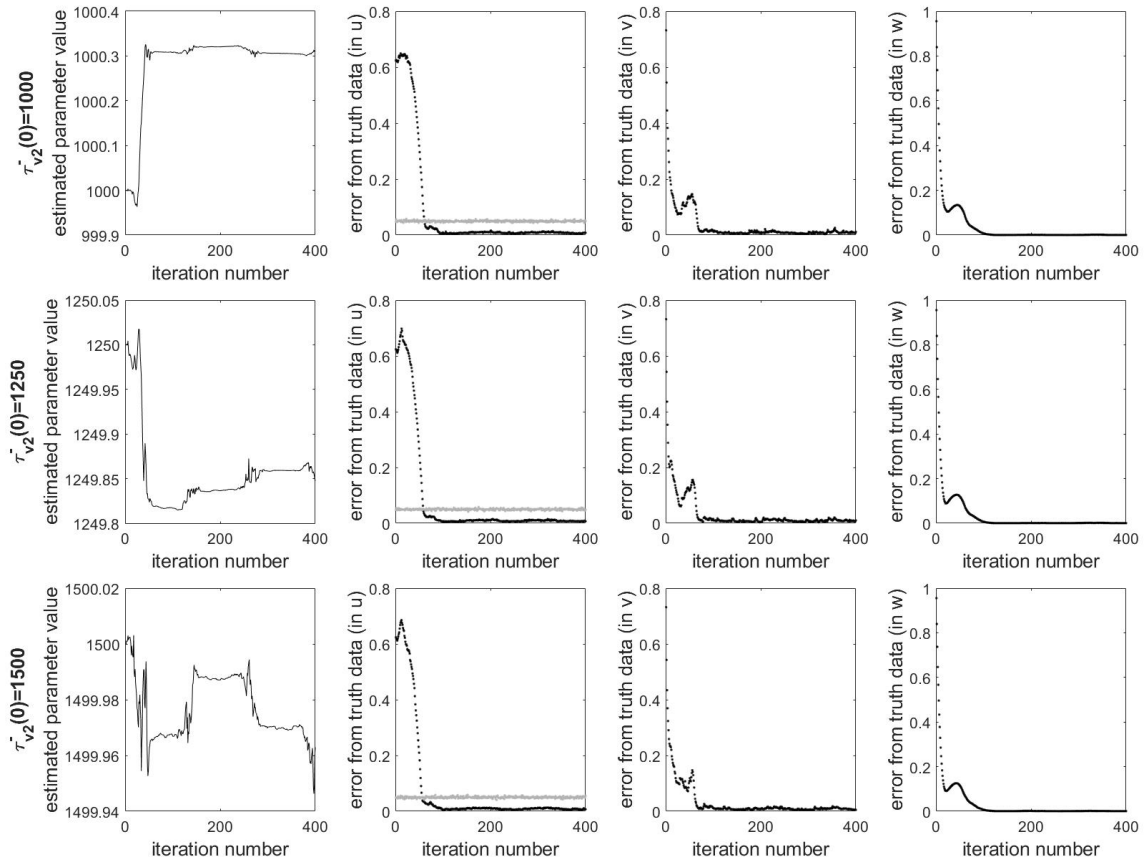

Figure 4: Illustration of the inability to estimate τ_{v2}^- regardless of initial guess. Each row corresponds to the estimation of τ_{v2}^- given a different initial guess of $\tau_{v2}^-(0)$. The leftmost plot in each row shows the estimation of τ_{v2}^-, and the right three plots show the RMSE of the state estimation of the 3 model variables, with a gray observation line present in the u figure for reference.

It can be seen in Figure 4 that τ_{v2}^- was not estimated well in the cases where the initial selection of the parameter value was not correct. In all three cases, the parameter changed by a maximum of 0.3 ms, making it unable to recover the correct value of the parameter over the course of assimilation.

The explanation for this is the insensitivity of the model to the parameter. As was shown in Figure 3, a change of 250 ms is not visible in the dynamics of the action potentials, meaning that 0.3 ms is even less noticeable. The LETKF bases its estimates of the covariance on how closely the ensemble of model states estimates the true state of the system at the current timestep. However, since the tiny perturbations of the parameter do not produce any noticeable differences in the dynamics (and thus the error) of the estimated state, the algorithm accepts whatever parameter value is

given to it and instead corrects the state estimation through the traditional LETKF process — with observational data guiding its trajectory over many iterations. As such, when tiny perturbations of the parameters do not reflect noticeable differences in the model propagation, the LETKF parameter estimation process fails to move the parameter much, causing its final estimated value to be approximately equal to the initial value given to the parameter, meaning that it was not estimated well.

It is worth noting that there are several steps in the algorithm in which stochasticity may impact the performance of the LETKF and parameter estimation; the randomness in the algorithm may be created in such a way that the algorithm is able to recover the true value of the parameter(s) being estimated even if the errors in the estimated value of the parameter(s) does not produce noticeable affects in the dynamics of the system. In the algorithm, we have randomness inherent in the selection of the initial ensemble members, the selection of additive inflation values, and the construction of observational data points. As such, despite having the same setup, different results may occur for both the effectiveness of the state estimation via the LETKF and the parameter estimation via the ETKF. We illustrate this in Figure 5.

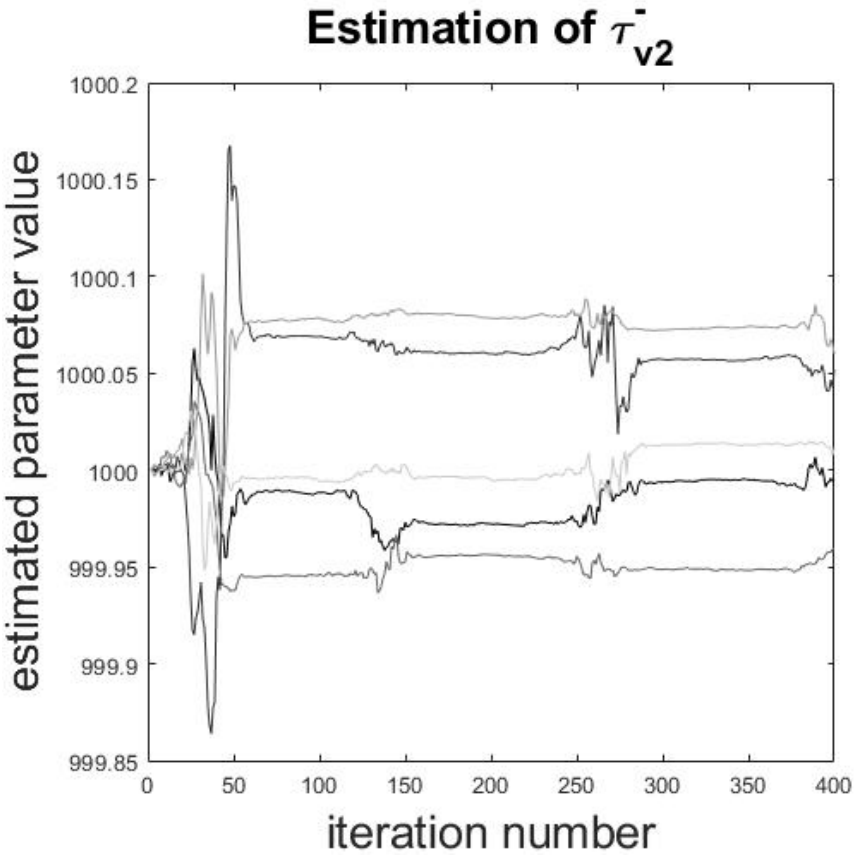

Figure 5: Illustration of the inherent randomness of the parameter estimation algorithm. Each line represents a different attempt at estimating τ_{v2}^- from the same initial conditions. The differences can be explained by pseudo-random numbers being used throughout the algorithm, including the generation of the initial ensemble members and the observational data. The true value of τ_{v2}^- is 1250, and the mean and standard deviation of the estimated value of the parameter is 1000.015 ± 0.041.

Due to the randomness present in our implementation of the algorithm, the results of both the state estimation and the parameter estimation are not deterministic. Normally, the results are very similar and differ by very little, as was the case in Figure 5. However, there are anomalous cases in which a parameter may be estimated much better in one simulation than another. We have run every simulation presented in this thesis several times and have selected cases that are largely representative of the general behavior of the algorithm. For example, there may be an anomalous case for the estimation of τ_{v2}^- in which the algorithm successfully recovers the parameter value from the same start conditions as present in Figure 5. However, we would not use this case as an example, and would instead elect to show a case such as the one shown in Figure 4.

It can be shown that parameters that affect dynamical changes over the time scale of assimilation are much more likely to be well estimated with the algorithm and will reliably recreate the actual value of the parameter, even in the presence of randomness in the algorithm. This can be seen in Figure 6, where we show the algorithm's ability to determine the true value of τ_d, despite the randomness.

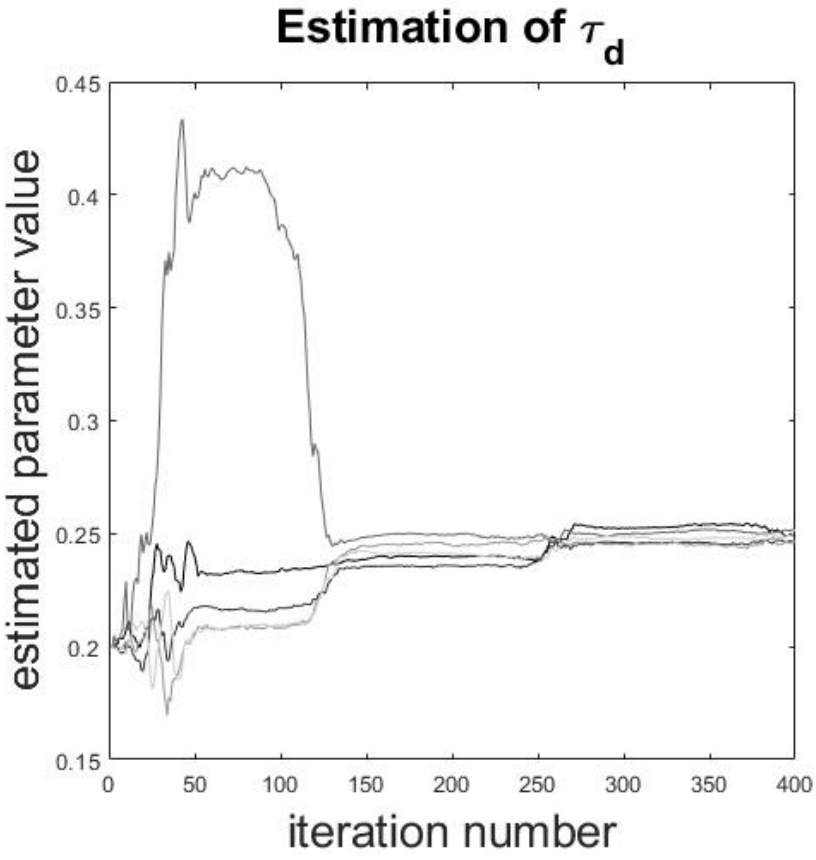

Figure 6: Illustration of the success of parameter estimation of τ_d regardless of randomness used in the algorithm. Each line represents a different attempt at estimating τ_d from the same initial conditions. All cases converge to the true value of 0.25, although the randomness affects both the convergence speed and the transient behavior of the parameter estimation before it reaches the true value. The mean and standard deviation of the estimated value of the parameter is 0.249 ± 0.003.

From the viewpoint of dynamics, the variables that produce the largest dynamical differences over the time scale of assimilation seem to be the most reliably estimated. In these cases, the error induced by incorrect estimation of a parameter at each step is much larger than for a parameter to which the model is not as sensitive. The larger, more apparent error causes the algorithm

to more readily add perturbations to the parameter value in an attempt to nudge it towards the correct parameter value. These apparent dynamical differences caused by the difference in parameter value cause the ensemble members to have a larger spread, as they all contain slightly different parameter values. This means that the perturbation matrix X^b has larger entries, since the ensemble members have greater variance from the mean.This causes the parameter to be estimated more reliably, as the error is more noticeable in the X^b perturbation matrix.

It is also worth noting that the wavefront of the action potential seems to be the strongest indicator of parameter estimation success. The model seems to be most sensitive to parameters whose value will affect the wavefront - whether that be shape, speed, location, or any other dynamical effect - and will be reliably estimated via this method. In the example of the estimation of τ_d from Figure 1, τ_d may not appear to affect the wave front, but rather the propensity for alternans and the speed of the wave. However, the wave speed will impact the location of the wavefront along the ring, or perhaps even cause the wave to collide with itself and be annihilated. This is an important part of the state to estimate for the LETKF, and so the impact of the parameter on the position of the wavefront causes the algorithm to reliably estimate τ_d, regardless of initial error in the guess.

We end this section by providing the results of the estimation of 4 sensitive parameters of the Fenton-Karma model, with initial values starting 20% below their true value. As they have strong effects on wave fronts in the model, the parameter estimation algorithm does a good job at recovering the true value of the parameters throughout the assimilation process. The results are displayed in Figure 7.

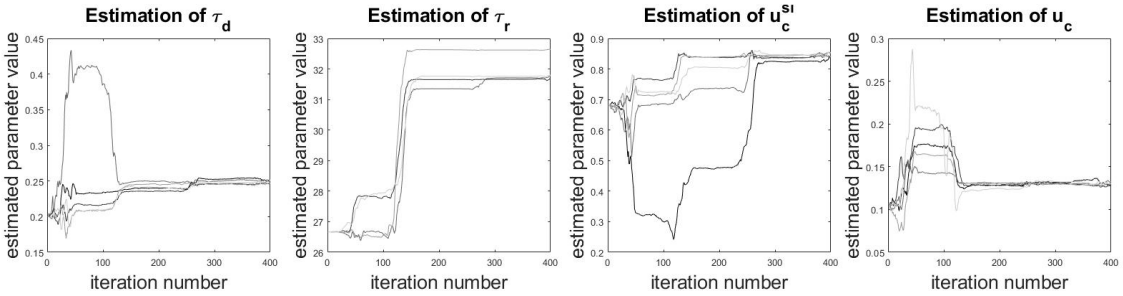

Figure 7: Illustration of the success of parameter estimation of certain parameters to which the model is dynamically sensitive. In the second panel, τ_r does not reach its true value, but it does get within a range where the dynamical differences are indistinguishable by eye. The other three cases show the parameter estimation converges to the true value, regardless of randomness in the algorithm.

This is further evidence that the dynamics induced by error in the parameter values strongly impacts the success of the parameter estimation algorithm. It also shows that these sensitive parameters can be reliably estimated, despite randomness inherent at multiple stages in each run of the algorithm.

III.3.2 Effect of Parameter Magnitude

In this section we further explore why this method of parameter estimation works better on some parameter choices than others, with the primary motivation being the importance of the magnitude of the parameter we are estimating. We also show experimental evidence that parameters are more likely to make larger relative changes of their values over the course of assimilation if their magnitudes are small in comparison to the rest of the parameters of the model. The explanation behind this phenomenon is that for large parameter values, the perturbations that get added during the algorithm are simply not large enough to notice bigger errors in the parameters. If the parameter is small, the perturbations added are relatively bigger, which means that the parameter estimation algorithm is more likely to change the value of the parameter significantly during the estimation process, in contrast to remaining stagnant, as is the case with many large, insensitive parameters, such as in Figure 5. These larger perturbations in the parameter estimation may increase the likelihood that the algorithm estimates the parameter successfully.

We first look at the case of a small parameter, u_v. The true value for u_v in the truth data is 0.04 (all variable values for truth data are provided in Table 2). We estimate from 0.32, a relative error in the parameter of 20%. We provide the dynamical differences incurred by incorrect values of u_v in Figure 8.

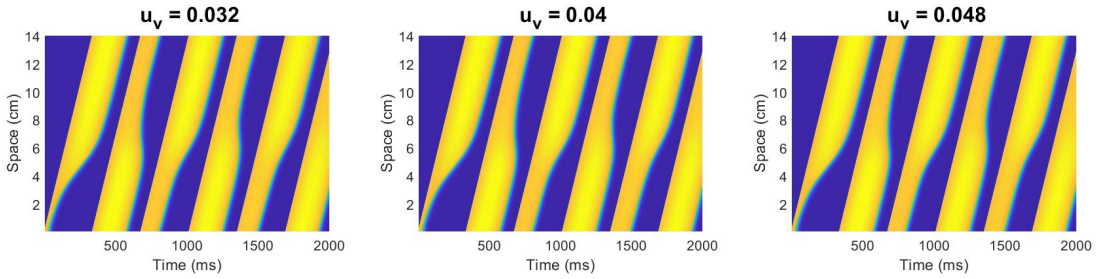

Figure 8: Illustration of the dynamical differences obtained by varying u_v.

Figure 8 illustrates that the dynamical differences caused by varying u_v are indistinguishable to the unaided eye. From this, and based upon experimental evidence shown in the previous section,

we can likely guess that this parameter will not be estimated well. We provide the results of the estimation of u_v in Figure 9.

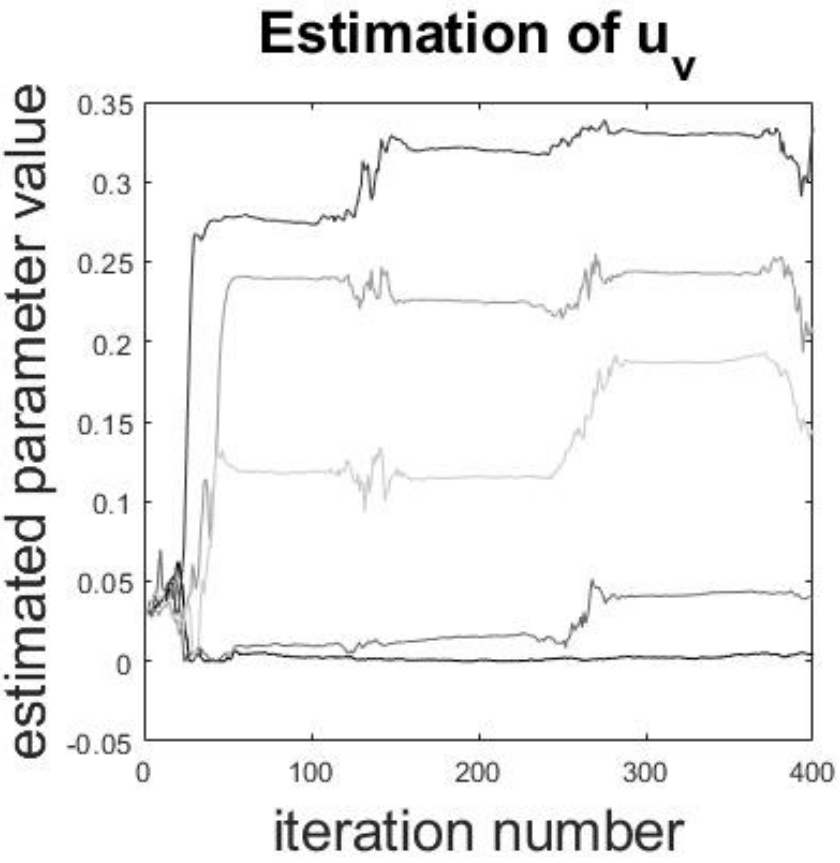

Figure 9: Illustration of the inability to estimate u_v reliably, as predicted. The multiple lines correspond to different attempts to estimate the parameter, but vary due to randomness in the algorithm. The true value of u_v is 0.04, and the mean and standard deviation of the estimated value of the parameter is 0.145 ± 0.133.

From Figure 9, we can deduce that u_v is not reliably estimated via the ETKF algorithm, as it only recovered the true value of the parameter in one run out of five attempts. However, the initial start point for the parameter was 0.032, and the parameter decreased by 0.032, and increased by 0.298, in some simulations. These correspond to relative differences of 100% and 931%, respectively.

In the estimation of large parameters using this method, the algorithm is much more likely to remain stagnant if there are not large dynamical differences - a phenomenon explored in the previous section. However, in the case of smaller parameters, such as u_v, the algorithm is more likely to change the value of the parameter by noticeable amounts in the search of the parameter.

This is due to the perturbations applied during the algorithm being large relative to the parameter value itself.

This can be seen by contrasting Figures 9 and 5. In Figure 5, τ_{v2}^{-} is being estimated from an initial value of 1000, but over the course of estimation, its value only changes by a maximum of .15, a relative change of 0.015%, despite the true value of the parameter being 1250. However, in Figure 9, as was previously mentioned, we obtain a much larger relative and absolute change in the value of the parameter.

Our hypothesis is that parameters with smaller magnitudes will in general be more prone to move larger distances from the initial value of the parameter, in an effort to find the true value of the parameter. The reasoning behind this hypothesis is that the algorithm is applying random perturbations to the parameter being estimated, and the ensemble. However, these perturbations are held at the same size, regardless of parameter being estimated. This means that the perturbations applied when estimating smaller parameters are relatively much larger than the parameters themselves, which in turn makes the error in the estimated parameter be overestimated during the algorithm. This overestimation then leads the algorithm to "self-correct" itself by varying the parameter value by relatively larger amounts in an effort to decrease the overestimated error in the parameter estimation.

These larger changes in the parameter, of course, die down after the state estimation problem has been solved by the LETKF. That is, the parameter changes will decrease in magnitude once the RMSE of the state estimation in all 3 variables has converged to low error. This does not mean that the error is no longer being overestimated or that the parameter has reached the correct value. This simply means that the state is being estimated sufficiently well, so that the Kalman Filter does not require changes to the parameter in order to maintain low RMSE in the state estimation problem; it has achieved a low enough RMSE where the Kalman Filter can continue to provide a good estimate of the state using only the numerical prediction model and the provided observational data.

We next look at another variable used in the Fenton-Karma model: k. This variable appears in the formulation of the slow inward current, J_{si}. The true value of this parameter is 10, and we estimate it from the value of 7.5. The dynamical differences caused by varying the parameter k can be seen in Figure 10.

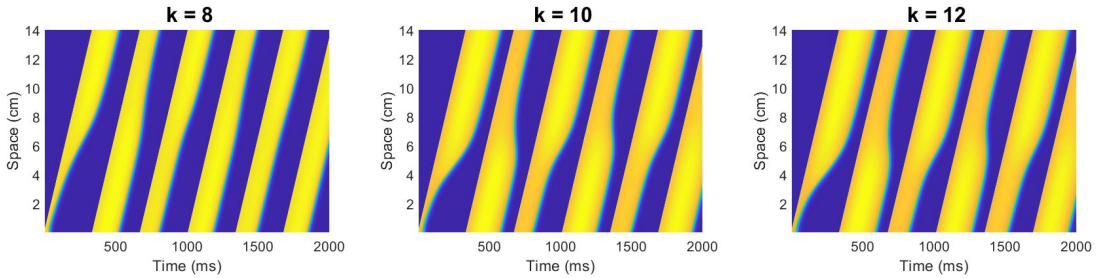

Figure 10: Illustration of the dynamical differences obtained by varying k.

Figure 10 shows that increasing k from 10 to 12 does not produce large effects in the dynamics of the wave on a macroscopic level. However, decreasing the value from 10 to 8 - as we shall do when estimating the parameter in the next figure - does produce noticeable dynamical differences. In particular, the waves tend to stabilize, removing almost all oscillations in the wavelength from the resulting space-time diagram. This removal of oscillations in the space-time diagram simply corresponds to the wave reaching a more stable position, removing most of the alternans. The results of parameter estimation from an initial value of 8 are shown in Figure 11. The inability of the algorithm to recover the true value of k is perhaps unexpected, as visible differences in the dynamics of the action potentials are present in the diagrams. The error in the value of this parameter is responsible for having an incorrect wavelength and much less pronounced alternans. However, the lack of larger differences in the morphology of the wave - in particular, differences in the wave front - prevents the algorithm from having pronounced error in the state estimation. Errors in the wavelength - when not severe - are too similar to the observations when the rest of the ring is being estimated well. Due to the overall good estimation of the state (outside of locations where the wavelength is off), the parameter will not make large changes, and is not estimated well.

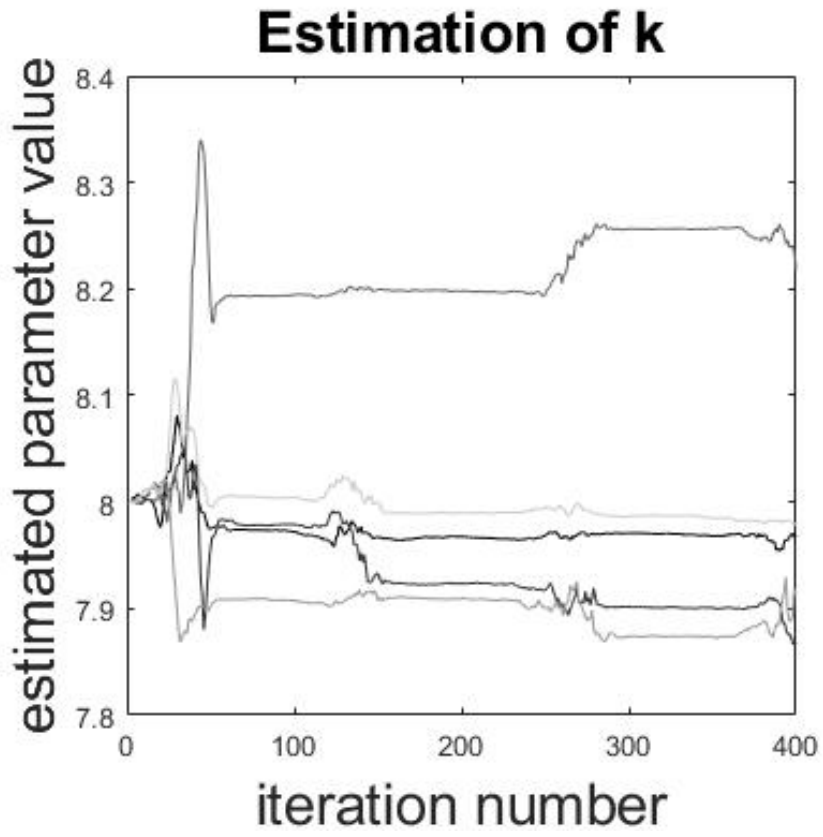

Figure 11: Illustration of the inability to estimate k. The multiple lines correspond to different attempts to estimate the parameter, but vary due to randomness in the algorithm. The mean and standard deviation of the estimated value of the parameter is 7.989 ± 0.136.

In this case, the parameter is varied by a change of up to 0.3, which is roughly the same change experienced in Figure 9. It can be seen that the final estimated parameter values often varies heavily from simulation to simulation, as was the case in Figure 9. This is the crux of the importance of parameter magnitude: the parameter value impacts how much the perturbations cause the estimated parameter value to change, which in turn leads to more divergent parameter estimation results from simulation to simulation. Experimentally, the dependence upon the magnitude of the parameters leads to more varied results in estimating low parameters and makes analysis of the effectiveness of the parameter estimation difficult, as any successful parameter estimation may simply be due to the randomness in the algorithm, much more than estimating larger parameters.

The dynamical differences in the parameter seem to take precedence in the order of importance

regarding the effectiveness of the parameter estimation. Small parameters to which the Fenton-Karma model is sensitive have been shown to converge to approximately the same value regardless of randomness. This can be seen in Figure 7, where the small parameters τ_d, u_c, and u_c^{si} are all reliably approximated despite having magnitudes less than one. However, also in Figure 7, we see that τ_r is pretty reliably estimated as well, as the model is moderately sensitive to the parameter. However, it is an order of magnitude larger than the other parameters in the figure, and we see the most variability in the results of its estimation. This variability can be accounted for in part by the magnitude being larger, and thus when it approaches the true parameter value, its perturbations grow relatively very small, and as such the algorithm will not correct the parameter value further, leaving the parameter value to remain unaffected. The magnitude of the parameter prevents it from fine-tuning to the exact value of the parameter; accuracy is achieved for the LETKF state estimation problem without obtaining a good parameter estimation.

III.3.3 Effect of Multiplicative Inflation

Multiplicative inflation is used to mitigate underestimation of the variance of the ensemble, which results from ensemble undersampling (using too few ensemble members when working with high-dimensional states) [25]. In the context of parameter estimation, this multiplicative inflation is applied to the augmented ensemble containing the parameter state(s) being estimated. In this manner, the augmented ensemble perturbation matrix X^b is inflated, meaning that the perturbations of the parameter(s) are increased. These additional perturbations serve to decrease the confidence of the ETKF in the selection of the parameter, as there is more variance of the parameter, thus making the algorithm more apt to change the value of the parameter over the course of estimation.

As such, the multiplicative inflation factor to be used is an important parameter in and of itself that can affect the successfulness of the parameter estimation. Throughout sections III.3.1 and III.3.2, we have been using a fixed value of 1.7 for multiplicative inflation, as this has been shown to be a generally good selection of inflation value. However, in this section, we will explore the impact that the multiplicative inflation value can have on the performance of both the state estimation problem of the LETKF and the parameter estimation of the ETKF.

First, we run simulations for the estimation of τ_d, as was done in Figure 6 with an inflation value of 1.7. We will now estimate the same parameter set with all of the same initial conditions, but with multiplicative inflation values of 1.2 and 1.9. Inflation values in general should be between 1

and 2 (corresponding to no inflation and doubling inflation, respectively), so these are good values to test the lower and upper limits of the inflation value. It is worth noting that a multiplicative inflation factor of 2 is not a bound on the value, but numerical problems have occurred in our particular model when increasing the value of multiplicative inflation past 2. The results are provided in Figure 12.

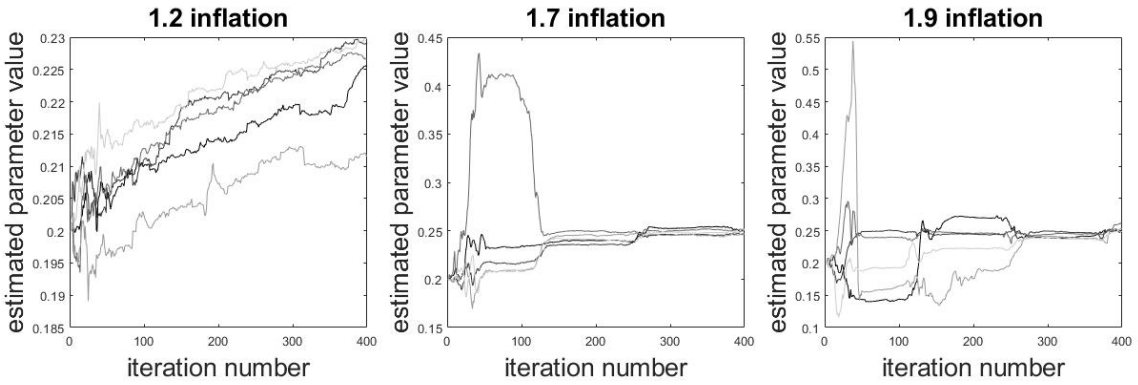

Figure 12: Illustration of the estimation of τ_d with different multiplicative inflation values. Each line represents a different attempt at estimating τ_d from the same initial conditions. The true value of τ_d is 0.25. The resulting estimated parameter values's mean and standard deviation are: 0.225 ± 0.007, 0.249 ± 0.003, and 0.252 ± 0.005, respectively. The multiplicative inflation factor of 1.7 provided the best results, both in closest mean value, and smallest standard deviation.

Figures 1 and 2 have already shown that the Fenton-Karma model is sensitive to τ_d. As such, based on section III.3.1, we would expect the parameter to be generally well estimated, regardless of randomness used throughout the algorithm in a given simulation.

Figure 12 demonstrates that the choice of multiplicative inflation can greatly influence the success of the estimation of parameters, even in the case where the model is sensitive to the parameters. In the case of a lower multiplicative inflation value of 1.2, we see that the parameter estimation does indeed correct towards the true value of 0.25, but it does not reach the true value over the assimilation interval of 400 iterations. This matches expectations: due to the lower multiplicative inflation factor, the variance in the parameter estimation is lower, so the algorithm tends to make smaller changes to the parameter value between iterations. The smaller changes to the parameter value then lead to slower convergence to the true parameter value, as the algorithm will not allow for larger "jumps" in the parameter value to more quickly reach the true value.

When the multiplicative inflation factor is 1.7, we see that the parameter estimation works

successfully in all instances, although the intermediary steps until convergence vary. This is much the same way that the estimation works in the case of a multiplicative inflation factor of 1.9. Due to the increased inflation value, the algorithm can make large adjustments to the parameter value while it is being estimated. This can be seen easily in that there are more sharp jumps in the value of the parameter estimation in the case of 1.9 multiplicative inflation. We also observe that the parameter spans a larger range (takes on a larger range of values over the course of estimation) over the course of multiple simulations when the inflation value is increased, as expected. In the case of 1.7 multiplicative inflation, the estimated parameter approximately ranges between 0.18 and 0.43, whereas it ranges between 0.12 and 0.54 when the multiplicative inflation is 1.9. Comparing the ranges the estimated parameter takes on over the course of multiple runs of algorithm, we see the increase of range when the multiplicative inflation factor is increased from 1.2 to 1.7 to 1.9 being 0.04, 0.25, and 0.42, respectively.

From Figure 12, we see that there may often be a "sweet spot" for the multiplicative inflation value for the estimation of a parameter that is not known a priori. In the case of τ_d, a parameter to which the model has been shown to be sensitive, we know that the algorithm will attempt to correct it quickly, as slow correction of the parameter leads to a poor state estimation by the LETKF. In order to recover the true value of τ_d quickly, the parameter estimation algorithm needs to have a large enough multiplicative inflation in order to perturb the value of the parameter large enough towards the correct value. But we do not want the multiplicative inflation factor to be too large, as the algorithm will tend to overshoot the true value of the parameter when the model is sensitive to that parameter. Overshooting during the estimation of the parameter is due to the parameter causing a high RMSE error in the state estimation problem, and larger changes to the value of the parameter are made during the estimation process in order to lower the RMSE. However, it is not known a priori how large of perturbations are required to add to the estimated parameter in order to obtain good estimation results, and so the algorithm often adds too much to the value of the parameter initially. So, the algorithm often overcorrects in an attempt to quickly minimize the RMSE of the state estimation. This overestimation can be minimized by a reduction of the multiplicative inflation, as can be see in the cases of 1.7 and 1.9 multiplicative inflation in Figure 12.

Figure 13 shows the same process for u_c, another parameter to which the Fenton-Karma model is sensitive (see Figure 7).

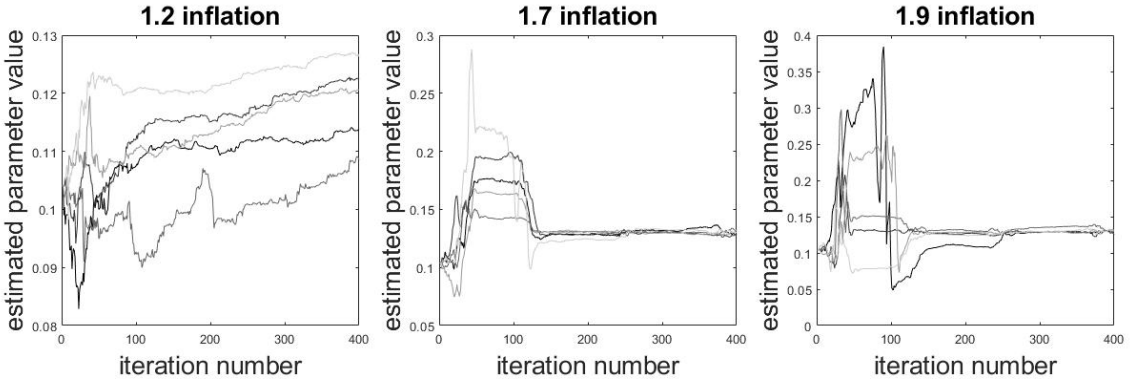

Figure 13: Illustration of the estimation of u_c with different multiplicative inflation amounts. Each line represents a different attempt at estimating u_c from the same initial conditions. The true value of u_c is 0.13. The resulting estimated parameter values's mean and standard deviation are: 0.118 ± 0.007, 0.130 ± 0.002, and 0.129 ± 0.003, respectively.

From this, we see the same general trends: a reduction in the multiplicative inflation factor results in slower (but steady) convergence to the true value of the parameter, and an increase in the multiplicative inflation results in larger jumps in the estimation of the parameter. These results are similar to the results of Figure 12. Similar results also hold for other parameters to which the Fenton-Karma model is sensitive, such as τ_r and u_c^{si}.

We now move our attention to the effect multiplicative inflation has on the success of the estimation of parameters to which the Fenton-Karma model is not sensitive. We first estimate τ_{v2}^-, which the Fenton-Karma model has been shown to be insensitive to in Figure 3. The results of the estimation of τ_{v2}^- at different inflation values are shown in Figure 14.

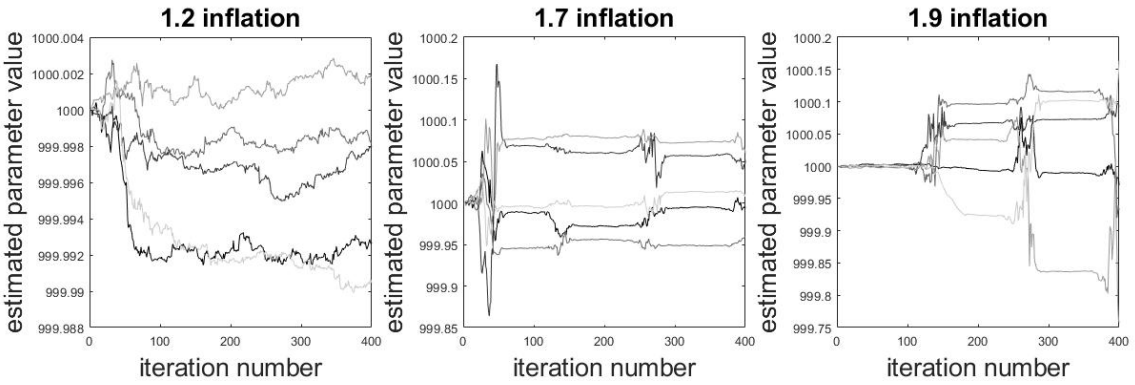

Figure 14: Illustration of the estimation of τ_{v2}^{-} with different multiplicative inflation amounts. Each line represents a different attempt at estimating τ_{v2}^{-} from the same initial conditions. The true value of τ_{v2}^{-} is 1250. The resulting estimated parameter values's mean and standard deviation are: 999.996 ± 0.005, 1000.015 ± 0.041, and 999.983 ± 0.157, respectively.

The same general trends can be seen in Figure 14: the increase of multiplicative inflation leads to an increase in the sharpness of jumps in the value of the estimated parameter, and leads to a further range of parameter values being tested although in all cases the parameter values changes very little (less than 0.1%). We also see that the parameter estimation does not seem to favor increasing or decreasing the value of the parameter; it simply perturbs the values of the parameter and checks to see if the state estimation is improved with the new value. As such, the successfulness of the parameter estimation for parameters that the Fenton-Karma model is not sensitive to such as τ_{v2}^{-} is highly dependent upon the randomness in the algorithm.

We also plot the results of the estimation of τ_{v2}^{-} with a multiplicative inflation value of 2.2 in Figure 15 to verify that this trend holds even beyond the recommended maximum value of multiplicative inflation, 2.

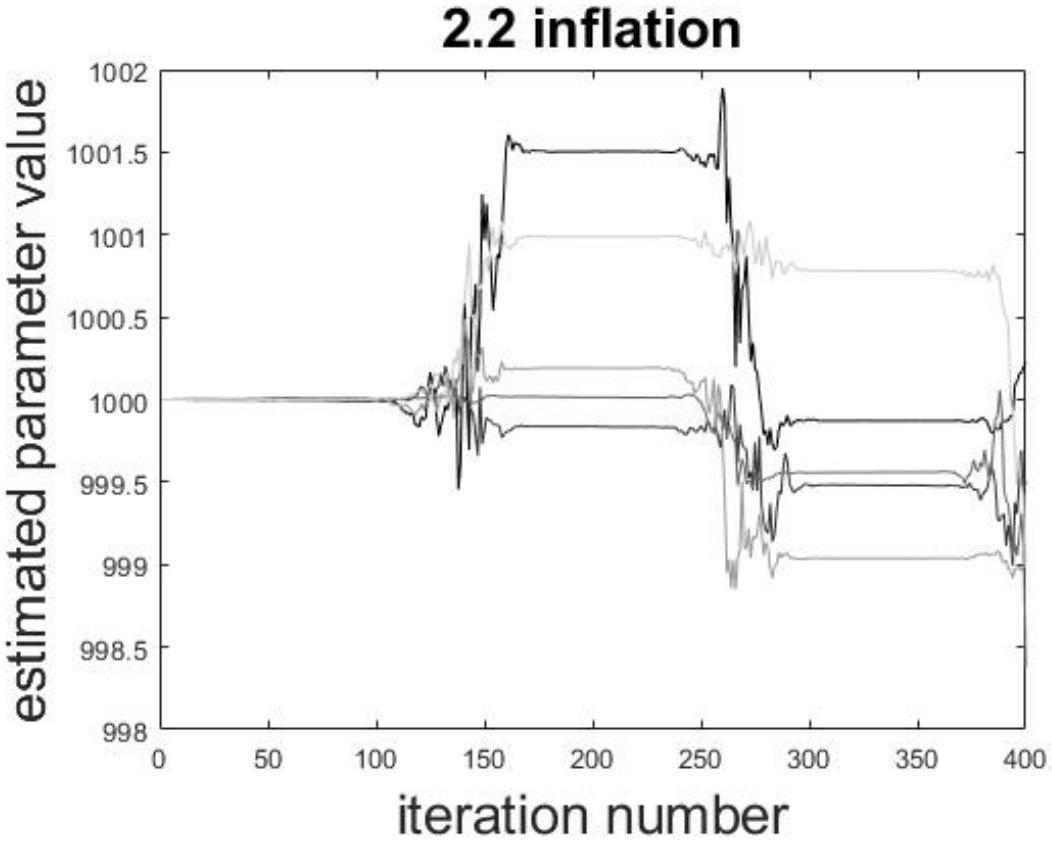

Figure 15: Illustration of the estimation of τ_{v2}^- with a multiplicative inflation value of 2.2. This shows the inability to estimate the parameter even when exceeding our initial multiplicative inflation maximum. The true value of τ_{v2}^- is 1250. The mean and standard deviation of the estimated value of the parameter is 999.228 ± 0.688.

The same trends continue, even past a multiplicative inflation value of 2. The parameter estimation is thus unable to reliably estimate certain parameters to which the model is insensitive, such as τ_{v2}^-, even with very large multiplicative inflation values. This multiplicative inflation is more than doubling the magnitude of the perturbations in the augmented X^b perturbation matrix, and it is still not enough to create a change in the parameter value larger than 1.5.

We conclude that the multiplicative inflation factor can impact the success of the parameter estimation algorithm, but in a much less important way than the impact of the ideas presented in section III.3.1. The multiplicative inflation factor impacts how large the magnitudes of changes to value of the estimated parameter are over the course of assimilation and often lead to overshooting the correct parameter value, making it take longer to estimate the parameter than it would with a

lower inflation value. Multiplicative inflation does not seem to play a large effect in whether or not the parameter estimation algorithm obtains the true parameter value, but rather how quickly it obtains the parameter value, and how drastic the parameter changes are from iteration to iteration. In this sense, multiplicative inflation is more about the speed and path of convergence towards the final estimated value of the parameter, and less about the likelihood of successful estimation.

III.4 Simultaneous Estimation of Two Parameters

We now show the results of estimating two parameters simultaneously. These parameters will be spatially and temporally homogeneous, as was the case in Section III.3. We first estimate two parameters to which the Fenton-Karma model was sensitive, τ_d and u_c^{si}, from an initial relative error of 20%. It can be seen in Figure 7 that both τ_d and u_c^{si} could be estimated well from these initial conditions when estimated separately. We show the dynamical effects of having a simultaneous error of 20% of both τ_d and u_c^{si} in Figure 16.

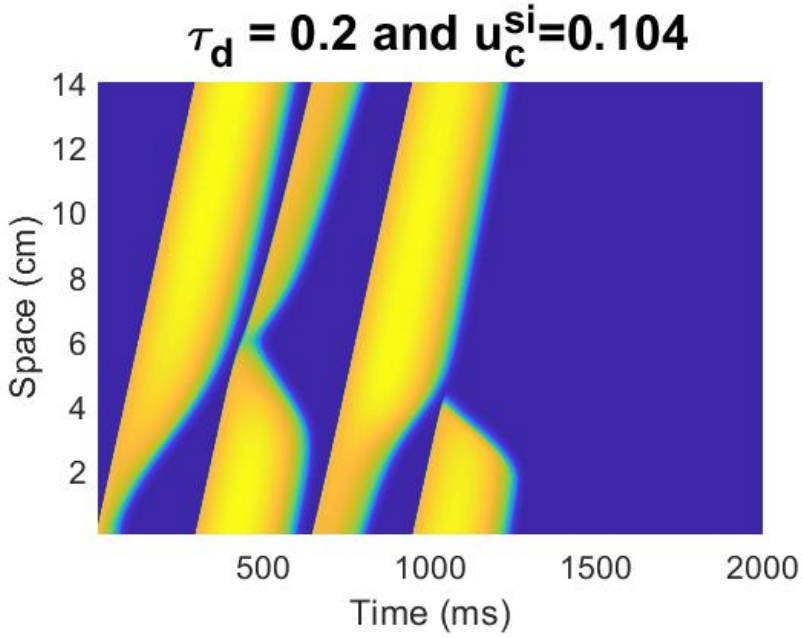

Figure 16: Illustration of the system dynamics when τ_d and u_c^{si} have relative errors of 20% and remain uncorrected. The true values of τ_d and u_c^{si} are 0.25 and 0.84, respectively.

Figure 17 shows the results of the simultaneous estimation of τ_d and u_c^{si} with a multiplicative

inflation value of 1.7.

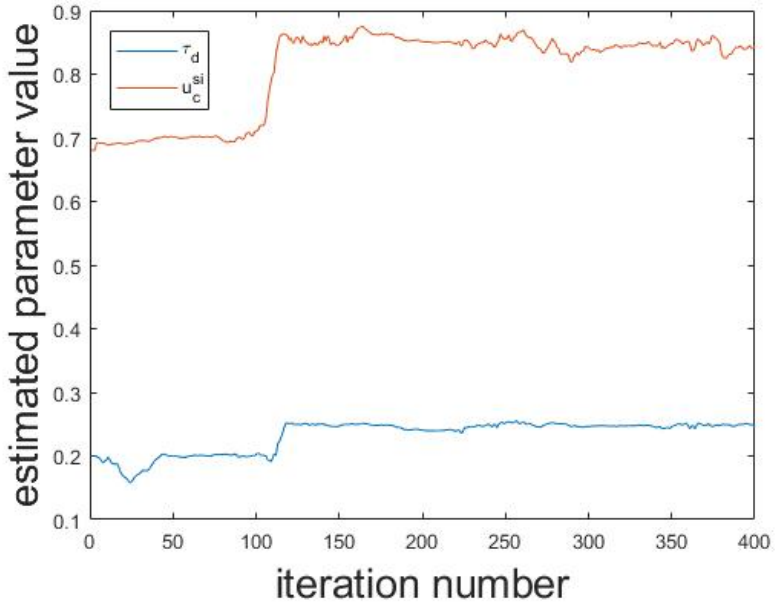

Figure 17: Illustration of the successful estimation of τ_d and u_c^{si} simultaneously. The Fenton-Karma model has been shown to be sensitive to both parameters and both parameters are reliably estimated simultaneously. The true values of τ_d and u_c^{si} are 0.25 and 0.85, respectively.

Both τ_d and u_c^{si} are estimated well, with their final estimated values being 0.2466 and 0.8433 compared to their true values of 0.25 and 0.85, respectively. A single instance of the parameter estimation algorithm is shown in this plot, but these results are representative of the behavior of the algorithm in general and match expectation that the algorithm can estimate both parameters successfully.

We next attempt to estimate one parameter to which the model is sensitive and one parameter to which the model is non-sensitive, τ_d and τ_{v2}^-. We have seen that τ_d is reliably estimated well when estimated alone, but τ_{v2}^- is not. Figure 18 shows the dynamical effects of having a simultaneous error of 20% of both τ_d and τ_{v2}^-, and Figure 19 shows the results of the simultaneous estimation of τ_d and τ_{v2}^- from an initial error of 20% with a multiplicative inflation factor of 1.7.

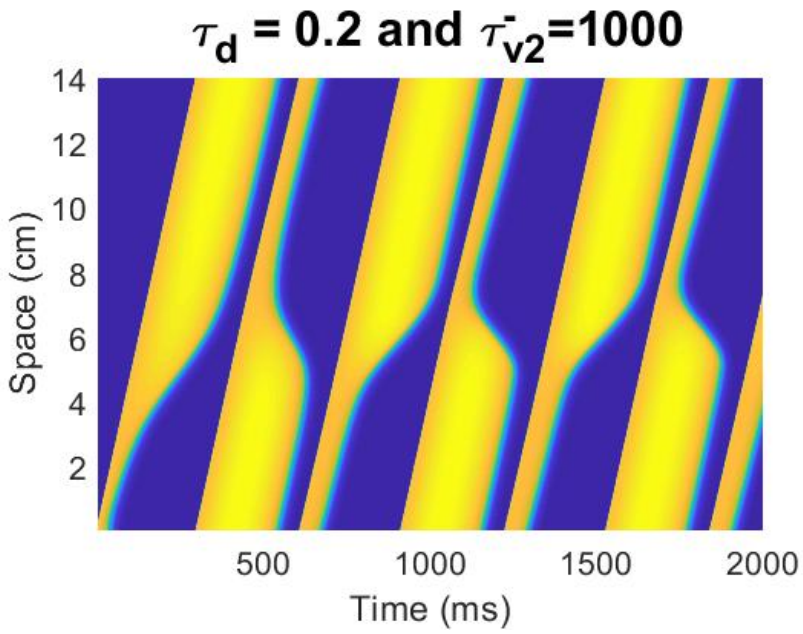

Figure 18: Illustration of the system dynamics when τ_d and τ_{v2}^- have relative errors of 20% and remain uncorrected. The true values of τ_d and τ_{v2}^- are 0.25 and 1250, respectively.

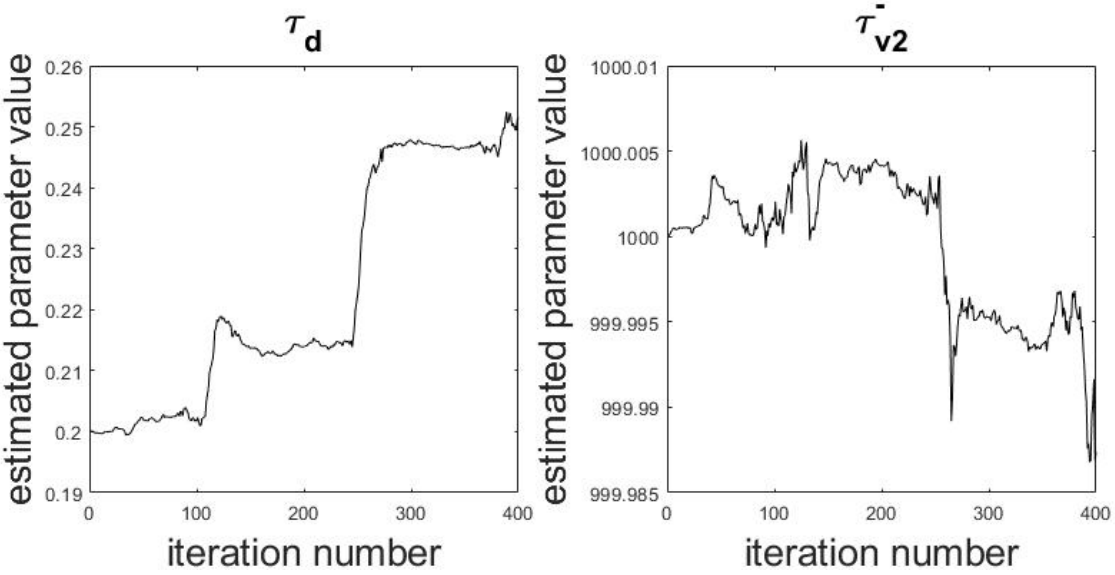

Figure 19: Illustration of the estimation of τ_d and τ_{v2}^- simultaneously. The parameter to which the model is sensitive, τ_d, is estimated well, and the parameter to which the model is not sensitive, τ_{v2}^-, is not estimated well. Both of these parameters were estimated simultaneously, but are displayed on separate plots due to their difference in scale. The true values of τ_d and τ_{v2}^- are 0.25 and 1250, respectively.

The algorithm estimates the parameter to which the model is sensitive well, and the parameter to which the model is not sensitive is estimated poorly, as was the case when they were estimated separately (see Figures 6 and 5, respectively). These results matches the expectation that the dynamically influential parameter will continue to get estimated well, as the error in the state estimation is more prominently reflected when there are errors in the estimated value of the parameter.

Finally, we consider the case of the simultaneous estimation of two parameters to which the model is not sensitive, τ_{v2}^- and u_v. These two parameters were shown to be estimated poorly when estimated separately in Figures 5 and 9 respectively. We show the dynamical effects of having a simultaneous error of 20% below the true value for both τ_{v2}^- and u_v in Figure 20 and present the results of the simultaneous estimation of τ_{v2}^- and u_v from 20% below the true value in Figure 21.

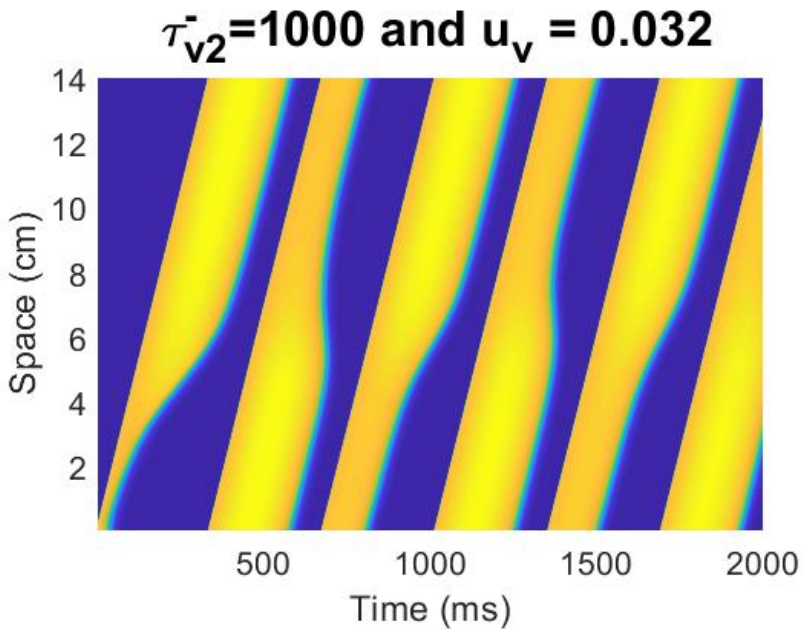

Figure 20: Illustration of the system dynamics when τ_{v2}^{-} and u_v have relative errors of 20% and remain uncorrected. The true values of τ_{v2}^{-} and u_v are 1250 and 0.04, respectively.

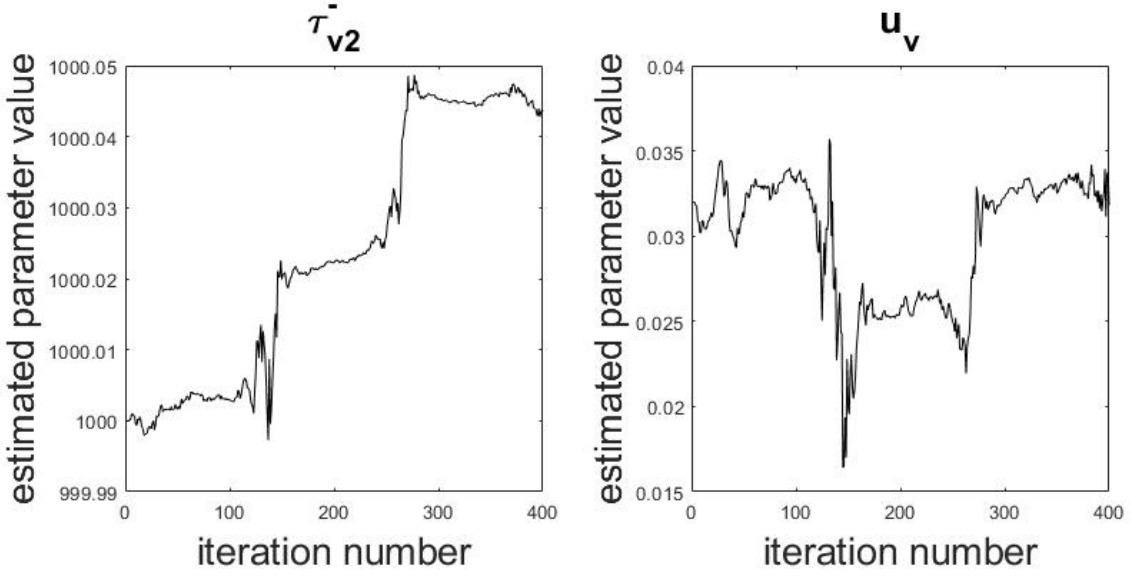

Figure 21: Illustration of the estimation of τ_{v2}^{-} and u_v simultaneously. The Fenton-Karma model is insensitive to both of these parameters, and they were unable to be estimated separately. The parameters were estimated simultaneously but are displayed on separate plots due to their difference in scale. The true values of τ_{v2}^{-} and u_v are 1250 and 0.04, respectively.

Neither τ_{v2}^{-} nor u_v is estimated well, matching the intuition gained from estimating them separately, as neither of the parameters was successfully estimated when done individually.

The results provided in this section all match up with what was found in section III.3. In the estimation of spatially and temporally homogeneous parameters, the most important factors that play a factor in the success of estimation are the dynamical sensitivity of the system to the parameters, the magnitude of the parameters, and the multiplicative inflation factor used in the estimation of the parameters. These results were well examined in the case of a single parameter being estimated, and short examples were provided for the case when two parameters are estimated simultaneously.

Due to the successful extrapolation of the results from one parameter to several, we expect that these three criteria will continue to be the biggest factors on the algorithm's success in other cardiac models and higher-dimensional regimes. Throughout this thesis, we have only considered the case of a one-dimensional ring geometry, but the problem can easily be extended to a two-dimensional or a three-dimensional problem (in fact, the biological motivation is studying the heart, which is naturally a three-dimensional setting). The simplification of the problem that we made in looking at the one-dimensional dynamics can be extended to higher-dimensional problems now that we have shown that the algorithm can be successful, and we have identified the factors that influence the success of the algorithm, which will provide a path forward to analyze higher-dimensional geometries.

There is an increased potential for incorrect estimation when multiple parameters are simultaneously. It is possible that an incorrect combination of parameter values will produce the same dynamics as the true parameter values. For example, in the case of the simultaneous estimation of two parameters, it is possible that another combination of parameter values (such as an increase in one parameter and a decrease in the other) will combine to produce the same dynamics as if the parameters had the true values used to generate the truth. In this case, the error could be equally low for an incorrect combination of estimated parameter values as for the true parameter values. This false correlation is a theoretical possibility of the algorithm when estimating multiple parameters simultaneously, although we have not encountered any examples of this happening.

We present one final example of the estimation of τ_d and u_c^{si} at different inflation values to illustrate the continued dependence on the parameter magnitude and inflation value in the case of simultaneous estimation of parameters. The simultaneous estimation of the two parameters with

different inflation values is shown in Figure 22.

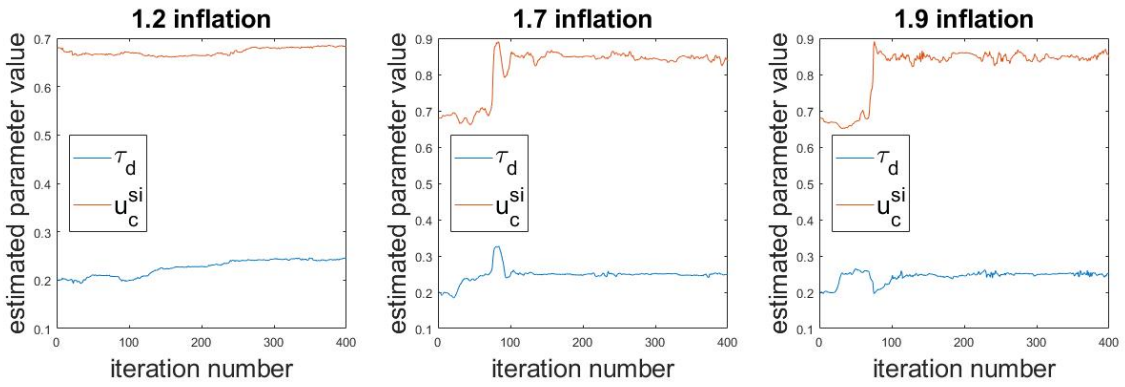

Figure 22: Illustration of the estimation of τ_d and u_c^{si} simultaneously. Both of these parameters have been shown to be sensitive in the model. The true values of τ_d and u_c^{si} are 0.25 and 0.85, respectively.

The two parameters are generally estimated well. The multiplicative inflation trend continues in this case: there are larger changes and more movement in the parameter values with higher multiplicative inflation factors. In fact, when a multiplicative inflation factor of 1.2 is used, u_c^{si} is not estimated correctly, but rather stays stagnant. The algorithm does manage to correct τ_d in this case though, so the error is reduced, and the state estimate converges well to the truth. As such, this multiplicative inflation value simply was not enough to cause the algorithm to estimate u_c^{si} accurately.

We also see that that in the case of a multiplicative inflation factor of 1.2, the parameter with the smaller magnitude (τ_d) is corrected, while the parameter with the larger magnitude is not. This agrees with our observations with estimation of a single parameter, as the inflation plays a larger role on the estimation of the smaller parameter.

With all of the provided examples and evidence, we see that the extension from the estimation of a single parameter to the estimation of multiple parameters is relatively straightforward. The behavior of the algorithm when estimating multiple parameters seems to be well predicted from the behavior of the algorithm when estimating a single parameter.

IV. Conclusion

IV.1 Results

In this thesis, we have focused on the effectiveness of parameter estimation using state-augmentation methods [24, 1] in the LETKF. We have used configuration settings for the algorithm that have been shown to be a reasonable choice for similar models in the weather-forecasting community [15]. We have generated synthetic truth data from a simulation of the Fenton-Karma model with a parameter set previously shown to generate complex dynamical states [9], and we have generated synthetic observation data points by sampling the truth and adding Gaussian noise to the sampled points to mimic observational error.

We have found a relationship between the effectiveness of parameter estimation and the effect of small variations of a parameter in the system's dynamics. Namely, the larger the dynamical differences produced by small variations in the value of the parameter, the more likely the parameter is to be estimated well using this method. The noticeable dynamical differences due to variations in the parameter value cause higher error in the RMSE of the state estimation, which causes the parameters to be adjusted more readily to the true value of the parameter. The sensitivity of the model to the parameter has been shown to be a reliable indicator of the success of parameter estimation in general in Figure 7. It is worth noting that due to the randomness inherent in the algorithm at multiple stages, the success or failure of parameter estimation is not guaranteed regardless of the sensitivity of the Fenton-Karma model to the parameter. Despite the dynamically sensitive τ_d being shown to be reliably estimated over five separate instances in Figure 6, there may be pernicious cases where the randomness is selected in such a way that the estimated value of the parameter does not converge to the true value of the parameter. However, these are atypical cases; more commonly, the dynamically influential parameters in the model are estimated reliably, regardless of randomness used during the algorithm. For all figures throughout this thesis, the algorithm has been run multiple times to determine typical behavior, and results have been selected to be representative of the normal output produced.

The magnitude of the parameter being estimated has also been shown to play an impact on the success of the estimation of the parameter. In the case of smaller parameter values, the perturbations in the X^b augmented background ensemble perturbation matrix are much larger relative to the parameter value. These larger perturbations cause the parameter estimation algorithm to make larger adjustments (relative to the parameter value) than if the parameter was

larger. This can be seen in Figure 9, where the value of the parameter u_v (which has a true value of 0.04) changes significantly, ranging from a decrease of 100% to an increase of 931%. This is in stark contrast to larger parameters such as τ_{v2}^{-} (see Figure 5), where the estimated value is only changed by a maximum of 0.015%. From this trend, we can conclude the parameters to which the model is sensitive and that are small in magnitude are especially likely to be estimated well; the sensitivity will cause numerically apparent errors in the estimation of the state, indicating when the estimated value is wrong and when it is correct, and the small magnitude will cause larger adjustments of the value of the parameter towards the true value, making it obtain the true value of the parameter more quickly.

We have examined the impact that multiplicative inflation can have upon parameter estimation. Figures 12, 13, and 14 show the general trends. As multiplicative inflation is increased, the estimated value of the parameter changes more significantly over the course of estimation. In cases where the model is not sensitive to the parameter, the increased perturbations to the estimated value of the parameter may allow the algorithm to approach or find the true value of the parameter, whereas smaller multiplicative inflation values may not have allowed the algorithm to find the true value of the parameter. An example of this can be found in Figure 12, where τ_d is not changed enough to reach the true value when the multiplicative inflation is 1.2, but it steadily tends towards the true value of the parameter. In cases where the model is sensitive to the parameter, these increased perturbations often lead to the an initial overcompensation in which the parameter value jumps past the true value of the parameter, and then has to change direction to approach the true value. Despite possibly overshooting the true value of the parameter during the course of estimation when the multiplicative inflation value is increased, the algorithm still manages to estimate the parameter well. This phenomenon can be seen in Figure 12 as well.

It has been found that the multiplicative inflation value largely does not impact the successfulness of the parameter estimation (except in cases where the parameter is insensitive and may need an extra "push" towards the true value of the parameter), but instead impacts the speed of convergence to the true parameter value, and the way it reaches the final estimated parameter value. As just discussed, the correct parameter value can often be overshot when the multiplicative inflation is increased, which leads to a different curve when plotted, such as in Figure 13. Multiplicative inflation also impacts the speed of convergence to the true parameter value. Parameters that the model has shown to be sensitive to seem to have an optimal level of multiplicative inflation associated with quickest convergence of the estimated parameter value to the true parameter value. The change in the speed of convergence can be easily seen in Figure 13, where the

lowest multiplicative inflation value leads to slow but steady estimation towards the true value of the parameter, but it is not reached within the time of assimilation, 2000ms. In the case of a medium value of multiplicative inflation, the true parameter values are overshot, as discussed previously, but are promptly recovered and successfully estimate the true value of the parameter within around 150 iterations, or 750ms. In the case of the highest multiplicative inflation, the overshooting is larger, and the algorithm overshoots again when it tries to correct itself from the first overcompensation, and all five instances of the parameter estimation converge to the true value within around 250 iterations, or 1250ms. From these results we see the lowest multiplicative inflation was not large enough, as the true value of the parameter was not recovered over the period of assimilation (2000ms); the medium multiplicative inflation was better as the parameter was estimated relatively quickly; and the highest multiplicative inflation was too large, and estimated the parameter value slower than the medium multiplicative inflation. From this, the speed of convergence of the algorithm is apparent, and there seems to be a "Goldilocks" situation in which there is an optimal multiplicative inflation value associated with a parameter whose value is not known a priori.

By combining results about the effect of the parameter's magnitude, and the multiplicative inflation used, we can also note the that "Goldilocks" multiplicative inflation is likely to be lower when the magnitude of the parameter is small. Since the parameter is small, the algorithm will make larger adjustments to the parameter value, which is the main purpose of increasing the multiplicative inflation. As such, since the relative changes in the estimated parameter value are larger due to the small magnitude of the parameter, the multiplicative inflation can likely be lowered to compensate for the increased volatility of the parameter estimation. This leads to the optimal multiplicative inflation value to maximize parameter estimation success and speed being lower for smaller parameters than larger parameters.

We have also found that these results found for a single parameter hold for the simultaneous estimation of multiple parameters as well. This is ideal, as this generally allows us to restrict our experiments and analysis to the estimation of a single parameter, and extrapolate the results from there. This is computationally easier and is also easier to analyze. The explanation of these effects, as discussed in Section III.3, are not as easy to reason through in a multiple parameter estimation case, as the parameters may be on different orders of magnitude in the X^b augmented background ensemble perturbation matrix, making the factors influence the parameters to differing scales. But we have shown that despite the more difficult analysis of multiple parameter estimation cases, it continues to exhibit the same phenomena as the single variable estimation case.

IV.2 Limitations

The results presented in this document have provided substantial information about the effectiveness of the state-augmentation method of parameter estimation on the Fenton-Karma, including providing three criteria that have been shown to be strong indicators of the performance of the parameter estimation algorithm. Despite the thorough account of the parameter estimation algorithm on various parameter sets in the Fenton-Karma model, several aspects have not been considered for the sake of complexity and scope. Below we discuss additional avenues of research that would be useful to pursue.

The sign of initial error in the estimation of the parameter was largely ignored in this document. That is, that initial value for each estimated parameter was than the true value of the parameter. In this sense, we have assumed that the initial error is negative; the true value is larger than the initial "guess" of the parameter value ($p_{guess} - p_{true} < 0$). Early on in this document it was noted that the performance of the parameter estimation algorithm was largely symmetric (i.e., an initial error of 20% will have approximately the same behavior for estimation, regardless of whether or not the error was over- or underestimated), and we elected to start with an underestimated initial value for all parameters we estimated from that point onward. See Section III.3.1 for a discussion of the performance of the algorithm with over- and underestimated initial values. Due to the algorithm generally performing in a symmetric manner, many parameters were not tested with initial error estimates being over- and underestimated. It is possible that unexpected behavior could result from switching the initial parameter values in some cases.

The magnitude of the initial error in the estimation of the parameter was also ignored. For the purposes of this thesis, all initial errors in the value of the parameter(s) being estimated were 20% of the true value of the parameter. Simulations were done on a subset of the parameters being estimated in which the magnitude of the initial error was increased (varying between 20, 40 and 50%). Results proved to worsen as the magnitude of the error was increased in a predictable fashion. It is possible that unpredictable behavior happens in the parameter estimation algorithm with varying magnitudes of initial error in certain scenarios, but it remains outside the scope of this project.

Finally, many parameter combinations have not been tested in the simultaneous parameter estimation algorithm. As there are 13 parameters listed for estimation in Table 2, there are 78 different combinations of two parameters to estimate simultaneously. As such, we have not tested all pairs of parameters in the Fenton-Karma model for simultaneous estimation,

and certain parameter pairs may provide unexpected results. Furthermore, we have not tested simultaneous estimation of parameters of three of more parameters at once, although this algorithm is flexible enough to accommodate more than two simultaneous parameters being estimated; simply appending the additional parameters to the ensemble (see Equation 9) allows for the estimation of more than two parameters at once. It is possible that a combination of parameter values that are not the true values provides equally low RMSE of the estimation of the state.

IV.3 Future Work

The work throughout this thesis has largely served to show the viability of this method of parameter estimation in a cardiac model and to establish heuristics about the success of the algorithm on various parameters. Due to this document serving as an initial proof of concept, the work leaves many avenues for future study.

First, altering the initial error in the value of the parameters being estimated is necessary in order to verify the criteria determined in this thesis are consistent despite initial error and are not artifacts of the model being used in the algorithm. Alteration of the initial error based upon direction (over- and underestimation) and magnitude will provide additional confidence that these criteria hold across all parameters in the model and are not a special case when the initial error is 20% below the true value of the parameter.

Second, additional parameter combinations should be chosen and tested for accordance with the three criteria determined in this document. Simultaneous parameters should be tested to determine whether they fit the expected heuristic behavior provided by this document and also whether the parameters being estimated behave in a manner similar to when they are being estimated alone. This behavior was captured in Section III.4 in which two parameters were estimated and the two parameters were estimated similarly to how they were estimated when estimated separately. This is ideal, as consitent estimation allows the parameters to be estimated separately and allows the extrapolation to simultaneous parameter estimation. In order to verify the consistency of parameter estimation regardless of the number of parameters being estimated simultaneously, more parameter combinations should be tested, and parameter combinations of more than two parameters should be chosen to verify this behavior continues beyond the simultaneous estimation of two parameters.

Finally, the H observation operator should be studied to determine the impact that the observations play upon the success of the estimation of the parameters. We used a uniform spacing of one

third of the ring observed by the H operator, which equates to observational data points available every three grid points, or 0.75mm. Two possible changes to the H operator that may prove fruitful for parameter estimation are to change the spacing of the observational points and have H represent a coarser resolution. If H is made to have non-uniform spacing, it is possible that it is more important to have observational data for certain areas of the ring (based upon the importance of capturing dynamics, it is most likely important to capture the wave front with the observations) and other areas can remain unobserved and maintain similar or better RMSE results. If H is made to have a coarser resolution, it could provide insight as to the extent with which the LETKF parameter estimation algorithm can recover the state and parameter values sufficiently. Our spatial and temporal resolutions used throughout this document were less precise than experimental optical mapping technologies [23], but clinical scenarios could benefit from knowledge about the resolution required in order to use this algorithm effectively. For optical mapping data in particular, modification of the H operator may be necessary in order to account for contributions from cells not on the tissue surface, as electrical signal is not purely from the surface of the tissue.

V. Bibliography

References

[1] Altuğ Aksoy, Fuqing Zhang, and John W. Nielsen-Gammon. "Ensemble-Based Simultaneous State and Parameter Estimation in a Two-Dimensional Sea-Breeze Model". In: *Monthly Weather Review* 134.10 (2006), pp. 2951–2970. DOI: 10.1175/MWR3224.1. URL: https://doi.org/10.1175/MWR3224.1.

[2] *Arrhythmia*. URL: https://www.nhlbi.nih.gov/health-topics/arrhythmia.

[3] Mina Attin and William T. Clusin. "Basic Concepts of Optical Mapping Techniques in Cardiac Electrophysiology". In: *Biological Research For Nursing* 11.2 (2009), 195–207. DOI: 10.1177/1099800409338516.

[4] G. W. Beeler and H. Reuter. "Reconstruction of the action potential of ventricular myocardial fibres". In: *The Journal of Physiology* 268.1 (1977), pp. 177–210. DOI: 10.1113/jphysiol.1977.sp011853.

[5] Emelia J. Benjamin et al. "Heart Disease and Stroke Statistics—2017 Update: A Report From the American Heart Association". In: *Circulation* 135.10 (2017), e146–e603. DOI: 10.1161/CIR.0000000000000485.

[6] Marc Courtemanche, Leon Glass, and James P. Keener. "Instabilities of a propagating pulse in a ring of excitable media". In: *Physical Review Letters* 70.14 (1993), pp. 2182–2185. DOI: 10.1103/physrevlett.70.2182.

[7] Geir Evensen. "The Ensemble Kalman Filter: theoretical formulation and practical implementation". In: *Ocean Dynamics* 53.4 (2003), pp. 343–367. DOI: 10.1007/s10236-003-0036-9.

[8] R T Faghih et al. "The Fitzhugh-Nagumo model: Firing modes with time-varying parameters & parameter estimation". In: *2010 Annual International Conference on the IEEE Engineering in Medicine and Biology* (2010). DOI: 10.1109/iembs.2019.5627326.

[9] Flavio Fenton and Alain Karma. "Vortex dynamics in three-dimensional continuous myocardium with fiber rotation: Filament instability and fibrillation". In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 8.1 (1998), pp. 20–47. DOI: 10.1063/1.166311.

[10] Alessio Gizzi et al. “Effects of Pacing Site and Stimulation History on Alternans Dynamics and the Development of Complex Spatiotemporal Patterns in Cardiac Tissue”. In: *Frontiers in Physiology* 4 (2013), p. 71. ISSN: 1664-042X. DOI: 10.3389/fphys.2013.00071. URL: https://www.frontiersin.org/article/10.3389/fphys.2013.00071.

[11] Thomas M. Hamill, Jeffrey S. Whitaker, and Chris Snyder. “Distance-Dependent Filtering of Background Error Covariance Estimates in an Ensemble Kalman Filter”. In: *Monthly Weather Review* 129.11 (2001), pp. 2776–2790. DOI: 10.1175/1520-0493.

[12] Todd J. Herron, Peter Lee, and José Jalife. “Optical Imaging of Voltage and Calcium in Cardiac Cells & Tissues”. In: *Circulation Research* 110.4 (2012), pp. 609–623. DOI: 10.1161/circresaha.111.247494.

[13] A L Hodgkin and A F Huxley. “A quantitative description of membrane current and its application to conduction and excitation in nerve”. In: *Journal of Physiology* 117.4 (1952), 500–544.

[14] M. J. Hoffman et al. “Reconstructing three-dimensional reentrant cardiac electrical wave dynamics using data assimilation”. In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 26.1 (2016). DOI: 10.1063/1.4940238.

[15] Brian R. Hunt, Eric J. Kostelich, and Istvan Szunyogh. “Efficient data assimilation for spatiotemporal chaos: A local ensemble transform Kalman filter”. In: *Physica D: Nonlinear Phenomena* 230.1-2 (2007), pp. 112–126. DOI: 10.1016/j.physd.2006.11.008.

[16] R E Kalman. “A New Approach to Linear Filtering and Prediction Problems”. In: *ASME-Journal of Basic Engineering*. D 82.1 (1960), pp. 35–45. DOI: 10.1115/1.3662552.

[17] Alain Karma. “Electrical alternans and spiral wave breakup in cardiac tissue”. In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 4.3 (1994), pp. 461–472. DOI: 10.1063/1.166024.

[18] Boris Y. Kogan et al. “The simplified FitzHugh-Nagumo model with action potential duration restitution: Effects on 2D wave propagation”. In: *Physica D: Nonlinear Phenomena* 50.3 (1991), pp. 327–340. DOI: 10.1016/0167-2789(91)90003-r.

[19] Hiroshi Koyama and Masahiro Watanabe. “Reducing Forecast Errors Due to Model Imperfections Using Ensemble Kalman Filtering”. In: *Monthly Weather Review* 138.8 (2010), pp. 3316–3332. DOI: 10.1175/2010mwr3067.1.

[20] Nicholas S. LaVigne et al. “Effects of model error on cardiac electrical wave state reconstruction using data assimilation”. In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 27.9 (2017), p. 093911. DOI: 10.1063/1.4999603. URL: https://doi.org/10.1063/1.4999603.

[21] W. Gregory Lawson and James A. Hansen. “Implications of Stochastic and Deterministic Filters as Ensemble-Based Data Assimilation Methods in Varying Regimes of Error Growth”. In: *Monthly Weather Review* 132.8 (2004), pp. 1966–1981. DOI: 10.1175/1520-0493.

[22] C H Luo and Y Rudy. “A dynamic model of the cardiac ventricular action potential. I. Simulations of ionic currents and concentration changes.” In: *Circulation Research* 74.6 (1994), pp. 1071–1096. DOI: 10.1161/01.res.74.6.1071.

[23] Florence Rothenberg et al. “Emerging patterns of cardiac conduction in the chick embryo: Waveform analysis with photodiode array-based optical imaging”. In: *Developmental Dynamics* 233.2 (2005), 456–465. DOI: 10.1002/dvdy.20338.

[24] Juan Jose Ruiz, Manuel Pulido, and Takemasa Miyoshi. “Estimating Model Parameters with Ensemble-Based Data Assimilation: A Review”. In: *Journal of the Meteorological Society of Japan. Ser. II* 91.2 (2013), pp. 79–99. DOI: 10.2151/jmsj.2013-201.

[25] Sanita Vetra-Carvalho. *The Ensemble Kalman filter Part II: Practicalities*. 2018.

[26] Arthur T. Winfree. “Varieties of spiral wave behavior: An experimentalist’s approach to the theory of excitable media”. In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 1.3 (1991), pp. 303–334. DOI: 10.1063/1.165844.