Shape Theoretic and Machine Learning Based Methods for Automatic Clustering and Classification of Cardiomyocytes Based on Action Potential Morphology

by

Giann Gorospe

A dissertation submitted to The Johns Hopkins University in conformity with the requirements for the degree of Doctor of Philosophy.

Baltimore, Maryland

October, 2017

Abstract

Stem cells have been a hot topic in the cardiology community for the last decade and a half. Ever since we learned how to differentiate cardiomyocytes from embryonic and induced pluripotent stem cells, there has been a lot of research devoted to the potential of utilizing these cardiomyocytes for regenerative medicine, drug model studies, and arrhythmogenesis analysis. However, while cardiomyocyte purification methods have advanced significantly, methods for the identification and isolation of specific types of cardiomyocytes, such as ventricular or pacemaking cells, have not seen the same progress. Among the different avenues for accomplishing this task, the electrophysiological one is of particular interest because every cardiomyocyte type generates a distinct signature known as an action potential. The current standard for analyzing the action potential of a cardiomyocyte is an expert-level subjective thresholding of specific features, such as action potential duration. However this approach does not transfer across datasets and does not scale with the increasing populations of cardiomyocytes.

In this thesis, ideas from the machine learning and shape analysis communities

ABSTRACT

are explored to develop new, automated methods for the analysis of cardiomyocytes based on their action potentials. These methods allow us to identify subpopulations of similar cardiomyocytes based on their action potential morphology, hypothesize the eventual chamber-specific fate of newly differentiated cardiomyocytes, and make effective comparisons between cardiomyocytes in drug and cell-line studies. The objective, scalable methods presented in this thesis present a new paradigm in performing analysis in high-throughput applications of cardiomyocytes via action potential morphology, and could be of large benefit to the cardiology and biology communities.

Commitee Members:

Dr. René Vidal, JHU BME (Primary Reader)

Dr. Leslie Tung, JHU BME (Secondary Reader)

Dr. Laurent Younes, JHU AMS

Acknowledgments

There are so many people I would like to thank for this thesis. First and foremost, thank you to my advisor, Rene, whose constant guidance helped me achieved this momentous feat. Thanks to my committee, Dr. Tung and and Dr. Younes, for continually providing additional insight that helped me guide my tests and developments. Thanks to Renjun Zhu, my primary contact for data, and primary collaborator throughout all of these experiments. Thanks to the labmates that went before me, Alvina, Ertan, Dheeraj, Avinash, Rizwan, Roberto, and Ehsan for helping me get settled in and my feet wet. Thanks to my lab contemporaries, Siddharth, Lingling, Chong, Evan, Manolis, Efi, Flori, Connor, and all of the interns and master students, who made every lunchtime enjoyable. Thanks to Dr. Miller for always spurning me to do great things, and Dr. Ratnanather for making walking the hallways an entertaining experience.

Thank you to my friends who traveled on this Ph.D. journey with me. I still remember when we ended up walking what felt like forever to get to that restaurant during our first retreat. Thanks for sharing your time with me as we all took on this

task in our own ways. Thanks for the frisbee, the late nights watching movies or playing games, and everything else that made the Ph.D. experience more engaging. Thanks to my friends outside of the Ph.D program, like the family at the Potomac Curling Club. When I needed reprieve from science, you guys were always there, and you have certainly helped me be a more complete person outside of this degree.

Dedication

This thesis is dedicated to my family, whose love motivates me to give my everything.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Human Stem Cell-Derived Cardiomyocytes

Since the turn of the century, stem cells have presented a new frontier in cardiology. Ever since the first stem cell derived cardiomyocytes were obtained in [1], the hope has been that these newly formed heart muscle cells would provide breakthroughs in modern medicine. In particular, two flavors of newly differentiated, immature cardiomyocytes have been the focus of a lot of study over the past 20 years. Human embryonic stem cell derived cardiomyocytes (hESC-CMs) are derived from human fetal tissue and are often lauded for their plasticity. This adaptability proved particularly helpful in highlighting the potential of stem cells to be used in regenera-

tive medicine back in 2007, when Laflamme et al. showed that hESC-CMs could be used to temporarily recover functional activity in regions of murine hearts that had been damaged due to myocardial infarction [2]. However, given the moral implications of the origin of these cells, they are increasingly less utilized. More common now are cardiomyocytes that come from reverse engineered stem cells obtained from other organs, such as the skin. These human induced pluripotent stem cell derived cardiomyocytes (hiPSC-CMs) have been shown to have less plasticity than their embryonic counterparts, but the resulting cardiomyocytes are still functionally similar to mature, or adult, heart cells. This similarity has led to even more interest in stem cells from the cardiology community, as researchers attempt to investigate the viability of hiPSC-CMs as models for testing the effects of drugs, as well as factors of arrhythmias [3, 4].

Since the groundbreaking work of [1], the methods for obtaining functional cardiomyocytes has vastly improved. For instance, multiple methods to isolate cardiomyocytes have been proposed [5–8], including methods that can also lead differentiated cardiomyocytes towards a specific cellular makeup indicative of a specific region of the heart [9, 10]. However, there are many applications in regenerative medicine where understanding the phenotype identity is paramount to maintaining functional compatibility and reducing the risk of arrhythmias. For instance, introducing pacemaking cells into the left ventricle of the heart might lead to asynchronous beating in that portion of the heart, leading to more harm than good. Despite this, identifying

the chamber specificity of an embryonic cardiomyocyte remains a challenge, as even commercially available cell lines do not guarantee the uniformity of chamber specific (ventricular, atrial, pacemaking) cell populations. Thus the question of identifying the phenotype of these embryonic cardiomyocytes still lies with the biologists at the laboratory.

One approach to identifying the phenotype of an embryonic cardiomyocyte is to identify proteins that are indicative of a specific chamber of the heart, and develop fluorescence markers for those proteins and methods to insert them into the cardiomyocytes [11]. The common example is the myosin heavy chain protein MHC-2v indicative of the ventricular phenotype. Since this approach is driven by the biology of adult cardiomyocytes, these protein biomarkers are regarded as strong indicators of chamber specificity. However, this approach requires not only the identification of markers, but also their modification to carry fluorophores and their careful insertion and uptake into target cardiomyocytes, all of which are labor intensive and invasive tasks. Additionally, the embryonic cardiomyocytes appear to present multiple chamber phenotypes, as eventual working chamber (atrial and ventricular) heart cells have been observed to show spontaneous beating patterns indicative of the pacemaking phenotype.

A much less invasive approach to analyzing cardiomyocytes is the electrophysiological approach. Specifically, ions are carefully taken in and released from a cardiomyocyte via the activity of various ion channels. An example cardiomyocyte with

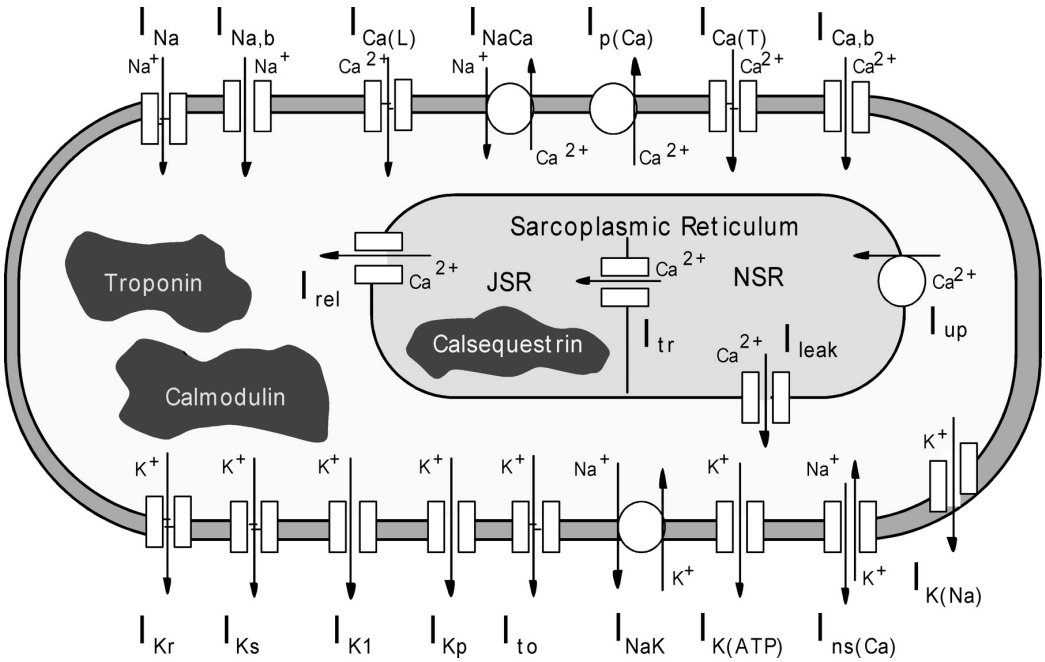

Figure 1.1: Ion channel depiction of a cardiomyocyte, taken from [12]

representative ion channels, as described by [12], is shown in Figure 1.1. The selectivity of ions that pass through the cell membrane generates a voltage potential between the inside of the heart cell and the outside. As a result of this potential and how it fluctuates throughout the routine functions of the heart, every cardiomyocyte generates a unique electrical signature, known as an action potential. Examples of action potentials from different regions of the heart are shown in Figure 1.2, as described in [13]. Notice that different regions produce distinct action potential shapes. As such, while this is a simplified feature of very complex processes inside a particular cardiomyocyte, it carries enough information to be discriminative of specific regions of the heart. Because of these reasons, the rest of this thesis will focus primarily on how to utilize the action potential of hSC-CMs to determine its phenotype.

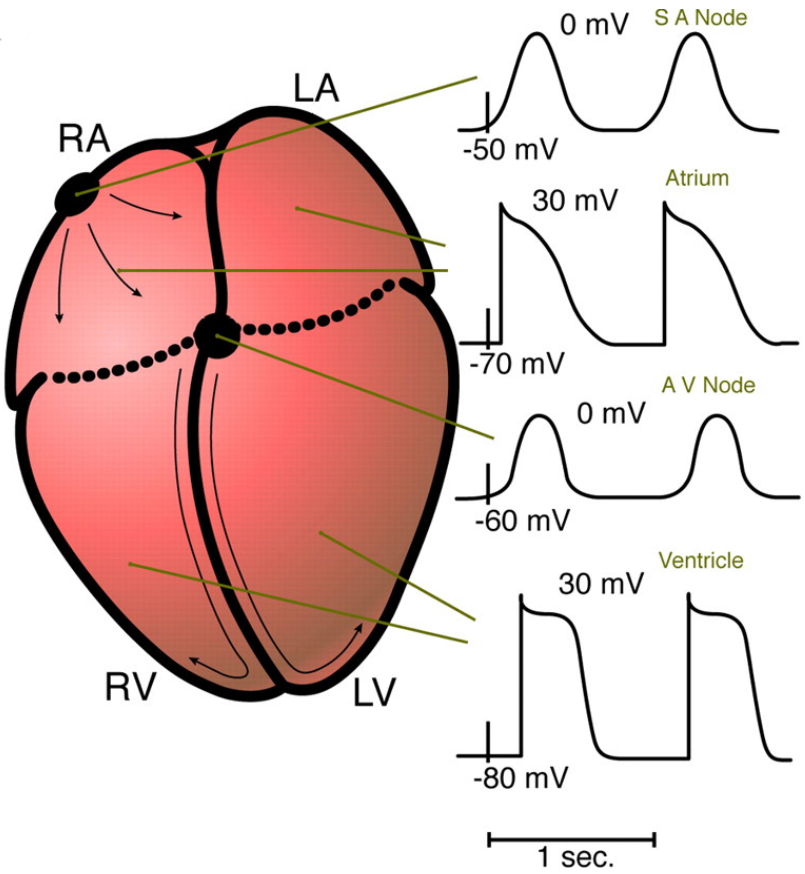

Figure 1.2: Action potentials from different heart regions, as described in [13]

1.2 Electrophysiological Methods for Differentiating Cardiomyocytes

This section describes methods for differentiating hSC-CMs based on the analysis of their action potential. In order to accomplish this, I first describe the typical phases of the action potential in more detail as well as methods for recording action potentials.

1.2.1 Cardiomyocyte Action Potential

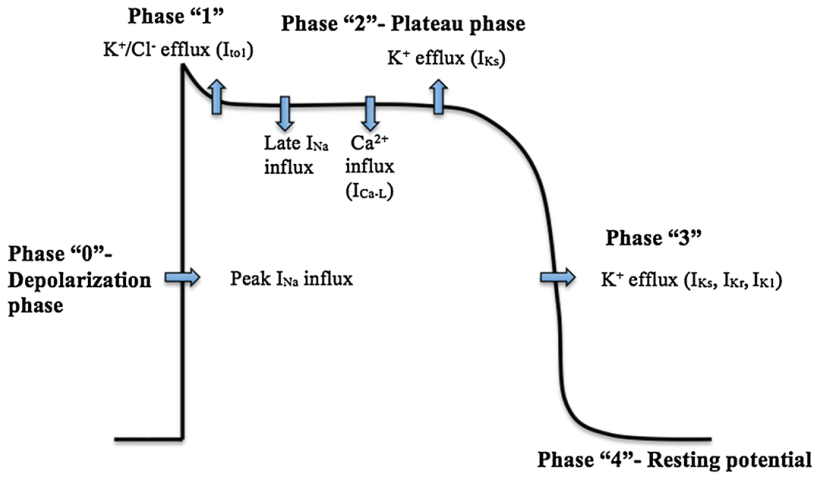

Figure 1.3: Phases of the cardiac action potential, with primary active currents during each phase labeled, courtesy of [14]

The cardiomyocyte action potential is typically described as consisting of 5 phases, as shown in the cardiomyocyte action potential from [14] in Figure 1.3. A large portion of the action potential cycle is spent at rest in phase 4, where the potential is constant, but the ion concentrations within and outside the cell are actively trying to achieve equilibrium between the inside and the outside of the cell. The action potential is typically initiated with a stimulus, and leads to a net influx of positive ions into the cell via primarily sodium channels, but also slower acting calcium channels. This phase 0 depolarization leads to the rapid increase in the membrane voltage. Eventually the sodium channels close and the potassium channels open, leading to potassium ions being expelled out of the cell. This generates the phase 1 "notch" of the action potential. Eventually, the potassium channels are countered by the slow influx of calcium and chlorine ions, leading to a temporary equilibrium in the mem-

brane potential. This plateau of voltage is phase 2 of the cardiac action potential cycle. In the final phase, eventually the calcium channels close and additional potassium channels open, expunging ions out of the interior of the cell. This leads to the repolarization of the cell membrane and a return to the resting potential.

1.2.2 Methods for Recording Cardiac Action Potentials

As the methods for isolating cardiomyocytes expanded, so, too, did the methods for obtaining action potential recordings. There are now two common methods to recording individual cardiomyocyte action potential signals. The long established way has been to record the action potential of an individual cell using a microelectrode [15, 16]. These electrode recordings have high signal to noise, but recording at individual sites may not be indicative of a population of cardiomyocytes, such as a cardiomyocyte monolayer or cardiomyocyte aggregates.

In order to get recordings from an entire population simultaneously, voltage sensitive dyes and optical mapping [17] have started to become more commonly utilized. The fluorescence in the dye oscillates with the change in membrane potential, and this fluorescence is captured by a standard charge-coupled device. Typically, the exact membrane potential is lost, but the scaled fluorescence still captures the general morphology of the action potentially up to scale. This allows one to record one ac-

tion potential from every pixel in the camera simultaneously, which has resulted in a growth in populations of action potential recordings from the order of tens to the order of hundreds and thousands. Depending on the FOV and resolution of the camera, each pixel can could contain the action potential recording from a single heart cell, or the average recording of the multiple cardiomyocytes contained in the pixel.

1.2.3 Current Methods for Analyzing Action Potential

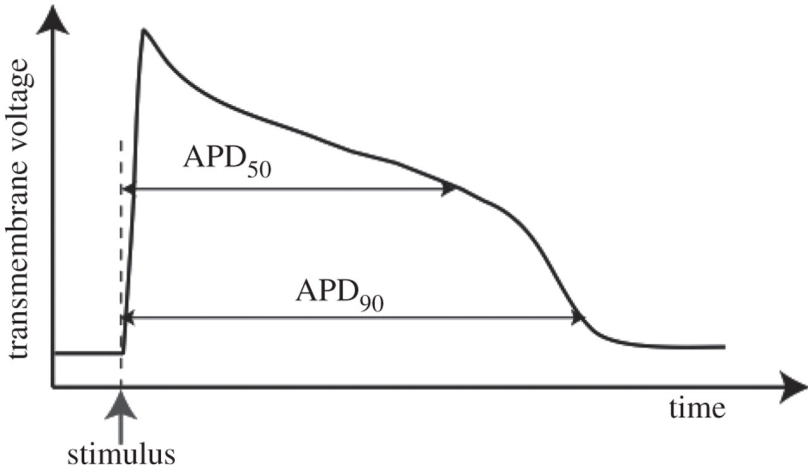

Figure 1.4: Example action potential showing measurements of APD_{50} and APD_{90}, as described in [18]

Lagging behind these key advances in cardiomyocyte differentiation and action potential recording in the last 15 years have been methods to analyze action potential data. The current practice is to obtain measurements of certain features of the action potential, such as the amount of time it takes for the action potential to recover $x\%$

from depolarization (commonly referred to as action potential duration, or APD_x), or the slope of the depolarization phase. An example of these features on a cardiac action potential is given in Figure 1.4, courtesy of [18]. Perhaps more egregious is that, to identify the phenotype, cardiologists have looked at these features and come up with subjective thresholds to attempt to identify cardiomyocytes as chamber specific cells. One such delineation [16] is shown in Figure 1.5 These thresholds typically do not translate from dataset to dataset, as they only encode a subset of the action potential signal specific to the dataset in question. Further, manual annotations of these measurements are intractible as newer methods produce ever increasing numbers of action potential data.

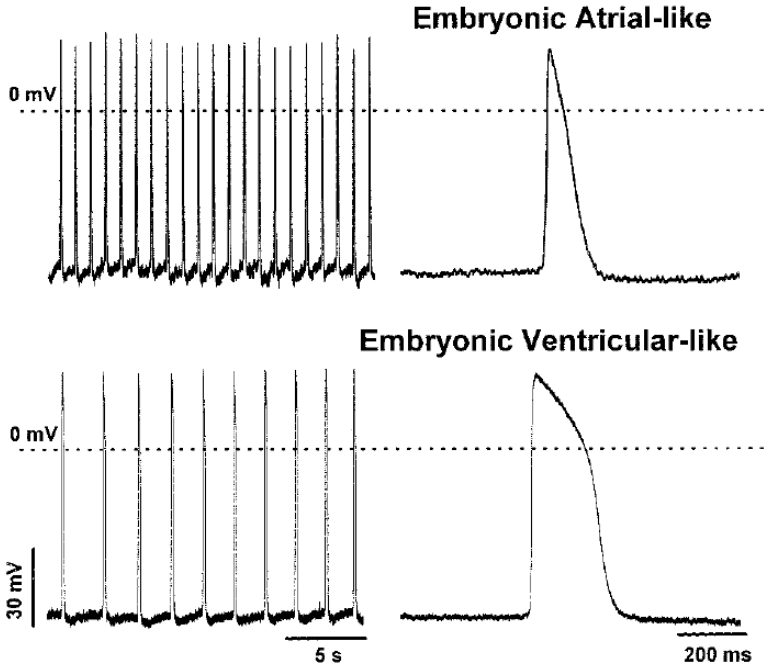

Figure 1.5: Expert Biologist Classification of embryonic action potentials, taken from [16]

1.2.3.1 Challenges

What appears to be largely missing from this analysis are well trained algorithms that can help biologists compare and discriminate the action potential recording data. In a perfect world, the algorithm unifies all the information from different studies, reducing the subjectivity of the manual classification by a biologist or cardiologist, and making it more efficient to analyze new action potentials as new recordings are generated. The process of training a computer to identify trends and patterns in data is known as machine learning [19]. Machine learning methods have already found their impact in other fields of biology, such as gene regulatory networks [20], biomedical image segmentation [21], and epidemiology [22]. Action potential data, being a one dimensional time signal, has analogs to other datasets where machine learning has already been utilized, such as speech recognition [23]. The premise of this thesis is that they can also be of some impact for the analysis of cardiac action potentials.

Additionally, there is much more information to be leveraged by observing the entire action potential, instead of just the standard biological features. But in order to do fully utilize this additional information, appropriate metrics between action potential signals need to be defined in order to be utilized by the machine learning algorithms. While standard metrics, such as the Euclidean distance or kernel methods, can be utilized, they can fail because the modes of variation captured by these metrics may not be indicative of the common variations in action potentials. An example of this is shown in Figure 1.6, where the query signal has a shape similar

to one phenotype, but might be misclassified because the raw potential values are closer to another phenotype. It is more appropriate to define metrics that have been tuned to action potential data and its variations, affording smaller penalty to common variations in action potentials while more strongly penalizing unlikely changes in action potential morphology. In addition, an appropriate metric should not only discriminate between cardiomyocytes in different regions of the heart, but also provide a link to connect action potentials at the immature state of cardiomyocytes, where chamber specificity is less clear, to action potentials of completely differentiated adult cardiomyocytes. The mathematics of defining such metrics is core topic of the shape theory community. By learning how to utilize the entire action potential morphology via shape theoretic methods, we can improve upon the standard biological features and potentially provide better insight into some mechanisms of cardiomyocytes.

As we have discussed previously, the action potential is the electrophysiological output of a complex nonlinear system. While some of the factors that influence the output are intrinsic to the cardiomyocyte, such as the number of sodium or potassium channels, external factors, such as the stimulus and its rate of occurrence, its connectivity to other cardiomyocytes, or the presence or absence of a drug that targets specific ion channels also play a vital role in the output of the system. These external factors could influence the estimation of the phenotype given by shape theoretic methods. Multiple action potential recordings of a cardiomyocyte under different environmental conditions may help alleviate these ambiguities regarding phenotypical

identity, but how we utilize this additional information to obtain as much information about the cardiomyocyte has not been throughly studied.

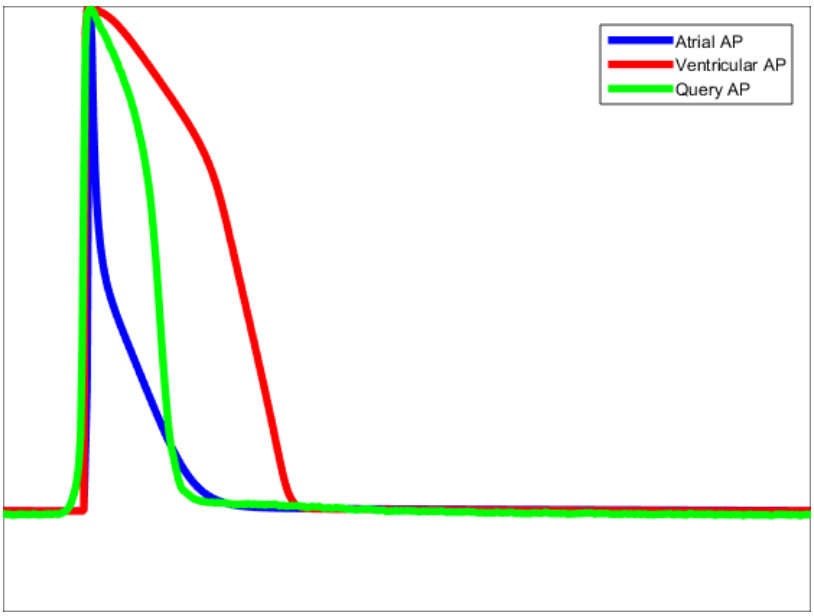

Figure 1.6: Example of how simple metrics may be insufficient to identify action potential shape similarities. The query AP (green) has a similar shape to the ventricular action potential (red), but membrane potential values closer to the atrial action potential (blue).

1.3 Thesis Contributions

The goal of this thesis is to improve the cardiologist's and biologist's study of individual and collections of cardiomyocytes by developing shape theoretic and machine learning tools and algorithms to analyze the action potential in an objective way. For this purpose, I have developed an automated framework for both separating populations of newly differentiated stem cell derived embryonic cardiomyocytes

into subgroups (clustering) and identifying chamber specificity of individual immature cardiomyocytes (classification). An additional tenet of this thesis is to be able to utilize the entire action potential morphology, instead of just a subset of handcrafted features. For this purpose I have adapted methods from the shape theory community to define a new metric on the space of action potentials that can be used to compare cardiomyocytes at the same or different stages of differentiation (e.g., to compare an action potential of an immature cardiomyocyte to that of a mature cardiomyocyte). Additionally, by considering the shape space induced by this metric, I have developed new methods for comparing cardiomyocytes based on the changes in the action potential morphology as the cardiac environment changes, allowing us to investigate the classification of cardiomyocytes at the functional level. Together, the hope is that this thesis will present a new paradigm for the analysis of stem cell derived cardiomyocytes, and improve high-throughput applications of populations of heart cells. To accomplish this I present work on three tasks:

1. In order to identify whether different populations of action potentials exist in large collections of cardiomyocytes and how they are distributed if they do exist, I develop a scheme that to identify how similar action potential shapes are and how to partition the data based on that similarity, based on technique known as spectral clustering, and a simple shape metric known as the Euclidean distance. By utilizing this algorithm, I show that some cell aggregates present with a single phenotype, while others present with potentially multiple phenotypes

(Chapter 2).

2. To determine whether these partitions lead to clinically relevant, chamber-specific (ventricular, atrial, etc.) phenotypes, I employ a simple classification method known as nearest neighbors to hypothesize the phenotype of an embryonic cardiomyocyte based on mature prototype action potentials. In order to use the whole action potential recording effectively across different levels of maturity, more sophisticated shape distances are required, as the simple Euclidean metric proves less effective at defining similarity between action potentials of different levels of maturity. For this purpose, I will go over deformable template theory and test a more sophisticated shape metric, called metamorphosis, to show whether or not one can identify the chamber specificity of an individual embryonic cardiomyocyte. Additionally, because it is more sophisticated metric, the cost to computing the metamorphosis distance could prove prohibitive in larger datasets. I alleviate this problem by developing an efficient implementation of the metamorphosis algorithm, which allows for one of the optimization variables to be solved in closed form when the others are fixed (Chapter 3).

3. In order to have a more comprehensive classification protocol, I expand from a single action potential recording of a given cardiomyocyte to multiple recordings of the same action potential under different conditions, such as the presence of a drug or the blocking of a specific ion channel. By looking at the changes in

the morphology of the action potential via the metamorphosis interpolation, I can identify behaviors unique to the query cardiomyocyte and potentially indicative of cardiomyocyte phenotype. But in order to do this, I need not only to determine features of the metamorphosis interpolation that can help isolate these behaviors, but a method for making appropriate comparisons of these features, as they are dependent on the starting action potential. The process of unifying this information at a common reference point is known as parallel transport, and can be approximated by a construction known as Schild's Ladder. I will show how to implement Schild's Ladder in the metamorphosis metric space, and how this improves our analysis of the cardiomyocytes based on changes in the action potential morphology. (Chapter 4).

A summary of the thesis and potential future experiments is discussed in Chapter 5.

Chapter 2

Automated Clustering of Action Potentials using the Euclidean Distance

The first step to developing tissues for drug discovery assays and therapeutic technology is assessing the homogeneity of differentiated populations of embryonic cardiomyocytes. In a perfect world, medical and bioengineering researchers would have populations of completely identical cardiomyocytes. Phenotypically identical heart cell tissues could potentially allow cardiologists to isolate the activity of a drug or treatment or make comparative assessments of therapies without having to account with the variability of the individuals in the population. This simplified model of *in vitro* cardiomyocytes could prove immensely helpful in identifying drug toxicities, and

testing new treatments.

However, this does not appear to be that simple in practice. While there are numerous methods for culturing embryonic cardiomyocytes, even within individual methods the resulting cardiomyocytes show varying electrophysiological properties. Now that we can measure the action potential at multiple locations simultaneously via optical mapping, this leads to additional questions regarding how prevalent these variations are from tissue to tissue, how these variations are arranged spatially throughout a tissue, and whether these variations lead to alternative outcomes to a given treatment. The last question will be tackled in Chapter 4, while the other two are the subject of this chapter.

In section 2.1, I will define the problem of clustering and one simple machine learning approach to handling this problem. In section 2.2 I will focus on the method of spectral clustering, a graph based approach to the clustering problem that has some nice robustness and optimization properties as well as how to evaluate the resulting partitions. In section 2.3, I will show how this method can be applied to action potentials. Some tests on a dataset of around 7000 action potentials and a discussion of their results follow.

2.1 K-Means Clustering

At the core of this scientific query is a much simpler question: How does we determine whether one action potential is similar to another action potential, and how does we use this similarity to partition a population of cardiomyocytes into groups based on the action potential? More generally, suppose I have data $\{x_i\}_{i=1}^N$, and would like to find a partitioning of the data so that data points in the same partition are more similar than data in different partitions. I know nothing about the partitions other than the fact that they may exist in the data. When the data can be visualized, as in Figure 2.1 it might be easier to see the partitions, but for higher dimensional data, such as action potentials, this may not be the case. This problem is generally called clustering in the machine learning literature.

One of the simplest ways to do this is via K-means clustering [24]. Given a Euclidean dataset $\{x_i\}_{i=1}^N$ in $\mathbb{R}^D$ and a pre-determined, for the time being, number of clusters, K, the goal of K-means is to identify K prototypes $\{\mu_k\}_{k=1}^K$ such that the sum of the squared distances of all the points dataset to any one of those K prototypes is minimized. Mathematically, let r_{ik} be the indicator of the group of the data point x_i, which takes the value 1 when data point x_i is assigned to the prototype μ_k, and 0 otherwise. With this notation, the K-means clustering algorithm aims to

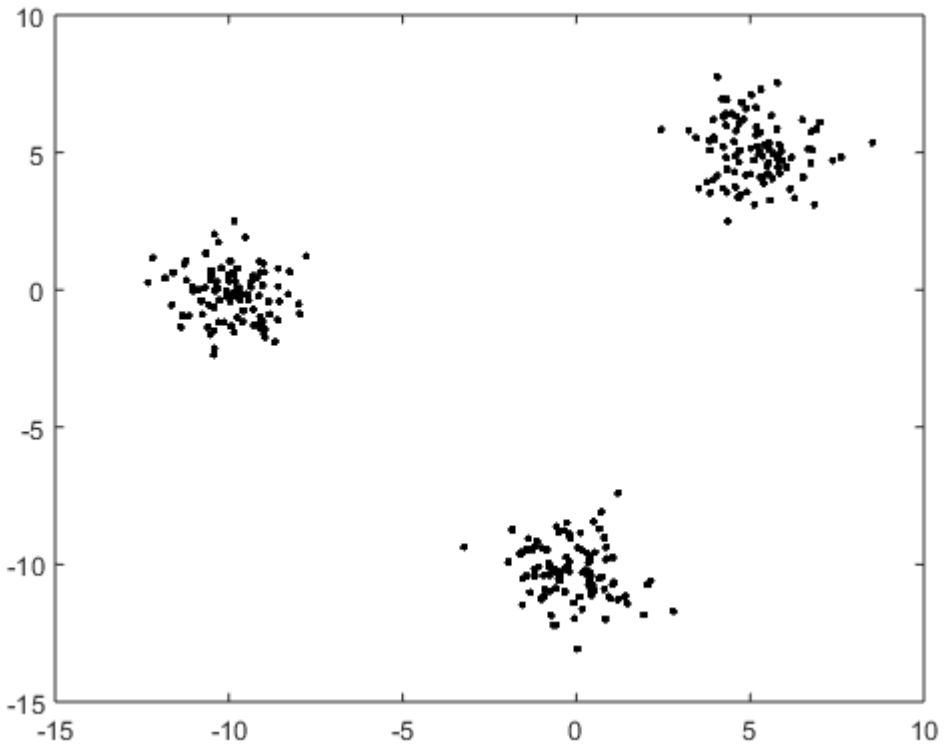

Figure 2.1: Population of data sampled from three Gaussians with different means adequately separated to see the groups clearly.

minimize the following objective function:

$$E(\{r_{ik}\}, \{\mu_k\}) = \sum_{i=1}^{N} \sum_{k=1}^{K} r_{ik} \|x_i - \mu_k\|^2, \tag{2.1}$$

where $\|\cdot\|$ is the standard Euclidean distance: $\|y\|^2 = \sum_{d=1}^{D} y_d^2$, and r_{ik} is constrained to take on the values 0 or 1. By itself, minimizing this function is ill posed, since the obvious choice in this case is to not assign any point to any prototype, or $r_{ik} \equiv 0$, making the function value 0. As a result, we have to add the constraint that every

point has to be assigned to one prototype, or $\sum_{k=1}^{K} r_{ik} = 1, \forall i$, which makes sense since the physical connotation of the problem is that every data point is assigned to one group.

In order to minimize this function E, the standard technique is to alternate between updating the assignments r_{ik} with the prototypes μ_k fixed, and updating the prototypes with the assignments fixed. Specifically, updating the assignments for a specific data point x_i with the prototypes fixed comes down to choosing the prototype with the smallest distance to the data point:

$$r_{ik}^{*} = \begin{cases} 1 & \text{if } k = \arg\min_l \|x_i - \mu_l\|^2 \\ 0 & \text{otherwise} \end{cases}. \tag{2.2}$$

When the assignments are fixed, the resulting update of the prototypes comes down to optimizing a quadratic equation in (2.1). The details are left to the reader, but taking the derivative of the function with respect to one of the prototypes and setting it to zero leads to the updates of the prototypes being the mean of the data points assigned to that prototype, hence the name K-means.

$$\mu_k^{*} = \frac{\sum_i r_{ik} x_i}{\sum_i r_{ik}}. \tag{2.3}$$

Alternating between these two updates decreases the value of the function until a local minimum is reached.

2.2 Spectral Clustering

One of the major limitations to K-means is that it imposes a structure on the partitions of the data that may be too restrictive in practice. Specifically, By using the Euclidean distance to the prototype to define the membership assignments, the resulting clusters are isotropic, or circular in nature, as shown in Figure 2.1. When the data does not follow this distribution, which might be the case with action potentials, or if each group is distributed differently, as in Figure 2.2, utilizing K-means typically ends in failure. As a result, one wants a clustering method that can be more robust to the distribution of the data and find clusters implicitly based on the density of samples. This can be achieved via a method known as spectral clustering. There are multiple versions of the spectral clustering algorithm [25, 26], each with their nuances, but they all operate under the same core tenets. For a complete overview of the spectral clustering algorithm, the reader is referred to the following review [27].

2.2.1 Spectral Clustering Algorithm

Instead of defining a prototype and connecting data points to prototypes, spectral clustering defines connections between the data points directly. This is done via imposing a graph structure on the data. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, w)$ be a weighted graph, comprised of a set of nodes $\mathcal{V}$, a set of edges between nodes $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ and weight function $w : \mathcal{E} \to [0, 1]$. The set of nodes $V = \{1, \ldots, N\}$ is defined so that it

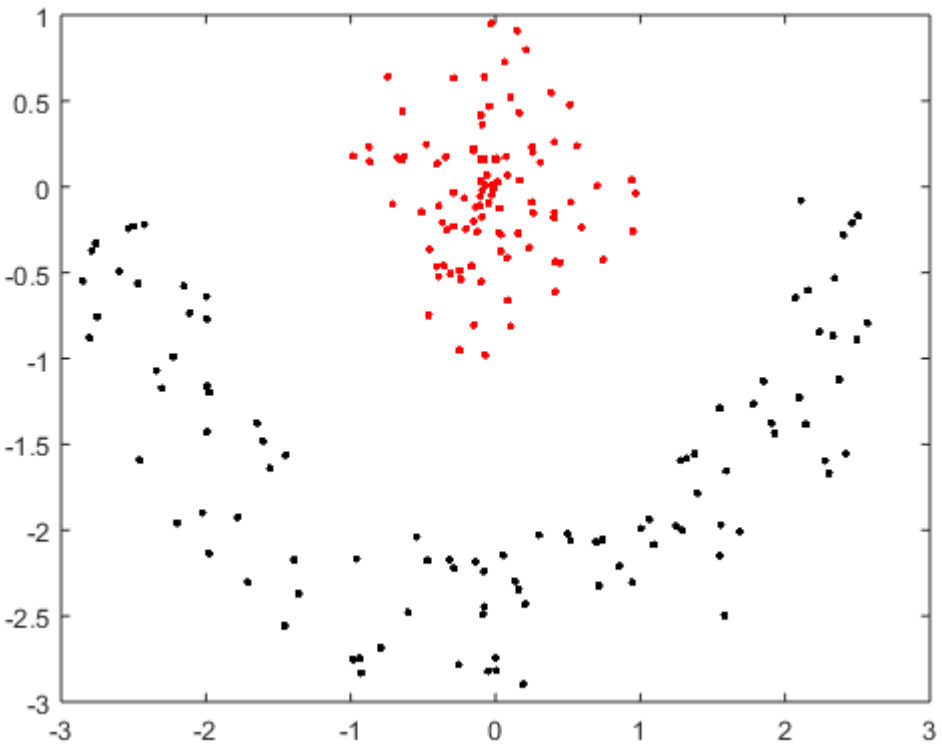

Figure 2.2: Collection of data sampled from two groups with different distributions. The red group is sampled via a Gaussian distribution, while the black group is sampled uniformly over the lower half circle.

corresponds with the set of data points, i.e., node i corresponds to data point x_i. When the data points x_i and x_j are similar to each other, $w(i, j)$ takes on a value closer to 1, and when the data points are not similar to each other, the weight is closer to 0. Since it does not make sense to put an edge between a node and itself, the weight corresponding to such an edge is 0.

In order to explain the spectral clustering algorithm concretely, a toy example of a graph generated from a dataset of 6 points is presented in figure 2.3. Assume, for

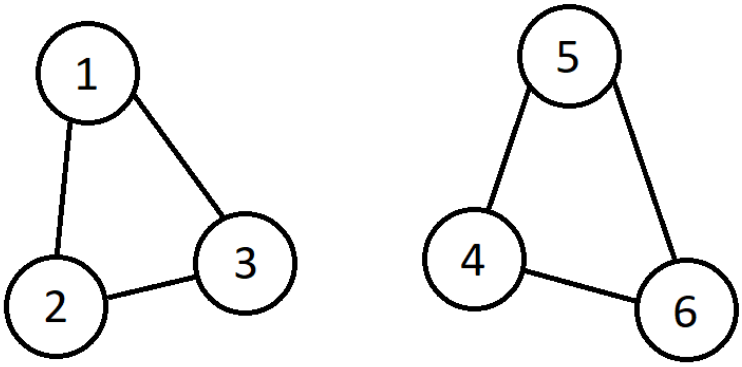

Figure 2.3: Toy graphical model generated from a population of 6 data points to explain spectral clustering.

the time being, that the connections drawn are full weight ($w = 1$) connections, and that there are no connections other than the ones drawn ($w = 0$). Let us define the matrix of weights $W_{ij} = w(i, j)$ for this toy example as:

$$W = \begin{pmatrix} 0 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 \end{pmatrix}. \tag{2.4}$$

How does W help us identify the groups of connected nodes? In order to continue, we need to define a second matrix. Let D be the degree matrix of the graph $\mathcal{G}$. D

is a diagonal matrix, and the value of the diagonal d_{ii} is the sum of the weights of all edges connected to node i. This can be computed by summing the weights on the rows of W ($d_{ii} = \sum_j w(i,j)$). For our toy example, the degrees of the six nodes are all 2.

In order to create K groups of nodes, we have to cut the graph into connected components by cutting the edges. Ideally we cut where the weights are the smallest. This is simple in our toy example because there is already no connection between the two groups of nodes. But in general, how do we define this penalty? Suppose we have a partition $\{\mathcal{V}_1, \mathcal{V}_2, \ldots \mathcal{V}_K\}$ of the nodes. Each of the nodes in the graph belongs to exactly one of the partitions. Then we can define a "cut" cost on a partition by looking at the weights that separate each group from the rest of the data:

$$\text{Cut}(\mathcal{V}_1, \ldots, \mathcal{V}_K) = \sum_{k=1}^{K} \sum_{i \in \mathcal{V}_k, j \notin \mathcal{V}_k} w(i,j). \tag{2.5}$$

Let us first look at the case where $K = 2$, and let f_i be the indicator function for whether node i is in group $\mathcal{V}_1$. Specifically, $f_i = 1$ when node i is in the group, and $f_i = 0$ otherwise. Given this, another way to write equation (2.5) is the following:

$$\text{Cut}(\mathcal{V}_1, \mathcal{V}_2) = \frac{1}{2} \sum_{i,j} w(i,j)(f_i - f_j)^2. \tag{2.6}$$

Expanding the quadratic and rearranging terms leads to the following result:

$$\text{Cut}(\mathcal{V}_1, \mathcal{V}_2) = \mathbf{f}^T (D - W)\mathbf{f}, \tag{2.7}$$

where $\mathbf{f}$ is the column vector made from concatenating all of the f_i's together. The matrix $L = D - W$ is typically called the Laplacian of the graph $\mathcal{G}$.

Finding the minimal cut cost partition is the equivalent to solving the following optimization problem:

$$\min_{\mathcal{V}_1, \mathcal{V}_2} \text{Cut}(\mathcal{V}_1, \mathcal{V}_2) = \min_{\substack{\mathbf{f} \\ f_i \in \{0,1\} \\ \mathbf{f} \neq \mathbf{0}, \mathbf{f} \neq \mathbf{1}}} \mathbf{f}^T L \mathbf{f}, \tag{2.8}$$

where $\mathbf{0}$, and $\mathbf{1}$ are the vectors of all 0's and 1's, respectively, and the additional two constraints are added to ensure that two groups are created. Obtaining the minimum of this problem is difficult given the combinatorial nature of the search space. One can find an approximate minimum by relaxing this constraint as follows:

$$\min_{\mathcal{V}_1, \mathcal{V}_2} \text{Cut}(\mathcal{V}_1, \mathcal{V}_2) \approx \min_{\substack{\mathbf{f} \\ \|\mathbf{f}\|_2^2 = 1 \\ \mathbf{f} \perp \mathbf{1}}} \mathbf{f}^T L \mathbf{f}. \tag{2.9}$$

L has some very nice properties. In particular, L is symmetric, as W, D are symmetric, and L is positive semi-definite, since $\mathbf{f}^T L \mathbf{f} = \frac{1}{2} \sum_{i,j} w(i,j)(f_i - f_j)^2 \geq 0$. Additionally, this is a well studied optimization problem in the linear algebra community. According to the Rayleigh-Ritz theorem, the minimizer $\mathbf{f}$ of (2.9) is the eigenvector

of L corresponding to its second smallest eigenvalue, as $\mathbf{1}$ is the eigenvector related to the eigenvalue 0, the smallest possible eigenvector since L is positive semi-definite.

Returning to the toy example (Figure 2.3), let us look at the graph Laplacian:

$$L = \begin{pmatrix} 2 & -1 & -1 & 0 & 0 & 0 \\ -1 & 2 & -1 & 0 & 0 & 0 \\ -1 & -1 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 & -1 & -1 \\ 0 & 0 & 0 & -1 & 2 & -1 \\ 0 & 0 & 0 & -1 & -1 & 2 \end{pmatrix}. \tag{2.10}$$

There are two interesting things about this L. First, 0 is an eigenvalue with multiplicity 2. To see the second important feature of L, let us look at the eigenvectors corresponding to the 0 eigenvalue. The first one is the vector of all ones, $\mathbf{1}$. The second one is $(1, 1, 1, -1, -1, -1)^T$. Notice that this vector identifies the membership of the two connected components in the graph, and thus we can separate the graph into the two groups easily.

The more interesting case is when signals are close to, but not exactly, connected components. To see this in the toy example, an edge is added between nodes 3 and 4, but this edge is made weaker than the other edges by reducing the weight. The modified graph is shown in Figure 2.4, where the dotted line reflects the weaker edge.

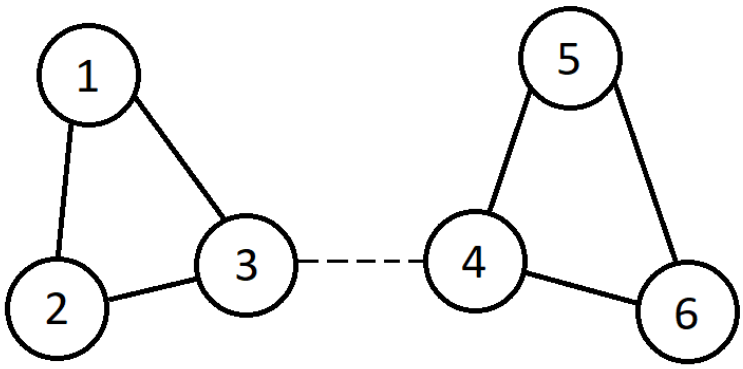

Figure 2.4: Modified Toy graphical model generated from a population of 6 data points. The dotted line indicates a weaker edge than the solid lines.

Specfically, let $w(3,4) = w(4,3) = 0.05$. The resulting Laplacian becomes:

$$L = \begin{pmatrix} 2 & -1 & -1 & 0 & 0 & 0 \\ -1 & 2 & -1 & 0 & 0 & 0 \\ -1 & -1 & 2.05 & -0.05 & 0 & 0 \\ 0 & 0 & -0.05 & 2.05 & -1 & -1 \\ 0 & 0 & 0 & -1 & 2 & -1 \\ 0 & 0 & 0 & -1 & -1 & 2 \end{pmatrix}. \tag{2.11}$$

Now 0 is only an eigenvector of multiplicity 1, with the eigenvector of all ones. So how do we identify the clusters? Let us look at the eigenvector of the second smallest eigenvalue (0.0326):

$$v_{0.0326} = (1, 1, 0.9674, -0.9674, -1, -1)^T. \tag{2.12}$$

While not the same as the previous example, it still roughly discriminates the two clusters. In fact, if one performs K-means clustering with $K = 2$ on this vector, the resulting clusters will be the two groups from the disconnected example. This is what spectral clustering does when the clusters are not completely disconnected. In order to extend this to $K > 2$ groups, additional assignment vectors are required. Instead of looking at one vector $\mathbf{f}$, a matrix $\mathbf{F}$ is generated by concatenating the $K - 1$ assignment vectors together. Following the same logic as in (2.9), one can approximate the minimizer $\mathbf{F}$ by finding the $K - 1$ smallest eigenvectors of the graph Laplacian L.

The problem with the cut cost in (2.5) is that it favors making a partition of small size, because that means fewer edges have to be cut. As a result, it is not a robust cut cost in the presence of outliers. In order to be more robust to this, other definitions of the cost of cutting the graph have been proposed. One approach is to scale the cuts by the total weight of the partition:

$$\mathrm{NCut}(\mathcal{V}_1, \ldots, \mathcal{V}_K) = \sum_{k=1}^{K} \frac{\sum_{i \in \mathcal{V}_k, j \notin \mathcal{V}_k} w(i,j)}{\sum_{i \in \mathcal{V}_k, m \in \mathcal{V}} w(i,m)}. \tag{2.13}$$

This avoids the problem of favoring small groups in a given partition because the numerator gets closer to the denominator as the groups get smaller, i.e., more of the total number of edges of a partition are separating edges, thus increasing the cost of the partition. At the extreme, when every data point has its own group, the

numerator and the denominator are identical, and the value of NCut is equal to the number of data points. The details are omitted, but one can find an approximation similar to (2.9) for (2.13) with a modified version of L. By replacing L with the normalized Laplacian $L_{sym} = I - D^{-\frac{1}{2}}WD^{-\frac{1}{2}}$, one can find a good approximation to the minimal NCut (2.13), as described in [27]. The complete spectral clustering algorithm is outlined in Algorithm 1.

Algorithm 1 Spectral Clustering Algorithm

Input: Weighted Graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, w)$, choice of Laplacian: unnormalized or normalized, number of clusters K.

Compute Weight matrix $W : W_{ij} = w(i,j), i,j, \in \{1,\ldots,|\mathcal{V}|\}$
Compute Degree matrix $D : D_{ii} = \sum_{j=1}^{|\mathcal{V}|} W_{ij}$
Compute Laplacian L:
if Unnormalized **then**
    $L = D - W$
else
    $L = Id - D^{-1/2}WD^{1/2}$
end if
Find the K eigenvectors v_1, v_2, $\ldots$, v_K corresponding to the K smallest eigenvalues $\lambda_1 < \lambda_2 < \ldots < \lambda_K$ of L. $\mathbf{V} = [v_1 v_2 \ldots v_K]$
$x_i = \mathbf{V}_{i\cdot}, i \in \{1,\ldots,|\mathcal{V}|\}$
if Normalized **then**
    $x_i \leftarrow \frac{x_i}{\|x_i\|}$
end if
Perform K-means on $\{x_i\}_{i=1}^{|\mathcal{V}|}$. Assign node i to group k if $r_{ik} = 1$.

Output: Assignments r_{ik}

2.2.2 Model Selection

Now that I can approximate the optimal partition of the data into K groups, it is important to be able to determine which K provides the best representation of the

data. In order to do that, a measure of the quality of a partition as a function of K is needed. Two measures are described below.

1) Normalized Cut Cost (NCC): The first measure that I consider is the value of the objective function of the NCut clustering method in (2.13) divided by the number of groups K:

$$\mathrm{NCC}(\mathcal{V}_1, \ldots, \mathcal{V}_K) = \frac{1}{K} \sum_{k=1}^{K} \frac{\sum_{i \in \mathcal{V}_k, j \notin \mathcal{V}_k} w(i,j)}{\sum_{i \in \mathcal{V}_k, m \in \mathcal{V}} w(i,m)}. \tag{2.14}$$

When the data is well separated into K clusters, the numerator will be very small, while the denominator will be large, because signal pairs of the same group will be well connected while signal pairs of different groups will be poorly connected. Therefore we expect the value of NCC to be small when K is properly chosen. When K is overestimated, the numerator increases and the denominator decreases, while when K is underestimated the opposite happens, favoring a small choice of K Dividing the ratio by K normalizes the cut costs across different values of K and makes them comparable. The optimal number of groups K^* will come from the minimal NCC.

2) Davies-Bouldin Index (DBI): The other measure we consider is a cluster distance between pairs of groups [28, 29]. Let $\mu_k(\tau)$ be the mean of the data points in group $\mathcal{V}_k$:

$$\mu_k(\tau) = \frac{1}{|\mathcal{V}_k|} \sum_{n_i \in \mathcal{V}_k} f_i(\tau). \tag{2.15}$$

Given two groups, $\mathcal{V}_k$ and $\mathcal{V}_{k'}$, the DBI distance between the groups is defined as

follows:

$$DBI(k, k') = \frac{S_k + S_{k'}}{M_{kk'}}, \tag{2.16}$$

where S_k and $S_{k'}$ are values for the dispersion or spread within each group, and $M_{kk'}$ is a measure between groups. For our work, S_k is defined as the average of the distances of the signals within one group to the average signal of that group:

$$S_k = \frac{1}{|\mathcal{V}_k|} \sum_{n_i \in \mathcal{V}_k} d(f_i(\tau), \mu_k(\tau)), \tag{2.17}$$

while the measure $M_{kk'}$ is the distance between the average signals of the two groups:

$$M_{kk'} = d(\mu_k(\tau), \mu_{k'}(\tau)). \tag{2.18}$$

For well separated, tight clusters, the dispersion for each group is very small, since the distances to the mean will be small, while the measure $M_{kk'}$ between groups will be large because the means are well separated from each other. Therefore, a lower DBI indicates better grouping. For K groups, the average DBI over all pairs of groups is chosen to be the group measure:

$$\overline{DBI} = \frac{1}{|\mathcal{K}|} \sum_{\mathcal{K}} DBI(k, k'), \tag{2.19}$$

where $\mathcal{K} = \{(k, k') : 1 \leq k < k' \leq K\}$.

2.3 Clustering of Cardiomyocyte Action Potentials

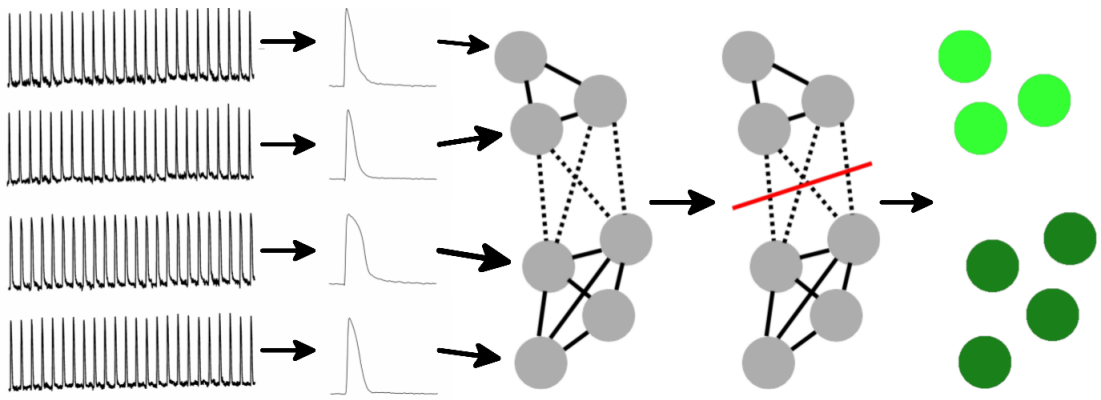

(a) Raw signals (b) Sample APs (c) Graph building (d) Graph cut (e) Final groups

Figure 2.5: Schematic drawing of the proposed method. (a) Action potential recordings are obtained from cardiomyocytes, either via optical mapping or electrode recording. (b) Representative action potentials are identified from these recordings. (c) Each action potential forms a node in a graph, and weighted edges are built based on the similarity between action potentials. (d) The graph is cut into groups based on where the edges are weakest. (e) The separated connected components form the groups of similar action potentials.

In this section, I show how the spectral clustering approach described in the previous section can be applied to the problem of clustering action potentials. The overall framework is illustrated in Figure 2.5. Let $f(\tau)$ be the voltage of an action potential derived from a cardiomyocyte. Let us define the domain of this time series signal as $[0, T]$, where T is the length of one beat cycle of the action potential. Although beat to beat variations exist in nature, including alternans behavior is apparent at shorter cycle lengths [30], in this thesis, we are going to operate under

the assumption that the action potential is identical from beat to beat. Thus $f(\tau)$ is a periodic function on the set $[0, T]$ Let $\{f_i\}_{i=1}^{P}$ be the set of action potentials in a population of P cardiomyocytes, and let K be the number of groups we hope to separate the data into.

As we have discussed already, we want to define a similarity $w(i, j)$ between any pair $f_i(\tau), f_j(\tau)$ of action potential signals. We do this indirectly by identifying in which ways they are different. A very simple way to do this is with the squared Euclidean distance: $d^2(f_i, f_j) = \int_0^T (f_i(\tau) - f_j(\tau))^2 d\tau$. When the f_i and f_j are identical, the distance is 0, and the distance grows as the two signals become more disparate from each other. Since we are looking at discrete samples of the action potential signal, this can be re-written as $d^2(f_i, f_j) = \sum_{m=1}^{M} (f_i(\tau_m) - f_j(\tau_m))^2$, where $0 < \tau_1 < \tau_2 < \cdots < \tau_m < T$ denote the sampling times. Given the distance d, a typical way to define the weight $w(i, j)$ between the nodes corresponding to $f_i(\tau)$ and $f_j(\tau)$ is to employ the Gaussian, or radial basis function, kernel:

$$w(i, j) = e^{(-\frac{d^2(f_i, f_j}{\sigma^2})}. \tag{2.20}$$

By using this kernel, we introduce an additional parameter, σ^2. This parameter helps determine how tightly the clusters are bound. When σ^2 is large, the denominator overpowers the numerator inside the exponential, pushing the weights closer to $e^0 = 1$, encouraging the data to aggregate into one group because the graph is fully connected

with weights close to 1. Conversely, when σ^2 is small, weights tend towards 0, as $\lim_{x\to\infty} e^{-x} = 0$. As a result, the data is much less connected, and spectral clustering will be encouraged to separate every data point into its own group. Thus, it is important to find a balanced σ^2 that prevents over and under partitioning. Since we are looking at the Euclidean distances between all pairs of action potentials in the dataset, an appropriate choice for σ^2 is the average distance between pairs:

$$\sigma^2 = \frac{1}{|\mathcal{E}|} \sum_{(i,j)\in\mathcal{E}} d^2(x_i, x_j), \tag{2.21}$$

where $|\mathcal{E}|$ is the number of edges in $\mathcal{G}$. Following the spectral clustering convention, there is no connection between a node and itself, so even though the distance between an action potential and itself is 0, we set $w(i,i) = 0$. Having defined the weights between nodes, we proceed with the spectral clustering algorithm 1.

2.4 Experiments and Results

The spectral clustering algorithm described in 2.3 was applied to a dataset of 6940 action potentials derived from individual cardiomyocytes from 9 different cell bodies via voltage sensitive dyes and optical mapping. In particular, the H9 line of human embryonic stem cells was differentiated into aggregates of cardiomyocytes known as embryoid bodies (hEBs) using a previously described hEB-based protocol [17]. The hEBs usually start beating around Day 9 after the initiation of differentiation. Beat-

ing areas of hEBs were mechanically dissected on Day 15-16 of differentiation, and plated on gelatin-coated plastic coverslips as cardiac cell clusters for optical mapping. Cardiac cell clusters were stained with $10\mu M$ voltage-sensitive dye di-4-ANEPPS (Invitrogen, Grand Island, NY), and $50\mu M$ myosin II inhibitor blebbistatin (Sigma-Aldrich, St. Louis, MO) was applied throughout experiments to inhibit motion. Action potentials were recorded using a MiCAM Ultima-L CMOS camera (100×100 pixels, $16\mu m$/pixel) at 500 frames per second (fps). A pair of platinum electrodes was used to deliver fixed 90 beat per minute (bpm) pacing to the clusters, and 16 second recordings containing multiple APs were obtained from each cluster. Figure 2.5(a) illustrates the recordings obtained with this protocol. Note that the resulting optical recordings are normalized so that the resting membrane potential is 0 and the maximum amplitude is 1.

The result of the optical mapping is an action potential recording at pixels of an image corresponding to the cell body. In particular, each pixel represents $16\mu m \times 16\mu m$ area of the cellular aggregate. Let $I(x, \tau)$ be the optical recording data at a specific pixel x. Ideally, the collection of these recordings over all of the aggregates of cells would make up our dataset $\{f_i(\tau)\}_i$. However, there is still significant noise in the optical mapping recording. In order to improve the signal-to-noise ratio, the follow pre-processing steps are performed. First, we smooth out the action potential recording across the cell body by averaging spatially across an aggregate. Cardiomyocytes located near each other are likely to have similar action potentials, so this

helps to magnify the signal quality of a particular pixel. To do this, we implemented a 5×5 boxcar averaging filter $g(x)$, and convolved this filter with the optical recording images:

$$I^{\bar{x}}(x, \tau) = g(x) * I(x, \tau). \tag{2.22}$$

Additionally, instead of looking at the entire 16 second recording of 24 cycles of the action potential, a representative action potential is obtained by averaging across the 24 cycles temporally:

$$I^{\overline{x\tau}}(x, \tau) = \frac{1}{24} \sum_{j=1}^{24} I^{\bar{x}}(x, \tau + (j-1)T), \tag{2.23}$$

with T being the cycle length of one action potential, which is $\frac{2}{3}$ s since the data is paced at 90 bpm. Both of these pre-processing steps improve the signal-to-noise ratio and make the action potential recordings amenable for analysis. However, given the conduction velocity of the embryoid body, the signals may not be completely aligned with each other. The APs from different averaged pixels were aligned by their activation time, which was calculated to be the time point with maximal upstroke slope:

$$I^{\overline{x\tau}a}(x, \tau) = I^{\overline{x\tau}}(x, \tau - \overset{*}{\tau}(x)), \tag{2.24}$$

where

$$\overset{*}{\tau}(x) = \underset{\tau \in [0,T)}{\arg\max} \frac{\partial I^{\overline{x\tau}}(x, \tau)}{\partial \tau}. \tag{2.25}$$

This alignment allows us to compare signals from different regions of the cell cluster that may be offset due to differences in activation time owing to electrical conduction of the AP wavefront through the cell cluster. Alignment is done post spatial averaging because the increased signal-to-noise ratio allows us to determine a more accurate $\overset{*}{\tau}(x)$. Experimentally we have found that the local misalignment is negligible compared to misalignment from different regions of the cell cluster. These aligned APs $\{I^{\overline{x\tau}a}(x,\tau)\}_x$ serve as our data elements $\{f_i(\tau)\}_i$ for the spectral clustering framework. Figure 2.5(b) illustrates typical representatives obtained by this method.

Before we look at the resulting partitions, let us take a look at the weight matrices generated by the data, to obtain some intuition about what should happen. Figure 2.6(a) shows the weight matrices for a subset of the populations. The silhouettes of the 9 cell bodies (gray regions) are shown on the top and rotated 90 degrees on the left. The subset of pixels being analyzed from each cell body is marked by the blue lines crossing each cluster. The similarity between the signals at pixel i in one and pixel j in another line is measured by the weight $w(i,j)$. Red indicates high weight and blue indicates low weight. Weights along the diagonal are 0 by convention of spectral grouping. In particular, let us investigate specific pairs of cell aggregates. Figure 2.6(b) details the AP similarity for a subset of cell aggregates 1 and 9. Since the weights relating these two clusters are blue (see the top right and bottom left areas of the submatrix), they should be in different groups. Figure 2.6(c) shows the AP similarity for a subset of cell aggregates 2 and 3. In contrast with 2.6(b), there are

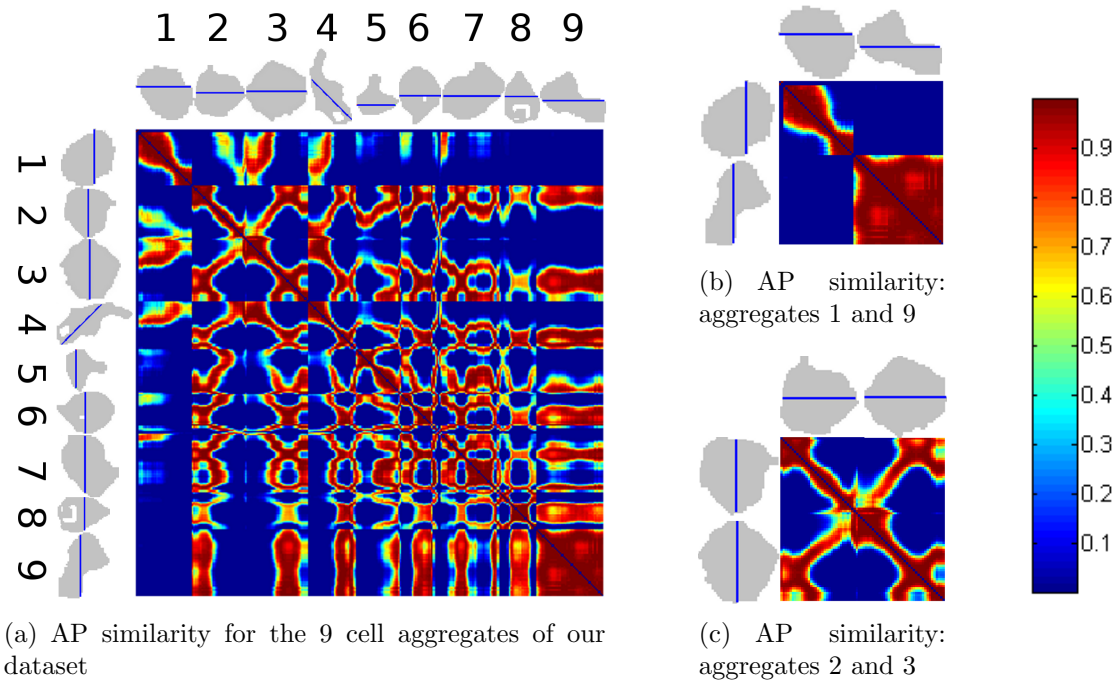

(a) AP similarity for the 9 cell aggregates of our dataset

(b) AP similarity: aggregates 1 and 9

(c) AP similarity: aggregates 2 and 3

Figure 2.6: Observed weight matrices for a subset of the hEB data. (a) Selected pixels from each cell body are indicated with the blue line. Weights are calculated using the Gaussian kernel 2.20 of Euclidean distance between the representative action potentials at the pixel locations. Recordings with high weight connections are indicated in red, low weight connections are indicated in blue. (b) Subset of the weight matrix isolating aggregates 1 and 9. The aggregates show high weight with among APs in the same aggregate, but low weight when compared to APs from the other cell body. (c) Subset of the weight matrix isolating aggregates 2 and 3. The left region of aggregate 2 shows low affinity for the right region, but high affinity for the right region of aggregate 3.

high similarities between the left part of aggregate 2 and the right part of aggregate 3 and vice versa, as well as low similarity between the left and right parts of each of the aggregates individually. The left side of cell aggregate 2 should group with the right side of cell aggregate 3 and the right side of aggregate 2 should group with the left side of aggregate 3.

Spectral clustering is performed to partition the dataset into 2, 3, and 4 groups.

(a) 2 Groups: NCC = 0.1457, DBI = 1.1655

(b) 3 Groups: NCC = 0.2756, DBI = 1.3203

(c) 4 Groups: NCC = 0.3883, DBI = 1.5863

Figure 2.7: Clustering visualization. The 9 cell aggregates are separated into 2, 3, and 4 groups and the group assignments are indicated by the color at the pixel location. Scale bars indicate 200 μm. The average action potential recording of each group of each of the partitions is shown on the right. Scale bar indicates 100 ms. Fitness is assessed by the NCC (2.14) and DBI (2.19) values.

The visualization of the clustering algorithm is presented in Figure 2.7. The first 9 columns show the results of applying spectral clustering to all the APs in our dataset for 2 (top row), 3 (middle row), and 4 (bottom row) groups with corresponding NCC and DBI measures listed underneath. Some cell clusters present with primarily one phenotype, while others present a mixture of phenotypes. Both NCC and DBI suggest that $K = 2$ gives the best grouping fitness. The average AP of each phenotype as determined by spectral grouping is shown on the last column. The average 2-group APs suggest different phenotypes, while for 3 and 4 groups, pairs of phenotypes (phenotypes 2 and 3 for 3 groups, phenotypes 1 and 2 and phenotypes 3 and 4 for 4 groups) have similar shapes.

To further evaluate the characteristics of the obtained groups, we also computed

the distributions of commonly used AP features [10, 16, 31, 32]. In particular, Figure 2.8 shows the distributions of APD_{30}, APD_{80}, triangulation ($APD_{90} - APD_{30}$), and normalized triangulation ($\frac{APD_{90}-APD_{30}}{APD_{90}}$) for the groups determined by spectral grouping and visualized in Figure 2.7. As a reminder, APD_x stands for action potential duration, which is the time from depolarization to $x\%$ repolarization (see Section 1.2.3). For the case of 2 groups, APD_{30} and APD_{80} histograms were fitted well as distinctly separated sub-populations. However, triangulation, a parameter indicative of a pro-arrhythmic substrate when prolonged, had a wide range in both groups, with substantial overlap. By normalizing triangulation to APD_{90}, we obtained a shape parameter that showed better discrimination of AP morphology between the two groups, and that the group with shorter APD_{30} and APD_{80} had greater triangulation. With increased number in groups (3 and 4 groups), we observed increased overlap in all AP features among different groups, especially in normalized triangulation. This increased overlap suggests that APs in different groups are similar in shape, thus increasing difficulty in discriminating the groups by such features. Overall, these results suggest that grouping the population based on standard AP features would be possible for 2 groups, while using the entire AP waveform is effective for 2, 3, or 4 groups.

There are lots of nice takeaways from this experiment. First, the weight matrix used for spectral clustering analysis also provides visualization of the similarity among APs within the dataset. The existence of both similar and dissimilar regions within

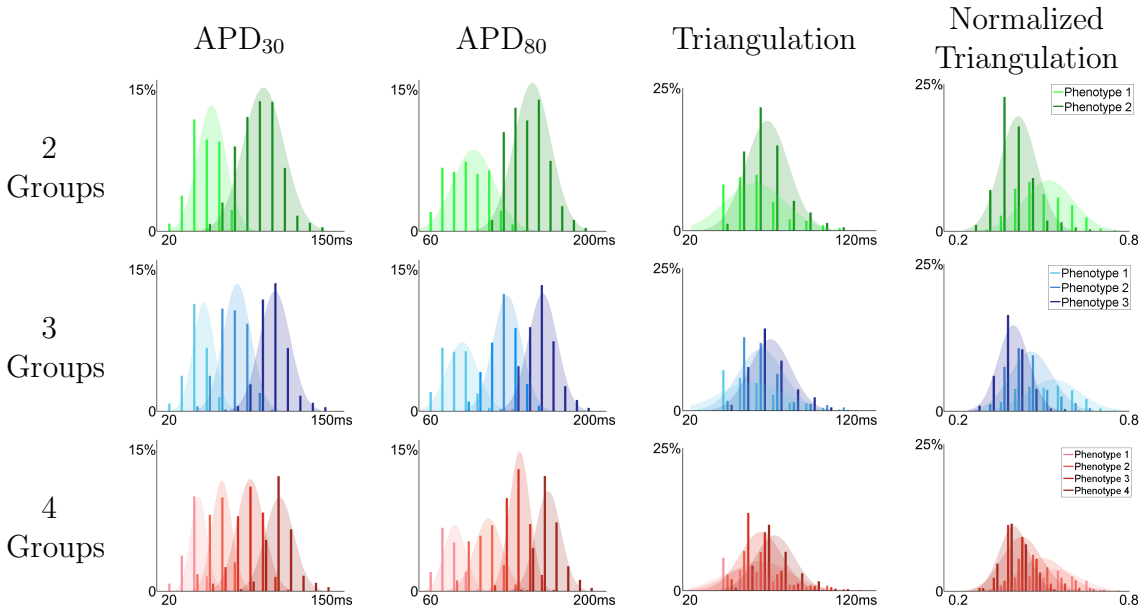

Figure 2.8: Histograms of standard biologist action potential features APD_{30}, APD_{80}, Triangulation ($\mathrm{APD}_{90} - \mathrm{APD}_{30}$), and normalized triangulation ($\frac{\mathrm{APD}_{90}-\mathrm{APD}_{30}}{\mathrm{APD}_{90}}$) (left to right) with respect to the 2, 3, and 4 (top to bottom) groups formed by the spectral clustering algorithm. APD_x is the action potential duration, the time from depolarization to $x\%$ repolarization.

some cardiac cell aggregates indicates that more than one phenotype may coexist within a single cluster, consistent with previous work [33]. Our results show that while there were cell aggregates in our dataset that expressed primarily one phenotype, the majority of the cell aggregates presented with multiple phenotypes. For the aggregates with more than one phenotype, even though the clustering method did not enforce spatial regularity, continuous regions of cell aggregates with only one phenotype were obtained. Further, we observed smoothly varying AP shape across the boundary separating phenotypes. This is partly because of the spatial averaging performed during preprocessing, but it primarily reflects that APs vary as a continuum from one phenotype to another. Thus, it is particularly important to develop

automated algorithms to maintain a consistent decision boundary across datasets. Perhaps more importantly, a major feature of the proposed framework is that the pieces are generalizable. We used optical mapping recordings for our dataset, but because the spectral grouping and fitness evaluation operate with a processed action potential signal, they are amenable to action potential recordings obtained from other techniques, and thus could be used to investigate variability in current methods for purification of specific phenotypes of hESC-CMs [6,10,34]. Additionally, more sophisticated measures other than the Euclidean distance can be implemented. One such measure will be discussed in Chapter 3. By defining better measures of similarity between action potentials, we can be more confident in the groups that are established via this method and identify more complex phenomena, such as sub-phenotype variability.

Chapter 3

Automated Classification of Action Potentials Using the Metamorphosis Distance

In the previous chapter, we observed how the Euclidean distance can be utilized to separate populations of embryonic cardiomyocytes into subpopulations based on the action potential morphology. A natural follow-up question is to ask whether or not these subpopulations are representative of separate phenotypes of cardiomyocytes. In particular, are these subpopulations indeed the precursors to chamber specific (atrial, ventricular, or pacemaking nodal) adult cardiomyocytes. In this chapter, we will discuss the challenges in answering this question in the action potential domain and introduce the metamorphosis metric as a way to address these challenges.

3.1 Action Potential-based Cardiomyocyte Classification

The problem of identifying the phenotype, or label, of an action potential, or any query data, is known in the machine learning community as classification. More specifically, given a collection of data $\{x_i\}_{i=1}^N \in \mathcal{X}$ and their corresponding labels $\{y_i\}_{i=1}^N \in \mathcal{Y}$,the goal of classification is to learn a classifier $c : \mathcal{X} \to \mathcal{Y}$ that minimizes a defined loss, or error, $l(c(x_i), y_i)$ of assigning each data point a label. In the ideal scenario, the classifier c not only correctly assigns the correct label to every data point in the collection, but also robustly extends to new, unseen data points. Examples of typical classification methods include support vector machines [35, 36], random forests [37], and, more recently, neural networks and deep learning [38, 39].

Defining Exemplars: In order to identify the phenotype of a query cardiomyocyte action potential, we need to first develop a database of action potentials for which the phenotype is known. In a perfect world, this would be a database of embryonic action potentials at the same maturation point as the cardiomyocyte in question. However, as we have discussed already, most works with phenotypically identified embryonic action potentials have developed that distinction subjectively via manual annotation, and do not readily transfer between studies and protocols.

Another approach is to find well classified data in an alternative, but related, domain, and hope to be able to utilize this data to infer something about the dataset in

question. Where can we find such a related dataset? Over the last 2 decades, concurrent with the growing experimental analysis of cardiomyocytes, the use of computational models for *in silico* analysis of action potential generation has shown increasing popularity. Evolving from the original work of Hodgkin and Huxley [40], these models attempt to replicate mathematically the cellular processes of the cardiomyocyte that lead to action potential generation. For adult atrial [41–43], ventricular [44, 45], and pacemaking nodal [46,47] phenotypes, individual systems of differential equations have been developed that show similar properties to physical cardiomyocytes with regards to pacing protocol, voltage-current recordings, and drug implementations. We can produce a population of variable action potentials of a particular phenotype by varying the ion concentrations and conductances of the models within a specified tolerance. Ideally we would utilize computational models for embryonic cardiomyocytes, but while computational models for embryonic cardiomyocytes exist [48], they are built on heuristics and have not been experimentally verified. Example outputs of the adult atrial models of [41] and the adult ventricular models of [45] are shown in Figure 3.1. Visually, while there is variability between the models of a specific phenotype, there's still a greater distinction between action potentials belonging to models of different phenotypes. Further, since these models were developed with the specific phenotype in mind, we have no ambiguity as to their labeling.

Given this labeled population of adult action potentials, how do we use it to infer the phenotype of an unknown embryonic action potential? Since this an embryonic

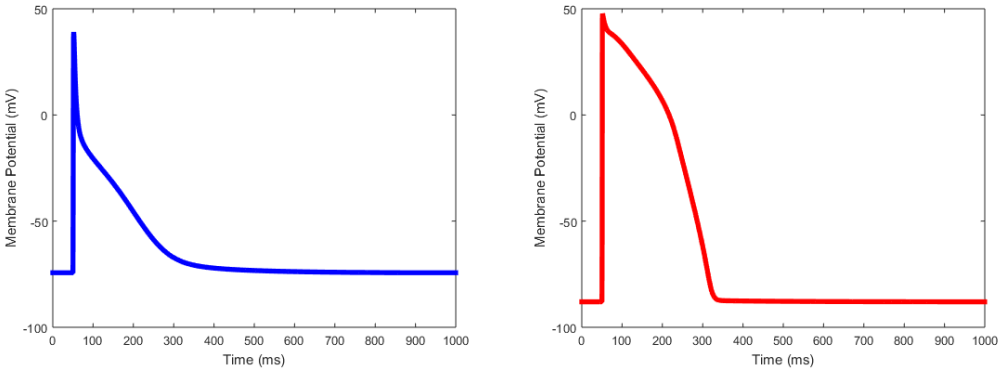

Figure 3.1: Prototypical Mature Atrial (left) and Mature Ventricular (right) action potential

heart cell, it will become like one of these mature action potentials, given the time to mature. We can think of the maturation process as producing an interpolation between the embryonic and mature action potentials, where consecutive interpolants are related by a small deformation. A reasonable hypothesis, then, is that an unknown embryonic action potential requires less energy to be deformed into its correct phenotype than alternative phenotypes. One way to measure this energy is the minimal amount of deformation required to warp the query action potential to a mature action potential of a particular phenotype. As we will see, the minimal amount of deformation provides a distance, that can be approximated by calculating the distances between the query and all of the exemplar prototypes, and picking the phenotype based on the phenotype of the cell that has the shortest distance to the test embryonic action potential. This classification process is commonly referred to as nearest neighbors. Since assigning a label based on one nearest neighbor is sensitive to outliers and the sampling of the training data spaces (e.g. one spurious outlier being

closer to the embryonic test action potential could alter the result of nearest neighbors), a more robust approach is to identify K neighbors, and then assign the test signal based on the mode of the labels of the neighbors. This classification process is commonly referred to as K-Nearest Neighbors (K-NN).

Regardless of the choice of number of neighbors, this methodology requires the choice of an appropriate distance between action potentials. One choice is the Euclidean distance we utilized to define the similarity between action potentials in chapter 2. However, this proves to be a poor selection of the distance in this problem. Recall Figure 1.6. Since the Euclidean distance is a point-wise measure between the two signals an intermediate signal is based on the difference between the two signals at any given point, it suggests that the query signal in green is closer to the atrial signal in blue than the ventricular signal in red. As a result, using nearest neighbors, the Euclidean distance would suggest that the query action potential is atrial, even though its shape more closely resembles the ventricular action potential. Additionally, we show the Euclidean interpolation between an embryonic and mature action potential in Figure 3.2. Notice that the Euclidean interpolation creates a bump along the repolarization phase of the action potential sequence. This behavior is not indicative of action potential morphology, and appears because we are using the Euclidean distance. This spurs us on to find an alternative choice for a measure between signals that, in particular, preserves the shape properties of the action potential.

Figure 3.2: Example Euclidean Interpolation of an embryonic to mature action potential

3.2 Continuous Metamorphosis

3.2.1 Basics of Deformable Templates

Given the variations observed in action potential morphology, a better approach is to consider the whole shape as a member of the family of shapes indicative of action potentials, i.e. as a member of a shape space, and learn how to define a metric on this space. The notions of shape spaces and shape theory are usually referred back to Kendall [49]. In that seminal work, he proposed the notion of a "shape" as what remains after we remove certain families of reversible transformations, such as translations, rotations, and scaling. Most interestingly, the resulting quotient space of shapes after removing these transformations can be analyzed to identify the difference between shapes via Riemannian geometry.

What if we want to extend the set of allowable transformations from rotations and translations to more complicated phenomena? D'Arcy Thompson suggested many centuries ago that related species of animals share form that can be connected via a smooth deformation of the key grid points [50]. The typical example is that of the fish shown in Figure 3.3. Given a reference fish, one is able to obtain similar, related

species of fish by applying an appropriate deformation on the reference. Grenander formalized these notions mathematically by defining the space of a particular pattern or shape as the orbit of a reference under the action of a specific group of transformations [51, 52]. Specifically, let $r : \mathcal{D} \to \mathbb{R}$ and $I : \mathcal{D} \to \mathbb{R}$ be the reference template and target shape, defined over a domain $\mathcal{D}$, respectively. Then the problem of shape matching can be thought of as finding the member of the group of transformations that takes the reference template closest to the target shape.

One such family of transformations is the group of diffeomorphisms, which is the set of functions $\phi : \mathcal{D} \to \mathcal{D}$ that are smooth and smoothly invertible. Under these conditions, an observed object $I(x)$ can be seen as a deformation of a reference object $r(x)$ with some error due to noise or other confounding factors ($I(x) = r(\phi(x)) + \epsilon$). This family has proven particularly useful in the field of computational anatomy, where this group of diffeomorphisms have been used to map regions of interest in organs to an atlas organ in order to develop biomarkers for disease [53, 54]. In the brain, for example, diffeomorphisms have been utilized to identify subcortical structures, such as the hippocampus, that are affected by neurological disorders like Alzheimer's Disease and the attention deficiency spectrum [55].

3.2.2 Continuous Metamorphosis Energy

Since the previous methods define an orbit of transformed shapes around a predefined template, if the target shape is not in the family of shapes spanned by trans-

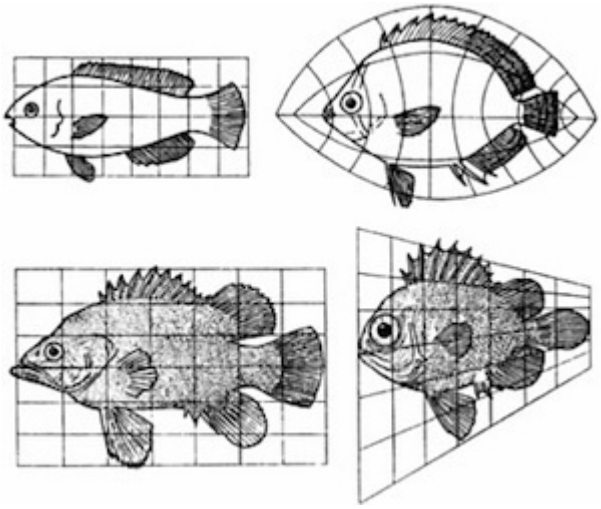

Figure 3.3: Typical example of homologous form from [50]

formations of the template, then these methods will fail to capture the morphology change accurately. In the diffeomorphism case, this is usually occurs when there is a subset of the range of the target signal that does not lie in the range of the template signal. Since we are only allowed to stretch and compress the domain, the resulting observed orbit will not capture the change in range of the target, and thus be attributed to noise. Given the evolving nature of ion channels and concentrations in cardiomyocytes from embryonic infancy to maturity along with inter-cellular variability, this is likely to happen when looking at the cardiac action potentials, and is indicative of electrophysiological phenomena that should not be discarded. How, then, do we incorporate changes in range into the deformable template framework?

The method of metamorphosis, as described in [56], extends deformable template theory, so that not only the target domain is deformed via a diffeomorphism, but also the target range is transformed as the reference evolves with the deformation.

As a consequence, instead of remaining in the orbit of one template, metamorphosis allows for a change of orbits in order to reach the target. In that sense, the Euclidean interpolation is one such metamorphosis transformation, where all of the evolution is done via modifying the template, and no deformation of the template is applied. Ideally, by finding the right combination of template transformation and diffeomorphic deformation, we can evolve an embryonic action potential into mature prototypes, in a way that preserves the shape of the intermediate action potentials.

More formally, let $f_0(\tau)$ and $f_1(\tau)$ be action potentials. As discussed in Chapter 2, assume f_0 and f_1 are periodic with cycle length T. The objective of metamorphosis is to find a family of functions $f(\tau, t)$ parameterized in $t \in [0, 1]$ that defines a path connecting f_0 and f_1, i.e., $f(\tau, 0) = f_0(\tau)$ and $f(\tau, 1) = f_1(\tau)$. Specifically, a metamorphosis defines a sequence of deformations $\phi(\tau, t)$ and templates $i(\tau, t)$ such that $f(\tau, t) = i(\phi^{-1}(\tau, t), t)$. A logical initial condition is that there is no deformation at the beginning of the interpolation, i.e., $\phi(\tau, 0) = id(\tau) = \tau$. This leads to the initial template being the starting action potential, i.e., $i(\tau, 0) = f_0(\tau)$. Note that when this template remains constant throughout the interpolation, i.e., $i(\tau, t) = i(\tau, 0) = f_0(\tau)$, then the metamorphosis is a diffeomorphic template family like those discussed previously.

How do we determine whether one metamorphosis between f_0 and f_1 is better than another metamorphosis? This requires defining an energy for a metamorphosis.

In metamorphosis, the energy is defined as follows:

$$E(\phi, i) := \int_0^1 \left(\left\| \frac{\partial \phi}{\partial t}(\tau, t) \right\|_\Phi^2 + \frac{1}{\sigma^2} \left\| \frac{\partial i}{\partial t}(\tau, t) \right\|_I^2 \right) dt, \tag{3.1}$$

where Φ and I are corresponding metrics for ϕ and i, respectively. In order to see how this is obtained, recall that metamorphosis defines a path between f_0 and f_1, so a logical choice for an energy is the arc length of the path. We calculate $\frac{\partial f}{\partial t}$ by following the chain rule:

$$\frac{\partial f}{\partial t}(\tau, t) = \frac{\partial i}{\partial \tau}(\phi^{-1}(\tau, t), t) \frac{\partial \phi^{-1}}{\partial t}(\tau, t) + \frac{\partial i}{\partial t}(\tau, t). \tag{3.2}$$

In order to handle $\frac{\partial \phi^{-1}}{\partial t}(\tau, t)$, we use the fact that $\phi(\phi^{-1}(\tau, t), t) = \tau$. Taking derivatives with respect to τ and t leads to the following pair of equations:

$$\frac{\partial \phi}{\partial \tau}(\phi^{-1}(\tau, t), t) \frac{\partial \phi^{-1}}{\partial t}(\tau, t) + \frac{\partial \phi}{\partial t}(\phi^{-1}(\tau, t), t) = 0, \tag{3.3}$$

$$\frac{\partial \phi}{\partial \tau}(\phi^{-1}(\tau, t), t) \frac{\partial \phi^{-1}}{\partial \tau}(\tau, t) = 1. \tag{3.4}$$

Combining the two equations leads to a conversion for $\frac{\partial \phi^{-1}}{\partial t}(\tau, t)$:

$$\frac{\partial \phi^{-1}}{\partial t}(\tau, t) = -\frac{\partial \phi}{\partial t}(\phi^{-1}(\tau, t), t) \frac{\partial \phi^{-1}}{\partial \tau}(\tau, t). \tag{3.5}$$

Substituting this into (3.2), we obtain:

$$\frac{\partial f}{\partial t}(\tau,t) = -\frac{\partial i}{\partial \tau}(\phi^{-1}(\tau,t),t)\frac{\partial \phi}{\partial t}(\phi^{-1}(\tau,t),t)\frac{\partial \phi^{-1}}{\partial \tau}(\tau,t) + \frac{\partial i}{\partial t}(\tau,t). \tag{3.6}$$

When $t = 0$, using the initial conditions described earlier, the above equation becomes:

$$\left.\frac{\partial f}{\partial t}(\tau,t)\right|_{t=0} = -\frac{\partial f}{\partial \tau}(\tau,0)\left.\frac{\partial \phi}{\partial t}(\tau,t)\right|_{t=0} + \left.\frac{\partial i}{\partial t}(\tau,t)\right|_{t=0}. \tag{3.7}$$

Since $\frac{\partial f}{\partial \tau}(\tau,0)$ is defined by the choice of template, (3.7) shows that we can evaluate the size of the infinitesimal change in the family f via infinitesimal changes in the deformation ϕ and template i. Furthermore, we can accumulate each infinitesimal step of the path in the same way by thinking about a corresponding metamorphosis starting at $f(\tau,\delta t)$ and finishing at $f(\tau,1)$. Thus, integrating over the entire path leads to (3.1).

Equation (3.1) requires us do define norms on both $\frac{\partial \phi}{\partial t}(\tau,t)$, and $\frac{\partial i}{\partial t}(\tau,t)$. Let $v(\tau,t) = \frac{\partial \phi}{\partial t}(\tau,t)$. This is typically called the flow field of the diffeomorphism. Given the periodicity in f, v is must also be periodic with period T. In order to ensure that the mapping ϕ is a diffeomorphism, this flow field needs to be smooth over τ. This is achieved by employing a norm that not only penalizes the magnitude of the function, but also the magnitude of the function's derivatives. For example, $\|v\|_V = \langle Lv, Lv\rangle$, where $Lv = v - \alpha\frac{\partial^2 v}{\partial \tau^2}$. Such norms are known as Sobolev norms. With respect to the norm on i, the Euclidean norm is an appropriate choice here, as we are looking at

infinitesimal changes in the template. Substituting these into (3.7) gives us:

$$E(v,f) = \int_0^1 \left(\|v(\cdot,t)\|_V^2 + \frac{1}{\sigma^2} \left\| \frac{\partial f}{\partial \tau}(\cdot,t) v(\cdot,t) + \frac{\partial f}{\partial t}(\cdot,t) \right\|_{L^2}^2 \right) dt. \tag{3.8}$$

Now that we have defined the energy for an individual metamorphosis, we can define the metamorphosis distance between two action potentials $f_0(\tau), f_1(\tau)$ as the smallest energy metamorphosis connecting the two signals:

$$d_{\mathcal{M}}^2(f_0,f_1) = \min_{\substack{v,f \\ f(\tau,0)=f_0(\tau) \\ f(\tau,1)=f_1(\tau)}} E(v,f)$$

$$d_{\mathcal{M}}^2(f_0,f_1) = \min_{\substack{v,f \\ f(\tau,0)=f_0(\tau) \\ f(\tau,1)=f_1(\tau)}} \int_0^1 \left(\|v(\tau,t)\|_V^2 + \frac{1}{\sigma^2} \left\| \frac{\partial f}{\partial \tau}(\tau,t) v(\tau,t) + \frac{\partial f}{\partial t}(\tau,t) \right\|_{L^2}^2 \right) dt. \tag{3.9}$$

3.2.3 Optimization of the Continuous Energy

How do we optimize this metamorphosis energy? Given the complicated nature of the energy, and the product of v and $\frac{\partial f}{\partial \tau}$ leading to some highly non-convex dynamics, our best approach is probably to alternate optimizing in one of the variables, while keeping the other fixed, until convergence [57]. The simplest way to do this is via alternating gradient descent. Fixing f and following the chain rule, variation of the energy with respect to $v(\tau,t)$ is the following:

$$\nabla_v E = \left(L^\dagger L + \frac{2}{\sigma^2} \left(\frac{\partial f}{\partial \tau}(\tau,t) \right)^2 \right) v(\tau,t) + \frac{2}{\sigma^2} \left(\frac{\partial f}{\partial \tau}(\tau,t) \frac{\partial f}{\partial t}(\tau,t) \right), \tag{3.10}$$

where L is a Sobolev norm operator ($Lv = v - \alpha\frac{\partial^2}{\partial\tau^2}v$, for example) and $L^\dagger$ is its adjoint. Obtaining the gradient for f is a bit more involved. Specifically, let us begin with the energy for a fixed $v(\tau, t)$:

$$E_v(f) = \int_0^1 \frac{1}{\sigma^2}\left\|\frac{\partial f}{\partial t}(\cdot,t) + \frac{\partial f}{\partial \tau}(\cdot,t)v(\cdot,t)\right\|_{L^2}^2 dt. \tag{3.11}$$

Computing the variation:

$$\begin{aligned}
(\nabla_f E_v; \tilde{f}) &= \lim_{\epsilon\to 0}\frac{1}{\epsilon\sigma^2}\left(\int_0^1\int_0^T\left((\frac{\partial f}{\partial\tau}v + \frac{\partial f}{\partial t}) + \epsilon(\frac{\partial\tilde{f}}{\partial\tau}v + \frac{\partial\tilde{f}}{\partial t})\right)^2 d\tau dt\right. \\
&\quad \left. - \int_0^1\int_0^T\left(\frac{\partial f}{\partial\tau}v + \frac{\partial f}{\partial t}\right)^2 d\tau dt\right) \\
&= \lim_{\epsilon\to 0}\frac{1}{\epsilon\sigma^2}\left(\int_0^1\int_0^T\left(2\epsilon(\frac{\partial f}{\partial\tau}v + \frac{\partial f}{\partial t})(\frac{\partial\tilde{f}}{\partial\tau}v + \frac{\partial\tilde{f}}{\partial t}) + \epsilon^2(\frac{\partial\tilde{f}}{\partial\tau}v + \frac{\partial\tilde{f}}{\partial t})^2\right)d\tau dt\right) \\
&= \frac{2}{\sigma^2}\int_0^1\int_0^T\left((\frac{\partial f}{\partial\tau}v + \frac{\partial f}{\partial t})\frac{\partial\tilde{f}}{\partial\tau}v + (\frac{\partial f}{\partial\tau}v + \frac{\partial f}{\partial t})\frac{\partial\tilde{f}}{\partial t}\right)d\tau dt \\
&= \frac{-2}{\sigma^2}\int_0^1\int_0^T\left(\frac{\partial}{\partial\tau}\left(v(\frac{\partial f}{\partial\tau}v + \frac{\partial f}{\partial t})\right) + \frac{\partial}{\partial t}\left(\frac{\partial f}{\partial\tau}v + \frac{\partial f}{\partial t}\right)\right)\tilde{f}\,d\tau dt.
\end{aligned} \tag{3.12}$$

Where the last equality is achieved via integration by parts, and the constraints that f, $\tilde{f}$, and v are periodic and that $\tilde{f}$ must keep the endpoints of the metamorphosis family the same ($\tilde{f}(\tau,0) = \tilde{f}(\tau,1) = 0$). Thus, the gradient is a sum of complicated partial derivatives:

$$\nabla_f E_v = -2(\frac{\partial}{\partial\tau}(v(\frac{\partial f}{\partial\tau}v + \frac{\partial f}{\partial t})) + \frac{\partial}{\partial t}(\frac{\partial f}{\partial\tau}v + \frac{\partial f}{\partial t})). \tag{3.13}$$

Given these two gradients, we can build an alternating descent algorithm to minimize the continuous energy (3.8).

3.2.4 Closed Form Updates for the Metamorphosis Family $f(\tau, t)$

Can we do any better than just obtaining the gradients? Ideally we would like a closed form solution for one of the variables when the other is fixed. Unfortunately, looking at the update for v first, solving this complex PDE for $v(\tau, t)$ is outside the scope of this thesis. Surprisingly, fixing $v(\tau, t)$ and solving for $f(\tau, t)$, however, can lead to something more tenable by making a change of variable [58]. Instead of looking at the optimization over the family $f(\tau, t)$, let us return to the sequence of templates $i(\tau, t)$. Let us first recall that $f(\tau, t) = i(\phi^{-1}(\tau, t), t)$ and $v = \frac{\partial \phi}{\partial t}$. We then have the relationship for the infinitesimal change in the template at any evolution time t:

$$\frac{\partial i}{\partial t}(\phi^{-1}(\tau, t), t) = \frac{\partial f}{\partial t}(\tau, t) + \frac{\partial f}{\partial \tau}(\tau, t) v(\tau, t) \tag{3.14}$$

This leads to an energy over i instead of f:

$$E_v(i) = \frac{1}{\sigma^2} \int_0^1 \int_0^T \left(\frac{\partial i}{\partial t}(\phi^{-1}(\tau, t), t) \right)^2 d\tau dt. \tag{3.15}$$

Note that $\phi^{-1}(\tau, t)$ is known because $v(\tau, t)$ is fixed, but let us decouple it from i by making the substitution $u = \phi^{-1}(\tau, t)$ (or $\phi(u, t) = \tau$). Then $d\tau = \frac{\partial \phi}{\partial u}(u, t)du$ and (3.15) becomes:

$$E_v(i) = \frac{1}{\sigma^2} \int_0^1 \int_0^T \left(\frac{\partial i}{\partial t}(u, t) \right)^2 \frac{\partial \phi}{\partial u}(u, t) du dt. \tag{3.16}$$

Now, we need to minimize this with respect to i. Taking the variation with respect to i:

$$\begin{aligned}
(\nabla_i E_v; \tilde{i}) &= \lim_{\epsilon \to 0} \frac{1}{\epsilon \sigma^2} \int_0^1 \int_0^T \left(\frac{\partial (i + \epsilon \tilde{i})}{\partial t}(u, t) \right)^2 \frac{\partial \phi}{\partial u}(u, t) du dt \\
&\quad - \frac{1}{\epsilon \sigma^2} \int_0^1 \int_0^T \left(\frac{\partial i}{\partial t}(u, t) \right)^2 \frac{\partial \phi}{\partial u}(u, t) du dt \\
&= \lim_{\epsilon \to 0} \frac{1}{\epsilon \sigma^2} \int_0^1 \int_0^T \left(2\epsilon \frac{\partial i}{\partial t} \frac{\partial \tilde{i}}{\partial t} + \epsilon^2 \left(\frac{\partial \tilde{i}}{\partial t} \right)^2 \right)(u, t) \frac{\partial \phi}{\partial u}(u, t) du dt \\
&= \frac{2}{\sigma^2} \int_0^1 \int_0^T \left(- \frac{\partial}{\partial t} \left(\frac{\partial i}{\partial t}(u, t) \frac{\partial \phi}{\partial u}(u, t) \right) \tilde{i}(u, t) \right) du dt.
\end{aligned} \tag{3.17}$$

Where the last equality is achieved by integration by parts, plus the fact that with a fixed $v(\tau, t)$, $i(u, 0)$ and $i(u, 1)$ are fixed in order to meet the constraint on $f(\tau, 0)$ and $f(\tau, 1)$. Any variation $\tilde{i}$ of i still needs to respect the constraint on f, which forces $\tilde{i}(u, 0) = \tilde{i}(u, 1) = 0$. So this means that $\nabla_i E = \frac{2}{\sigma^2} \big(- \frac{\partial}{\partial t} \big(\frac{\partial i}{\partial t}(u, t) \frac{\partial \phi}{\partial u}(u, t) \big) \big)$. In order

to find the optimum, we set this gradient to 0:

$$\begin{aligned} 0 &= \frac{\partial}{\partial t}\left(\frac{\partial i^*}{\partial t}(u,t)\frac{\partial \phi}{\partial u}(u,t)\right) \\ \implies c(u) &= \left(\frac{\partial i^*}{\partial t}(u,t)\right)\frac{\partial \phi}{\partial u}(u,t) \\ \implies c(u)\int_0^{t'} \frac{1}{\frac{\partial \phi}{\partial u}(u,t)}dt &= \int_0^{t'}\left(\frac{\partial i^*}{\partial t}(u,t)\right)dt = i^*(u,t') - i^*(u,0). \end{aligned} \tag{3.18}$$

Setting $t' = 1$ allows us to solve for $c(u)$:

$$c(u) = \frac{i^*(u,1) - i^*(u,0)}{\int_0^1 \frac{1}{\frac{\partial \phi}{\partial u}(u,t)}dt}. \tag{3.19}$$

Combining the last two equations:

$$i^*(u,t') = i(u,0)\frac{\int_{t'}^1 \frac{1}{\frac{\partial \phi}{\partial u}(u,t)}dt}{\int_0^1 \frac{1}{\frac{\partial \phi}{\partial u}(u,t)}dt} + i(u,1)\frac{\int_0^{t'} \frac{1}{\frac{\partial \phi}{\partial u}(u,t)}dt}{\int_0^1 \frac{1}{\frac{\partial \phi}{\partial u}(u,t)}dt}. \tag{3.20}$$

Therefore, the optimal family of templates, $i^*(u,t)$, when the diffeomorphisms $\phi(u,t)$ are fixed, is a weighted sum between the templates at the beginning and the end of the metamorphosis, and can be computed in closed form. Notice that when the there is no deformation of the template, i.e., $\phi(u,t) = id(u) = u \forall t$, (3.20) simplifies to linear interpolation. For more complex v, we obtain $\phi(u,t)$ by integrating along t (Recall $v(\tau,t) = \frac{\partial \phi}{\partial t}(\tau,t)$). More simply, $\phi(u,t)$ is determined by putting a particle at time point u at evolution time 0, and observing its location at evolution time t, given the flow fields v. Doing this for all time points u allows us to calculate $\frac{\partial \phi}{\partial u}(u,t)$. Once

the family of templates is known, combining them with the deformations allow us to determine the metamorphosis family (Recall $f(\tau, t) = i(\phi^{-1}(\tau, t), t)$).

3.3 Discrete Metamorphosis

In the previous section, we defined metamorphosis in the continuous domain and calculated the gradients in order to determine a scheme to minimize the metamorphosis energy. Further, we showed how to solve for metamorphosis the family $f(\tau, t)$ in closed form given a fixed deformation velocity field $v(\tau, t)$. But how do we go about transferring that into an algorithm implementable by a computer. The obvious approach is to approximate the differential equations in (3.10) and (3.20) with discrete difference schemes in order to implement the gradients calculated previously. The discretized updates, however, are not guaranteed to minimize the continuous energy. Therefore, this thesis will follow the alternative approach of finding an appropriate discrete approximation to the energy in (3.8), then determine how to minimize the discretized energy. This ensures that we develop methods to minimize the given objective and lead to minimal geodesics.

3.3.1 Discrete Metamorphosis Energy

As previously discussed, the distance between two action potentials comes down to solving the following optimization problem:

$$d^2_{\mathcal{M}}(f_0, f_1) = \min_{\substack{v,f \\ f(\tau,0)=f_0(\tau) \\ f(\tau,1)=f_1(\tau)}} \int_0^1 \left(\|v(\cdot,t)\|^2_V + \frac{1}{\sigma^2} \left\| \frac{\partial f}{\partial \tau}(\cdot,t)v(\cdot,t) + \frac{\partial f}{\partial t}(\cdot,t) \right\|^2_{L^2} \right) dt. \quad (3.21)$$

Let us first discretize the interpolation parameter t at evenly spaced intervals, namely $t_s = \frac{s}{S}, s = 0, \ldots, S$. First, we need to identify how to discretize the right summand of the integral in equation 3.9. Note the following Taylor Series expansion:

$$\begin{aligned} f(\tau + \Delta t v(\tau,t), t + \Delta t) &= f(\tau,t) + \frac{\partial f}{\partial \tau}(\tau,t)(\Delta t v(\tau,t)) + \frac{\partial f}{\partial t}(\tau,t)(\Delta t) + \mathcal{O}(\Delta t^2) \\ &= f(\tau,t) + \Delta t \left(\frac{\partial f}{\partial \tau}(\tau,t)v(\tau,t) + \frac{\partial f}{\partial t}(\tau,t) \right) + \mathcal{O}(\Delta t^2). \end{aligned}$$

Thus as $\Delta t \to 0$,

$$\frac{f(\tau + \Delta t v(\tau,t), t + \Delta t) - f(\tau,t)}{\Delta t} \to \left(\frac{\partial f}{\partial \tau}(\tau,t)v(\tau,t) + \frac{\partial f}{\partial t}(\tau,t) \right). \quad (3.22)$$

Letting $\Delta t = \frac{1}{S}$, we can approximate the energy in equation 3.9 by the following sum:

$$d^2_{\mathcal{M}}(f_0, f_1) \approx \min_{\substack{v,f \\ f(\tau,0)=f_0(\tau) \\ f(\tau,1)=f_1(\tau)}} \sum_{s=0}^{S-1} \left(\|v(\tau,t_s)\|^2_V + \frac{1}{\sigma^2} \left\| \frac{f(\tau + \frac{1}{S}v(\tau,t_s), t_{s+1}) - f(\tau,t_s)}{\frac{1}{S}} \right\|^2_{L^2} \right) \quad (3.23)$$

Instead of carrying this $\frac{1}{S}$ term everywhere in the calculation, note that we can multiply the entire objective by $\frac{\frac{1}{S^2}}{\frac{1}{S^2}}$, take the $\frac{1}{S}^2$ term in the numerator inside the norm of v, and bring $\frac{1}{S}$ into v by substituting $\frac{1}{S}v(\tau, t_s)$ by $v(\tau, t_s)$. As a result,

$$d^2_{\mathcal{M}}(f_0, f_1) \approx \min_{\substack{v,f \\ f(\tau,0)=f_0(\tau) \\ f(\tau,1)=f_1(\tau)}} S^2 \sum_{s=0}^{S-1} (\|v(\tau, t_s)\|^2_V + \frac{1}{\sigma^2}\|f(\tau + v(\tau, t_s), t_{s+1}) - f(\tau, t_s)\|^2_{L^2}) \tag{3.24}$$

Now we can focus on τ. Discretization in τ is done by sampling the signal along an evenly spaced fixed grid, i.e., $\tau_j = \frac{j}{J}T, j = 0, \ldots, J-1$. Including the discretization in τ leads to the final discretized energy:

$$d^2_{\mathcal{M}}(f_0, f_1) \approx \min_{\substack{v,f \\ f(\tau,0)=f_0(\tau) \\ f(\tau,1)=f_1(\tau)}} S^2 \sum_{s=0}^{S-1} (\|v(\tau_j, t_s)\|^2_{V_d} + \frac{1}{\sigma^2}\|f(\tau_j + v(\tau_j, t_s), t_{s+1}) - f(\tau_j, t_s)\|^2_{l^2}), \tag{3.25}$$

where V_d defines a discretized Sobolev norm, i.e. $\|v\|^2_{V_d} = \langle L_d v, L_d v\rangle$, with L_d a discretized approximation of the Sobolev differential operator L. It should be noted that this optimization problem is nonconvex due to the composition of f and v in the second term. Thus it is possible to get caught in a local minimum depending on the initialization and optimization procedure. We make the choice of initializing the metamorphosis with the Euclidean interpolation ($f(\tau_j, t) = (1-t)f_0(\tau_j) + tf_1(\tau_j)$) with no flow ($v(\tau_j, t_s) = 0, j = 0, \ldots, J-1, s = 0, \ldots, S-1$), as it is a special case of metamorphosis. We enforce the constraint on $f(\tau, 0)$ and $f(\tau, t_S)$ by fixing them to be equal to $f_0(\tau)$ and $f_1(\tau)$, respectively. This leaves the set of flow fields

$\{v(\tau_j, t_s), j = 0, \ldots, J-1, s = 0, \ldots, S-1\}$ and the internal interpolants of the metamorphosis family $\{f(\tau_j, t_s), j = 0, \ldots, J-1, s = 1, \ldots, S-1\}$ as the independent variables to be optimized over.

3.3.2 Optimization of the Discrete Energy

Let $E_d = S^2 \sum_{s=0}^{S-1} (\|v(\tau_j, t_s)\|_{V_d}^2 + \frac{1}{\sigma^2} \|f(\tau_j + v(\tau_j, t_s), t_{s+1}) - f(\tau_j, t_s)\|_{l^2}^2)$. The simplest approach to obtaining a local minimum of (3.25) is via gradient descent. Given the complex interactions of v and f in the right summand of the integral, it is not obvious how to descend in a joint v and f direction. A more tenable approach is to fix one set of variables, optimize over the other set, and then switch roles and repeat this process until convergence. This general process is described in Algorithm 2. In order to do this, we need to be able to calculate the gradients with respect to $v(\tau_j, t_s)$ with $f(\tau_j, t_s)$ fixed and vice versa.

Calculating the gradient with respect to $f(\tau_j, t_s)$. When $v(\tau_j, t_s)$ is fixed the only term in 3.25 that needs to be optimized over $f(\tau_j, t_s)$ is the second term: Let $(E_d)_v(f) = \sum_{s=0}^{S-1} \frac{1}{\sigma^2} \|f(\tau_j + v(\tau_j, t_s), t_{s+1}) - f(\tau_j, t_s)\|_{l^2}^2 = \sum_j^J (f(\tau_j + v(\tau_j, t_s), t_{s+1}) - f(\tau_j, t_s))^2$. The concern then is that position of the sample $\tau_j + v(\tau_j, t_s)$ may not lie on the fixed grid, and thus has to be approximated by the samples on the fixed grid. We do this via linear interpolation: Let $\mathbf{v}_s$ be the vector containing the samples of $v(\cdot, t_s)$, i.e. $(\mathbf{v}_s)_j = v(\tau_j, t_s)$, and let $N_{\mathbf{v}_s}$ be the matrix the approximates the values of $f(\tau_j + v(\tau_j, t_s), t_{s+1})$ via the samples $f(\tau_j, t_{s+1})$: $(N_{\mathbf{v}_s}) f(\tau_j, t_{s+1}) \approx f(\tau_j +$

Algorithm 2 Discrete Metamorphosis Optimization

Input: Template Signal $f_0(\tau)$, Target Signal $f_1(\tau)$, Balance Parameter σ, Number of Evolution Intermediates S, Sobolev Operator L_d, Update frequency n, Convergence parameter δ, and Maximum number of Iterations M.

Initialization
$m = 0$, $d_{-1} = \infty$, $K = L_d^{-1}$.
for $s = 0, \ldots, S$ **do**
    $w(\tau_i, t_s) \equiv 0$, $v(\tau_i, t_s) = Kw(\tau_i, t_s) \equiv 0$
    $f(\tau_i, t_s) = \frac{S-s}{S} f_0(\tau_i) + \frac{s}{S} f_1(\tau_i)$. $\mathbf{f}_s^T = [f(\cdot, t_s)]$
end for
$d_0^2 \leftarrow \sum_{s=0}^{S-1} \|w(\tau_i, t_s)\|_{l_2}^2 + \frac{1}{\sigma^2} \|N_{v_s} \mathbf{f}_{s+1} - \mathbf{f}_s)\|_{l_2}^2$
while $|d_{m-1} - d_m| > \delta$ **AND** $j < M$ **do**
    $m \leftarrow m + 1$
    for $s = 0, \ldots, S-1$ **do**
        Perform gradient descent to update $w(\tau_i, t_s)$ using (3.31).
        $v(\tau_i, t_s) \leftarrow \mathbf{R}(Kw(\tau_i, t_s))$,
        $N_{v_s} \leftarrow$ (3.26).
    end for
    for $s = 1, \ldots, S-1$ **do**
        Perform gradient descent to update $f(\tau_i, t_s)$ using (3.28).
    end for
    $d_m^2 \leftarrow \sum_{s=0}^{S-1} \|w(\tau_i, t_s)\|_{l_2}^2 + \frac{1}{\sigma^2} \|N_{v_s} \mathbf{f}_{s+1} - \mathbf{f}_s)\|_{l_2}^2$
end while

Output: Metamorphosis family $\mathbf{f}_s$, Metamorphosis velocity $v(\tau_i, t_s)$.

$v(\tau_J, t_s), t_{s+1})$. Specifically:

$$(N_{\mathbf{v}_s})_{mn} = \begin{cases} \frac{\tau_m + v(\tau_m, t_s)}{\Delta\tau} - (n-1) & n-1 <= \frac{\tau_m + v(\tau_m, t_s)}{\Delta\tau} < n \\ n + 1 - \frac{\tau_i + v}{\Delta\tau} & n <= \frac{\tau_m + v(\tau_m, t_s)}{\Delta\tau} < n+1 \\ 0 & \text{otherwise} \end{cases} . \tag{3.26}$$

This allows us to write the second term in vector notation as:

$$(E_d)_v(f) = \sum_{s=0}^{S-1} \frac{1}{\sigma^2} \| N_{\mathbf{v}_s} \mathbf{f}_{s+1} - \mathbf{f}_s \|_{l^2}^2, \tag{3.27}$$

where $\mathbf{f}_s$ is the column vector that contains the samples of $f(\cdot, t_s)$. It is easy to calculate from here the gradient with respect to $\mathbf{f}_s$ as:

$$\frac{\partial (E_d)_v}{\partial \mathbf{f}_s} = \frac{2}{\sigma^2} N_{\mathbf{v}_{s-1}}^T (N_{\mathbf{v}_{s-1}} \mathbf{f}_s - \mathbf{f}_{s-1}) - \frac{2}{\sigma^2} (N_{\mathbf{v}_s} \mathbf{f}_{s+1} - \mathbf{f}_s). \tag{3.28}$$

Calculating the gradient with respect to $v(\tau_j, t_s)$. Let L_d be the discretized version of the Sobolev norm operator L, and recall that $\|v(\cdot, t_s)\|_{V_d}^2 = \langle L_d v(\cdot, t_s), L_d v(\cdot, t_s) \rangle$. L_d can be generated via finite differences, but since we are assuming the action potential is periodic, we are going to utilize Fourier theory to generate the discretized differential operator. With the following choice of operator $L_d v = v - \alpha \frac{\partial^2 v}{\partial \tau^2}$, we obtain:

$$\begin{aligned} u = L_d v &= v - \alpha \frac{\partial^2}{\partial \tau^2} v \\ \implies \mathcal{F}(u) &= \mathcal{F}(v) - \alpha (j\omega)^2 \mathcal{F}(v) \\ \implies L_d v &= \mathcal{F}^{-1}((1 + \alpha \omega^2) \mathcal{F}(v)). \end{aligned} \tag{3.29}$$

Thus L_d can be computed by taking the Fourier Transform, scaling it by the frequency, and taking the inverse transform. Instead of operating on v, however, and carrying this operator throughout the calculations of the gradient, we act on the dummy

variable $w(\tau,t) = L_d v(\tau,t)$ and let $K = L_d^{-1}$. It follows from the previous calculation that for the choice of operator described earlier, we have $K = \mathcal{F}^{-1}\frac{1}{1+\alpha\omega^2}\mathcal{F}$ and is still easily computable. Making this substitution into (3.25) leads to the following for a fixed metamorphosis f:

$$(E_d)_f(w) = \sum_{s=0}^{S-1} \|w(\tau,t)\|_{l^2}^2 + \frac{1}{\sigma^2}\|f(\tau + Kw(\tau,t_s),t_{s+1}) - f(\tau,t_s)\|_{l^2}^2. \tag{3.30}$$

It is important to note that with the metamorphosis family f fixed, the updates to the flow fields w are decoupled from each other. More specifically, following the chain rule:

$$\frac{\partial E_f}{\partial w(\tau,t_s)} = 2w(\tau,t_s) + \frac{2}{\sigma^2}K(f(\tau+Kw(\tau,t_s),t_{s+1}) - f(\tau,t_s))\frac{\partial f}{\partial \tau}(f(\tau+Kw(\tau,t_s),t_{s+1})) \tag{3.31}$$

3.3.3 Closed Form Updates for the Discrete Metamorphosis Family $f(\tau_j, t_s)$

We have shown how to calculate the gradient of the discrete energy with respect to both $v(\tau_j,t_s), j = 0,\ldots,J-1, s = 0,\ldots,S-1$ and $f(\tau_j,t_s), j = 0,\ldots,J-1, s = 1,\ldots,S-1$. Just as we did in the case of the continuous energy, we would like to use these gradients to try to solve for the optimum of some of the variables while others are fixed. Our hope is that by finding a closed form solution for the optimum, we can

speed up the metamorphosis calculation, allowing us to work with larger populations of embryonic cardiomyocytes.

The derivation in section 3.2.4 inspires us to seek a potential discrete analog to the closed form solution for $f(\tau, t)$. More specifically, we seek to find an appropriate closed form update for $\mathbf{f}_s = f(\cdot, t_s), s = 1, \ldots, S-1$ for a fixed $\mathbf{v}_s$ that minimizes the discrete energy defined in 3.25. We begin by considering the gradient of the discrete energy with respect to $f(\tau_j, t_s)$ (Equation 3.28). Let $N_s = N_{\mathbf{v}_s}$ and $\mathbf{z}_s = N_s \mathbf{f}_{s+1} - \mathbf{f}_s$. Then at the optimum, (3.28) becomes:

$$\mathbf{z}_s = N_{s-1}^T \mathbf{z}_{s-1}, s = 1, \ldots, S-1. \tag{3.32}$$

Then, for any s between 1 and $S-1$:

$$\mathbf{z}_s = N_{s-1}^T N_{s-2}^T \ldots N_0^T \mathbf{z}_0 = R_{0,s}^T \mathbf{z}_0, \tag{3.33}$$

where $R_{0,s} = N_0 \ldots N_{s-1}$, or more generally, $R_{a,b} = N_a N_{a+1} \ldots N_{b-1}$, for $a < b$. Note that $R_{a,b} R_{b,c} = R_{a,c}$ for $a < b < c$. It follows from the definition of $\mathbf{z}_s$ that

$$\mathbf{f}_s = N_s \mathbf{f}_{s+1} - \mathbf{z}_s = N_s \mathbf{f}_{s+1} - R_{0,s}^T \mathbf{z}_0, \tag{3.34}$$

Let us look at what this equation means, starting from $s = S - 1$:

$$\begin{aligned}
\mathbf{f}_{S-1} &= N_{S-1}\mathbf{f}_S - R^T_{0,S-1}\mathbf{z}_0 \\
\mathbf{f}_{S-2} &= N_{S-2}\mathbf{f}_{S-1} - R^T_{0,S-2}\mathbf{z}_0 \\
&= N_{S-2}N_{S-1}\mathbf{f}_S - (N_{S-2}R^T_{0,S-1} + R^T_{0,S-2})\mathbf{z}_0 \\
\mathbf{f}_{S-3} &= N_{S-3}\mathbf{f}_{S-2} - R^T_{0,S-3}\mathbf{z}_0 \\
&= N_{S-3}N_{S-2}N_{S-1}\mathbf{f}_S - (N_{S-3}N_{S-2}R^T_{0,S-1} + N_{S-3}R^T_{0,S-2} + R^T_{0,S-3})\mathbf{z}_0 \\
&= R_{S-3,S}\mathbf{f}_S - \big(\sum_{i=S-3}^{S-1} R_{S-3,i}R^T_{0,i}\big)\mathbf{z}_0
\end{aligned}$$

Following this logic, we have:

$$\mathbf{f}_s = R_{s,S}\mathbf{f}_S - \big(\sum_{i=s}^{S-1} R_{s,i}R^T_{0,i}\big)\mathbf{z}_0 = R_{s,S}\mathbf{f}_S - A_{s,S}\mathbf{z}_0, \tag{3.35}$$

where $A_{l,m} = \sum_{i=l}^{m-1} R_{l,i}R^T_{0,i}$. Every $\mathbf{f}_s$ is a function of $\mathbf{f}_S$ and $\mathbf{z}_0$. We can replace this dependence on $\mathbf{z}_0$ by looking at this equation when $s = 0$:

$$\mathbf{f}_0 = R_{0,S}\mathbf{f}_S - A_{0,S}\mathbf{z}_0 \implies \mathbf{z}_0 = A^{-1}_{0,S}(R_{0,S}\mathbf{f}_S - \mathbf{f}_0). \tag{3.36}$$

Substituting this back into equation (3.35):

$$\mathbf{f}_s = A_{s,S}A^{-1}_{0,S}\mathbf{f}_0 + (R_{s,S} - A_{s,S}A^{-1}_{0,S}R_{0,S})\mathbf{f}_S. \tag{3.37}$$

Thus the optimal metamorphosis interpolants $(\mathbf{f}_s, s = 1, \ldots, S-1)$ for the discretized energy (3.24) is again a weighted combination of the initial signal $\mathbf{f}_0$, and the target $\mathbf{f}_S$. Further, based on the following theorem, this update leads to the same interpolants as in (3.20) in the continuum limit:

Theorem 1. *Let $\tau_j = \frac{j}{J}T$ for $j = 0, \ldots, J$ and cycle length T, and $t_s = \frac{s}{S}$ for $s = 0, \ldots, S$. Let $\{f(\tau_j, t_s), j = 0, \ldots, J\}_{s=0}^{S} = \{\mathbf{f}_s\}_{s=0}^{S}$ be set the metamorphosis interpolants derived from the discrete interpolant update* (3.37)*:*

$$\mathbf{f}_s = A_{s,S}A_{0,S}^{-1}\mathbf{f}_0 + (R_{s,S} - A_{s,S}A_{0,S}^{-1}R_{0,S})\mathbf{f}_S$$

and let $f_c(\tau, t)$ be the continuous metamorphosis family of interpolants derived from $i(u, t)$ given by the (3.20)*:*

$$i(u,t) = i(u,0)\frac{\int_t^1 \frac{1}{\frac{\partial\phi(u,t')}{\partial u}}dt'}{\int_0^1 \frac{1}{\frac{\partial\phi(u,t')}{\partial u}}dt'} + i(u,1)\frac{\int_0^t \frac{1}{\frac{\partial\phi(u,t')}{\partial u}}dt'}{\int_0^1 \frac{1}{\frac{\partial\phi(u,t')}{\partial u}}dt'}$$

with $i(\tau, t) = f_c(\phi(\tau, t), t)$ and $u = \phi^{-1}(\tau, t)$. Then, as $J \to \infty$, $\tau_j \to \tau$ and $\mathbf{f}_s \to f(\tau, \frac{s}{S})$. That is,

$$\lim_{J\to\infty} \mathbf{f}_s = f_c(\tau, \frac{s}{S}). \tag{3.38}$$

Proof. First, we show that the discrete version of the deforming templates $\mathbf{i}_s$, given $\mathbf{f}_s$ is $\mathbf{i}_s = R_{0,s}\mathbf{f}_s$. By definition, $\mathbf{i}_0 = \mathbf{f}_0$. Notice that: $\mathbf{i}_1 = \mathbf{i}_0 + \mathbf{z}_0 = \mathbf{f}_0 + N_{v_0}\mathbf{f}_1 - \mathbf{f}_0 = N_{v_0}\mathbf{f}_1$. In order to calculate $\mathbf{i}_2$, we need resample $\mathbf{z}_1$ at sample points $\tau_j + v(\tau_j, t_0)$ in order

to correspond to the samples τ_i at i_1. This can be done by multiplying $\mathbf{z}_1$ on the left by N_{v_0}. Then:

$$\begin{aligned}\mathbf{i}_2 &= \mathbf{i}_1 + N_{v_0}\mathbf{z}_1 \\ &= N_{v_0}\mathbf{f}_1 + N_{v_0}N_{v_1}\mathbf{f}_2 - N_{v_0}\mathbf{f}_1 \\ &= R_{0,2}\mathbf{f}_2. \end{aligned} \tag{3.39}$$

Following this logic for subsequent s gives us $\mathbf{i}_s = R_{0,s}\mathbf{f}_s$. Given this, the discrete update can be re-written as follows:

$$\begin{aligned} R_{0,s}^{-1}\mathbf{i}_s &= A_{s,S}A_{0,S}^{-1}\mathbf{i}_0 + (R_{s,S} - A_{s,S}A_{0,S}^{-1}R_{0,S})R_{0,S}^{-1}\mathbf{i}_S \\ \implies \mathbf{i}_s &= R_{0,s}A_{s,S}A_{0,S}^{-1}\mathbf{i}_0 + (Id - R_{0,s}A_{s,S}A_{0,S}^{-1})\mathbf{i}_S, \end{aligned} \tag{3.40}$$

where the last equality is achieved as $R_{0,s}R_{s,S}R_{0,S}^{-1} = R_{0,S}R_{0,S}^{-1} = Id$. Before we continue, let us look at $R_{0,s}A_{s,S}$:

$$R_{0,s}A_{s,S} = R_{0,s}\sum_{k=s}^{S-1} R_{s,k}R_{0,k}^T = \sum_{k=s}^{S-1} R_{0,k}R_{0,k}^T = \sum_{k=0}^{S-1} R_{0,k}R_{0,k}^T - \sum_{k=0}^{s-1} R_{0,k}R_{0,k}^T = A_{0,S} - A_{0,k}. \tag{3.41}$$

Thus, putting this back in the previous equation, we get:

$$\mathbf{i}_s = (Id - A_{0,s}A_{0,S}^{-1})\mathbf{i}_0 + A_{0,s}A_{0,S}^{-1}\mathbf{i}_S. \tag{3.42}$$

We have now re-written the discrete update equation (3.37) in terms of i_s. We now allow $J \to \infty$. Naturally, $\mathbf{i}_0 \to i(\cdot, 0)$, $\mathbf{i}_S \to i(\cdot, 1)$, and $\mathbf{i}_s \to i(\cdot, \frac{s}{S})$. To see what happens to the operator $A_{0,s}$, we start with $R_{0,s}$. Since N_{v_s} relates samples at time t_{s+1} to samples at time t_s, as $J \to \infty$, N_{v_s} represents the deformation map between $[t_s, t_{s+1}]$. And since $R_{0,s} = N_{v_0} \dots N_{v_s - 1}$:

$$\lim_{J \to \infty} R_{0,s}\mathbf{i} = i(\phi(\cdot, \frac{s}{S})) = i(\phi_{\frac{s}{S}}(\cdot)). \tag{3.43}$$

For ease of notation, from now on we will drop the notation $\frac{s}{S}$ from ϕ. Now, finding the limit of the adjoint of $R_{0,s}$ as $J \to \infty$ will equal the adjoint of the operator on the right. The adjoint of $i(\phi(\cdot))$ can be calculated via a dummy integration:

$$\int_0^T i(\phi(\tau))\hat{i}(\tau)d\tau = \int_0^T i(u)\hat{i}(\phi^{-1}(u))\Big|\frac{\partial \phi^{-1}}{\partial u}(u)\Big|du. \tag{3.44}$$

Thus $\lim_{J \to \infty} R_{0,s}^T\mathbf{i} = i(\phi^{-1}(\cdot))\Big|\frac{\partial \phi^{-1}}{\partial(\cdot)}(\cdot)\Big|$. Combining the two operators together:

$$\begin{aligned} \lim_{J \to \infty} R_{0,s}R_{0,s}^T\mathbf{i} &= i(\phi^{-1}(\phi(\cdot)))|\frac{\partial \phi^{-1}}{\partial(\cdot)}(\phi(\cdot))| \\ &= i(\cdot)\left|\frac{1}{\frac{\partial \phi}{\partial(\cdot)}(\cdot)}\right| \end{aligned} \tag{3.45}$$

The final equality can be seen by taking looking at the derivative of $u = \phi^{-1}(\phi(u))$.

Since $A_{0,s} = \sum_{i=0}^{s-1} R_{0,s} R_{0,s}^T$ it follows that:

$$\lim_{J\to\infty} A_{0,s}\mathbf{i} = i(\cdot) \int_0^{\frac{s}{S}} |\frac{1}{\frac{\partial\phi}{\partial(\cdot)}(\cdot)}| dt. \tag{3.46}$$

Replacing $A_{0,s}$ and $A_{0,S}$ with this limit and doing some basic math leads to (3.20). □

As nice as all of this linear algebra and variational calculus is, there is little benefit if it does not lead to something that is easy and efficient to implement on a computer. The next question to tackle is how to implement the closed form update in equation (3.37). Looking at (3.37), the quantities of particular interest to efficiently calculate are $A_{s,S}$. Attempting to solve it by its current definition ($A_{l,m} = \sum_{i=l}^{m-1} R_{l,i} R_{0,i}^T$) would require us to store all of linear operations $R_{i,j}$, which is a large computational burden. So let us first aim to simplify this calculation:

$$\begin{aligned} A_{s,S} &= \sum_{i=s}^{S-1} R_{s,i} R_{0,i}^T \\ &= R_{s,s} R_{0,s}^T + \sum_{i=s+1}^{S-1} N_s R_{s+1,i} R_{0,i}^T \\ &= R_{0,s}^T + N_k A_{s+1,S}, \end{aligned} \tag{3.47}$$

This final equation is much better, because we can recursively generate $A_{s,S}$, given $A_{s+1,S}$, and $R_{0,s}$, which is a much smaller set than $R_{i,j}$. Moreover, by letting $B_{s,S} = A_{s,S} A_{0,S}^{-1}$ and $C_{0,s}^T = R_{0,s}^T A_{0,S}^{-1}$, we obtain the recursion $B_{s,S} = C_{0,s}^T + N_s B_{s+1,S}$. This

leads to the following system up backwards and forwards updates:

$$\mathbf{f}_s = B_{s,S}\mathbf{f}_0 + (R_{s,S} - B_{s,S}R_{0,S})\mathbf{f}_S, \tag{3.48}$$

$$R_{s,S} = N_s R_{s+1,S}, R_{S,S} = Id, \tag{3.49}$$

$$C_{0,s}^T = N_s^T C_{0,s-1}^T, C_{0,0} = A_{0,S}^{-1}, \tag{3.50}$$

$$B_{s,S} = C_{0,s}^T + N_s B_{s+1,S}, B_{S,S} = 0, \tag{3.51}$$

with Id being the identity matrix. Using these updates allows us to avoid multiple calculations of $A_{0,S}^{-1}$. Specifically, $A_{0,S}^{-1}$ can be calculated by the definition of A and the aggregation of $R_{0,s}$, which can be kept in order to calculate $C_{0,s}$. As the optimization gets closer to the optimum, and $v(\tau, t)$ changes by smaller amounts, $A_{0,S}^{-1}$ can be updated less often to speed up the optimization.

The updated process is given in Algorithm 3. The update of the velocity remains the same (3.31), and we again initialize with the Euclidean interpolation. There are two potential avenues for improving this algorithm. The first, which we have discussed previously, is to utilize the non-homogeneous differential equation in (3.10) to determine a closed form update for the velocity $v(\tau, t)$ for fixed metamorphosis interpolants $f(\tau, t)$. The other is to find a first order approximation to update the matrices $R_{i,j}$ and $A_{l,m}$ for a particular change in the velocity fields $\Delta v(\tau, t)$. As we get closer to the optimum, $v(\tau, t)$ changes very little, and thus $N_k, R_{i,j}$, and $A_{l,m}$ should also be changing by a small amount. Finding an appropriate $\Delta N_k, \Delta R_{i,j}$, and $\Delta A_{l,m}$

for a given $\Delta v(\tau, t)$ would alleviate a large number of recalculations during every update of f and would drastically improve the performance of the metamorphosis algorithm. This sensitivity analysis is part of future research directions.

3.4 Experiments and Results

3.4.1 Single Cell Recording Dataset

In order to evaluate the potential of the metamorphosis distance for action potential classification, we compare Algorithm 2 and Algorithm 3 on the dataset generated by [16]. The dataset of [16] consists of 16 atrial and 36 ventricular cells, which were manually labeled using biological characteristics of the APs, as described by the authors. Since the embryonic signals were spontaneously paced, we used the preprocessing steps described in [59] to adjust their cycle length to 1 second. We then generated 10 mature atrial and 10 mature ventricular protoype action potentials using the atrial model of [41] and the ventricular model of [45], respectively, by modulating the conductances in the ion channels between 80% and 120% of the value of the base model. These prototypes were paced to have a cycle length of 1 second to match the embryonic data. All signals were then normalized so that the resting membrane potential has voltage 0, and the amplitude has voltage 1.

We computed the metamorphosis distance from each of the embryonic cardiomyocytes to each one of the mature prototypes. We selected a balance parameter $\sigma = 0.3$,

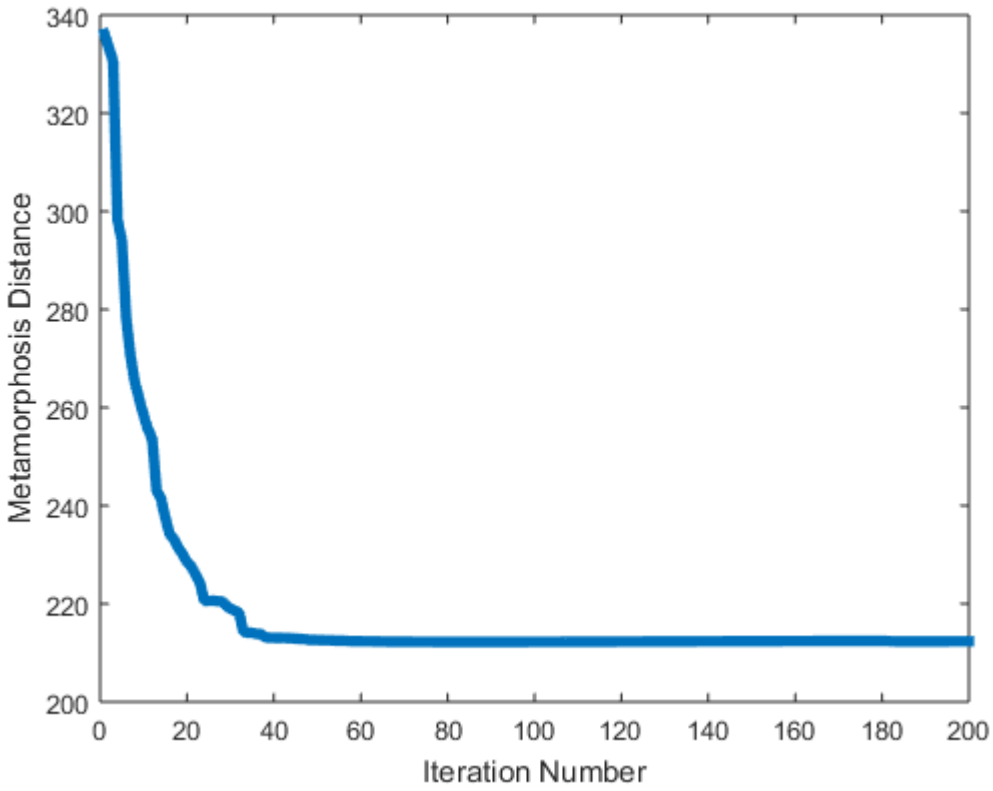

Figure 3.4: Metamorphosis Distance as a function of iterations of Algorithm 3

number of interpolants between template and target as 3 ($S = 4$), and the linear operator $L_d(\cdot) = id(\cdot) - \alpha\Delta(\cdot)$, with $\alpha = 8$. We iterated both Algorithm 2 and Algorithm 3 until they reached convergence or 300 iterations. Generally, all 300 iterations are unnecessary, as the distance tends to converge much sooner, as shown in Figure 3.4.

Figure 3.5 compares the method of Algorithm 2 and two variants of Algorithm 3 (with $A_{0,S}^{-1}$ updated each iteration or every 10 iterations) in terms of the final interpolations and the distances they produce. We see that there is very little difference between the three interpolations, and that the three distances are approximately

equal. The main difference is that our method with $A_{0,S}^{-1}$ updated every 10 iterations requires about half the number of iterations than the other methods. Notice that the metamorphosis method acts to preserve the shape of action potential throughout the entire hypothesized interpolation.

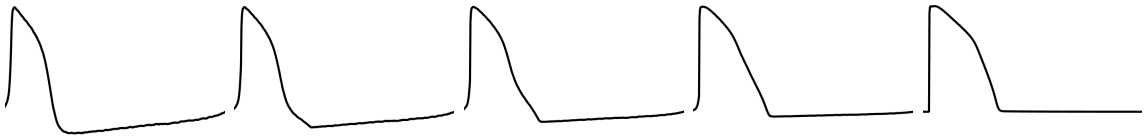

(a) Euclidean Interpolation: $d^2_{\mathcal{M}} = 376.7744$

(b) Algorithm 2: Alternating Gradient Descent in v and f: $d^2_{\mathcal{M}} = 206.8240$, Iterations: 74

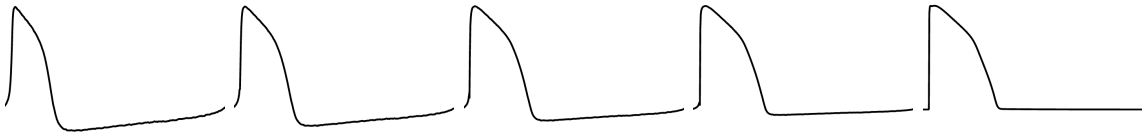

(c) Algorithm 3: $A_{0,S}^{-1}$ updated every iteration: $d^2_{\mathcal{M}} = 208.2602$, Iterations: 72

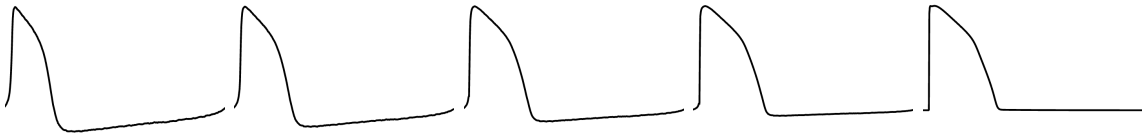

(d) Algorithm 3: $A_{0,S}^{-1}$ updated every 10 iterations: $d^2_{\mathcal{M}} = 207.6308$, Iterations: 38

Figure 3.5: Five samples of the evolution $f(\tau, t_s)$ for $t_s = 0, 0.25, 0.5, 0.75, 1$ showing on columns 1-5, respectively, from left to right, generated by Euclidean interpolation (row 1), Algorithms 2 (row 2) and 2 versions of Algorithm 3 (row 3, 4). The resulting distance between $f(\tau, 0)$ and $f(\tau, 1)$ as well at the number of iterations needed to reach convergence are also displayed.

Table 3.1 compares the Euclidean distance and three methods in terms of their

Table 3.1: Comparison of the metamorphosis algorithms in terms of classification performance and computation time on the Kamp dataset [16].

	Euclidean Distance	Algorithm 2	Algorithm 3 ($n = 1$)	Algorithm 3 ($n = 10$)
1 NN Atrial	**16**/16	13/16	14/16	14/16
1 NN Ventricular	29/36	**36**/36	**36**/36	**36**/36
3 NN Atrial	**16**/16	13/16	14/16	14/16
3 NN Ventricular	29/36	**36**/36	**36**/36	**36**/36
Avg. Comp Time (s)	$<$ **1**	17.0181	15.8265	12.0815

classification performance and computation time on the entire dataset. As previously discussed, classification is done with the 1 nearest neighbor (NN) and 3 NN classifiers, meaning we classify an AP based on the class of the closest 1 or 3 mature prototypes. While the Euclidean distance is indeed the fastest to compute, the metamorphosis distance gives a better classification rate. Moreover, we see that using the closed form update of Algorithm 3 provides a better classification rate at a reduced computation time. Further, we investigated the misclassified atrial cells obtained from Algorithm 3. Specifically, notice from Figure 3.6 that the shape of one of the misclassified atrial cells resembles more a ventricular shape than an atrial one. We observed a similar result for the other misclassified atrial cell. The findings on this dataset suggest that our algorithms provide a significant improvement over the current standard.

3.4.2 Optical Data

Given our success on the small patch clamp dataset, we utilized Algorithm 3 on the much larger dataset consisting of 9 cell aggregates discussed in chapter 2, with

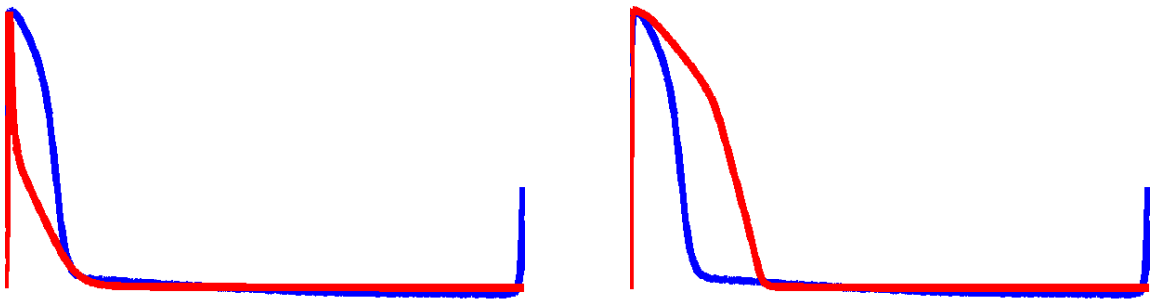

(a) Closest atrial

(b) Closest ventricular

Figure 3.6: Misclassified embryonic atrial AP (blue) and closest mature model APs (red)

action potentials recorded using the optical mapping technique of [17]. The number of APs in each cell cluster ranges from 400 to 1000, and the total number of APs in the dataset equals 6940. Mature prototypes were generated using the same computational models as in the previous experiment. The signals were paced at a rate of 1.5 Hz (cycle length of $\frac{2}{3}$ seconds), and also normalized to have resting potential voltage 0, and maximum voltage amplitude of 1.

We computed the metamorphosis distance using Algorithm 3 with $n = 1$ from each AP in the dataset to each one of the mature prototypes using the same parameters as in the patch clamp experiment. The algorithms were run in 2 8-core computer nodes with 8 hyperthreaded 2.3 GHz CPUs per node. The total time to complete the analysis on the entire dataset was 13 hours, with individual cell clusters taking between 50 and 80 minutes.

Figure 3.7 compares the classification results obtained by a 1-NN classifier with

the Euclidean distance versus the metamorphosis distance for each one of the 9 cell clusters. We omit the 3 NN results as the results are identical. The blue color indicates areas where the APs were classified as atrial, and red indicates areas that were classified as ventricular. While some of the cell clusters present with only 1 phenotype, the majority of the cell clusters present with both phenotypes in varying concentrations, confirming our clustering results in Chapter 2.

(a) Euclidean 1-NN

(b) Metamorphosis 1-NN

Figure 3.7: 1-NN classification results for each of the pixels in each of the cell aggregates of the optical dataset. Blue indicates that the action potential in that pixel is closer to the atrial phenotype, while red indicates the action potential is closer to the ventricular phenotype.

We further investigate the heterogeneous cell clusters. For a pair of these clusters, we show the cell cluster labeling with the overlay of traces for each class in Figure 3.8. For the metamorphosis classification, the two classes show distinct shapes, and they are similar to those described by [16] for embryonic atrial-like and embryonic ventricular-like (Figure 1.5). In comparison, the Euclidean classification fails to capture the distinction between the phenotypes. This confirms that the metamorphosis model is a suitable automated counterpart to manual classification by biologists.

More importantly, it suggests that the metamorphosis model could be used to reliably assess the phenotype statistics of populations of APs. As a consequence, the metamorphosis model may prove insightful to a growing collection of methods that have been derived to isolate a particular phenotype of embryonic cardiomyocytes [9,10,34].

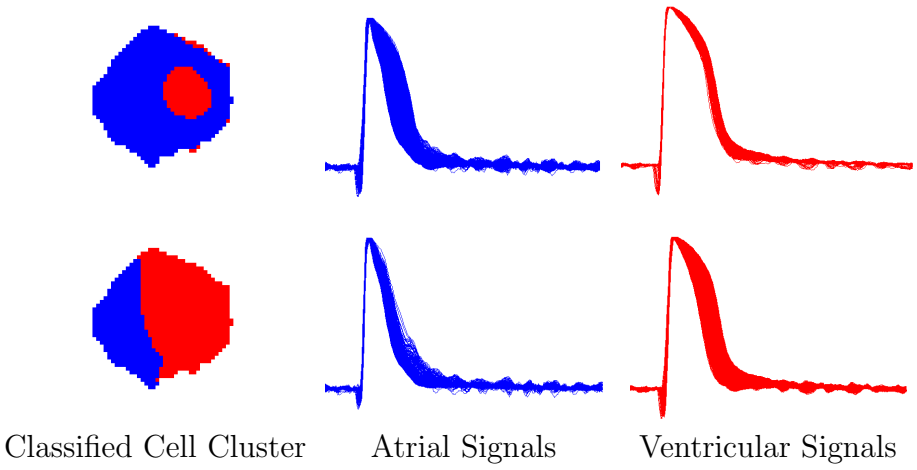

Figure 3.8: Euclidean (top) and Metamorphosis (bottom) cell cluster classifications and their corresponding action potentials. The distribution of phenotypes given by the metamorphosis distance more closely resembles the expert delineation (Figure 1.5) than the distribution given by the Euclidean distance.

Algorithm 3 Discrete Metamorphosis via Direct Image Computation

Input: Template Signal $f_0(\tau)$, Target Signal $f_1(\tau)$, Balance Parameter σ, Number of Evolution Intermediates S, Sobolev Operator L_d, Update frequency n, Convergence parameter δ, and Maximum number of Iterations M.

Initialization
$m = 0$, $d_{-1} = \infty$, $K = L_d^{-1}$.
for $s = 0, \ldots, S$ **do**
    $w(\tau_i, t_s) \equiv 0$, $v(\tau_i, t_s) = Kw(\tau_i, t_s) \equiv 0$
    $f(\tau_i, t_s) = \frac{S-s}{S} f_0(\tau_i) + \frac{s}{S} f_1(\tau_i)$. $\mathbf{f}_s^T = [f(\cdot, t_s)]$
end for
$d_0^2 \leftarrow \sum_{s=0}^{S-1} \|w(\tau_i, t_s)\|_{l_2}^2 + \frac{1}{\sigma^2} \|N_{v_s} \mathbf{f}_{s+1} - \mathbf{f}_s)\|_{l_2}^2$
while $|d_{m-1} - d_m| > \delta$ **AND** $j < M$ **do**
    $m \leftarrow m + 1$
    for $s = 0, \ldots, S-1$ **do**
        Perform gradient descent to update $w(\tau_i, t_s)$ using (3.31).
        $v(\tau_i, t_s) \leftarrow \mathbf{R}(Kw(\tau_i, t_s))$,
        $N_{v_s} \leftarrow$ (3.26).
        for $k = 0, \ldots, S-1$ **do**
            $R_{0,k} \leftarrow N_{v_0} \ldots N_{v_{k-1}}$, $R_{k,S} \leftarrow N_{v_k} \ldots N_{v_{S-1}}$.
        end for
        if $mod(m, n) = 0$ **then**
            $A_{0,S} \leftarrow \sum_{j=0}^{S-1} R_{0,j} R_{0,j}^T$, Calculate $A_{0,S}^{-1}$.
        end if
        $C_{0,0}^T \leftarrow A_{0,S}^{-1}, B_{S,S} = 0$
        for $k = 0, \ldots, S-1$ **do**
            $C_{0,k}^T \leftarrow N_{v_k}^T C_{0,k-1}^T$
        end for
        for $k = S-1, \ldots, 0$ **do**
            $B_{k,S} \leftarrow C_{0,k}^T + N_{v_k} B_{k+1,S}$
            $\mathbf{f}_k \leftarrow B_{k,S} \mathbf{f}_0 + (R_{k,S} - B_{k,S} R_{0,S}) \mathbf{f}_S$
        end for
    end for
    $d_m^2 \leftarrow \sum_{s=0}^{S-1} \|w(\tau_i, t_s)\|_{l_2}^2 + \frac{1}{\sigma^2} \|N_{v_s} \mathbf{f}_{s+1} - \mathbf{f}_s)\|_{l_2}^2$
end while

Output: Metamorphosis family $\mathbf{f}_s$, Metamorphosis velocity $v(\tau_i, t_s)$.

Chapter 4

Automated Clustering and Classification based on Schild's Ladder in the Metamorphosis Metric Space

With the metamorphosis distance, we have provided a method to help predict the phenotype of embryonic cardiomyocytes based on action potential morphology. However, as evidenced by the optical mapping data experiment, it might still be difficult to distinguish, at the decision boundary in particular, what makes a cell more likely to be one phenotype over another. Furthermore, there is some literature that has suggested that the AP shape alone may not be sufficient to identify the

phenotype of a cardiomyocyte [60]. Recall from Figure 1.1 that the action potential is the electrophysiological output of a complex nonlinear system. While some of the factors that influence the output are intrinsic to the cardiomyocyte, such as number of sodium or potassium channels, external factors, such as the stimulus and its rate of occurrence, its connectivity to other cardiomyocytes, or the presence or absence of a drug that targets specific ion channels (Recall Figure 1.3) also play a vital role in the output of the system. The authors of [60] suggest that by taking a snapshot action potential of the system, those factors are effectively being discarded, and it might be those factors that are influencing phenotype classification. By leveraging various external conditions, one should be better able to identify the internal characteristics of the cardiomyocyte and better ascertain its phenotype.

This suggests that multiple action potential recordings of a cardiomyocyte under different environmental conditions may help alleviate these concerns regarding phenotype identity. Let us suppose that we have multiple observations of a cardiac action potential at different external conditions. One approach to identify the phenotype of the query cardiomyocyte is to classify each of the individual observations, and then aggregate the outputs via a voting or other aggregation system to assess the phenotype of the heart cell. This approach is simple, and we have provided a framework to do the first part of this strategy in Chapter 3. However, this requires training data at each of the observed conditions, which could lead to a very large computational overhead.

A more sophisticated approach would be to analyze the trajectory of the action potential with respect to the alteration of external conditions. Since the different phenotypes of cardiomyocytes are intrinsically governed by different sets of differential equations, it is not unreasonable to believe that, under an appropriate external condition, the cardiomyocytes of different phenotypes will exhibit different reactions to the change in the environment. Thus subjecting cells to a common change in the environment will provide insight into the intrinsic features of a cardiomyocyte that should be shared amongst cardiomyocytes of the same phenotype. This belief is already being utilized in anti-arrhythmic drugs that target fibrillation in one chamber over another [61]. By identifying the discriminating factors of drug responses across phenotypes of adult cardiomyocytes, we should be able to transfer that insight to the embryonic cardiomyocytes in order to augment the classification procedure.

Besides phenotype classification, this analysis can help answer additional scientific inquiries. Identifying drug sensitive/resistant populations is of particular interest in the cardiac community. For example, Braam and colleagues suggested that there was a repolarization reserve, or redundancy, linked to the slow acting potassium rectifier that has an effect on how resilient cardiomyocytes are to alterations in the fast acting potassium rectifier channel [62]. By observing the trajectories of the action potential with respect to blockages in fast rectifier, we should be able to verify this observation. Providing additional methods for assessing the drug response of a population may prove helpful in isolating other such characteristics of cardiomyocytes.

As methods for isolating cardiomyocyte populations continue to improve, another question worth asking along the lines of the previous drug screening examples is whether cardiomyocytes derived by these specific cell lines and protocols react similarly to a treatment or drug or not, or how a certain cell line interacts with different classes of drugs. Some commercial cell lines have performed some analysis of the reaction of their products reaction to certain commonly investigated drugs [4,63], but this is still largely unexplored electrophysiologially. Comparatively assessing whether two cell lines differ in their reaction to a drug, and where in the action potential morphology the lines differ, provides information about the mechanistic differences of the derived cardiomyocytes that biologists can use to identify which cells will be better suited for a particular treatment strategy. Additionally, there has been preliminary work in utilizing machine learning techniques based on random forests to identify certain beating rate modulators by their effect on the action potential [64]. Expanding upon these ideas could provide insight into learning a prototypical response to a specific drug, and utilizing those prototypes to identify what the effect of a newly developed drug therapy has on these particular tissues of cardiomyocytes.

Currently, similar to the classification problem discussed in the previous chapter, the discriminating factors of these drug studies are the action potential features such as duration, amplitude, and slope. Not only does this analysis discard most of the signal, which could provide more insight into the machinery of the drug action, but also it is not clear how appropriate these comparisons are. A comparison based on

raw change in features does not transfer across members of the population: a change of 100 ms in APD_{50} means different things if the original APD_{50} was 300 ms vs 450 ms. The simplest way around this is to look at the percentage of change in the APD features, which is unit-less, but even this is still sensitive to initial value of the feature.

The metamorphosis algorithm from Chapter 3 not only provides us a distance from which to assess similarity between action potentials, but also defines a trajectory describing the transformation between the template and the target. In particular, the velocity fields that define the diffeomorphism and the infinitesimal changes in the template provided by this trajectory may provide additional information that discriminates one cell from another. Utilizing these metamorphosis features may provide enough depth to be able to characterize not only the effect of a treatment, but also the subtleties that differentiate various treatments or individuals.

However, if the standard action potential features were sensitive to the initial condition of the action potential, the sensitivity of these metamorphosis features are far worse. Given the temporal structure of the action potential and the corresponding metamorphosis features, a velocity field that acts at a time τ might mean very different things depending on which phase of the action potential cycle the cardiomyocyte is in. How then, do we make these metamorphosis features comparable in order to identify meaningful functional differences in cardiomyocytes? In this chapter, I will describe how differential geometry provides a solution to the problem, and provide a novel implementation of this solution in metamorphosis space.

4.1 Parallel Transport for Metamorphosis

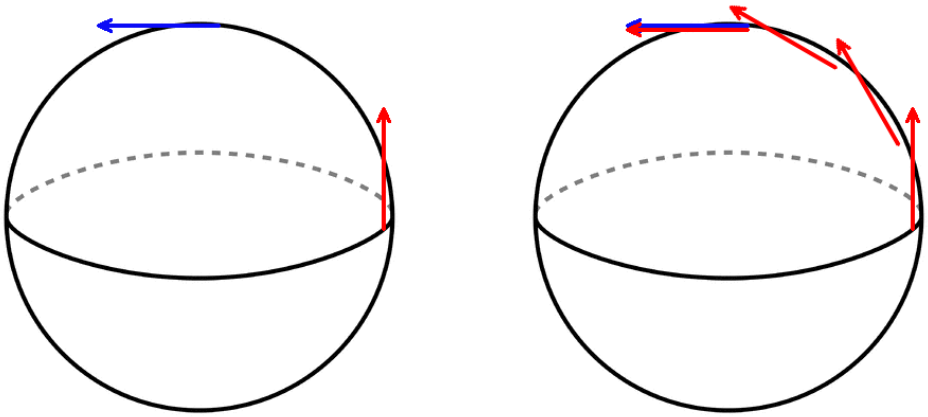

(a) Erroneous Comparison (b) Appropriate Comparison

Figure 4.1: Comparing observations on the sphere. (left) The two vectors appear different at face value, but this difference stems from the fact that they are encoding information from different parts of the sphere, (right) By moving the red vector to the location of the blue vector while maintaining its tangent information, one can see that the vectors correspond to similar behaviors

One of the consequences of identifying the metamorphosis metric as a method for action potential analysis is that action potentials lie in a curved manifold. In order to see, in another way, how features can be influenced by where in the manifold the feature is coming from, let us first look at a much simpler manifold, like a sphere (Figure 4.1). The blue and red vectors appear to be different; they point in different directions. However, they both appear to encode the same tangent action. Under the normal Riemannian metric on the sphere, the analog of the red vector at the location of the red vector is a vector that points to the left, preserving its tangent information. Once we look at the two vectors at the same point on the circle, it is clear that the

vectors correspond to similar actions.

Following the example of the sphere, in order to compare different trajectories of drug activity, we should attempt to carry the trajectories to a common frame of reference in a way that preserves the information of each trajectory. In the differential geometry literature, the notion of preserving the information in the trajectory is defined by requiring the final trajectory at the new frame of reference to be parallel to the original trajectory at its original starting point. Specifically, if V_0 is a vector field at a point $s(t_0)$, and we have a differentiable curve $s(t), t \in [0, 1]$, then $V(t)$ is parallel at time point $s(t)$ to V_0 if:

$$\nabla_{\dot{s}(t)} V(t) = 0, \tag{4.1}$$

where ∇_y is the covariant derivative in the direction y, and $\dot{s}(t)$ is the (time) derivative of the $s(t)$. This is appropriately known as parallel transport [65]. In the spherical example (Figure 4.1), in order to transport the vector from the equator of the sphere to the north pole, under the standard spherical metric, the vector has to rotate while traveling along the arc of the circle. As a result, we can see that the resulting transported vector is indeed similar to the other vector defined at this frame of reference. Unfortunately, parallel transport in the space of action potentials with the metamorphosis metric is not that simple.

4.1.1 Defining an Appropriate Atlas, or Frame of Reference

The first step of parallel transport is picking an atlas, or a frame of reference where all members of the population are going to be compared. A simple choice for this is a random member of the population. On the plus side, this requires no computational overhead, which in larger populations could prove useful. One major concern is the biased nature of this selection. If the randomly chosen reference is an outlier of the population, then the resulting parallel transport may include features of the outlier that are not indicative of the change in the environment, but rather the behavior of the outlier, potentially harming the resulting analysis.

A typical choice for a reference is the "center" of the data. This is typically found using the Fréchet mean. Given a collection of data points $\{x_i\}$ and a distance metric d on the population space, the Fréchet mean $\bar{x}$ of the population is the point that minimizes the total square distance to population. Mathematically, $\bar{x} = \arg\min_p \sum_i d^2(p, x_i)$. It is not difficult to show that the well known arithmetic mean is the Fréchet mean under the Euclidean distance.

How do we calculate the Fréchet mean in the space induced by the metamorphosis

metric $d_{\mathcal{M}}$? Let us start with a collection of P action potentials $\{\tilde{f}_i\}_{i=1}^P$:

$$\begin{aligned}\tilde{f} &= \arg\min_f \sum_i^P d^2_{\mathcal{M}}(f, \tilde{f}_i) \\ &= \arg\min_f \sum_i^P \min_{\substack{v_i, f_i \\ f_i(\tau,t_0)=f \\ g_i(\tau,t_S)=\tilde{f}_i}} \sum_{s=0}^{S-1} (\|v_i(\tau,t_s)\|_V^2 + \frac{1}{\sigma^2}\|f_i(\tau + v_i(\tau,t_s), t_{s+1}) - f_i(\tau,t_s)\|_{l^2}^2).\end{aligned} \tag{4.2}$$

We have an optimization problem over f, v_i, and $\tilde{f}_i$. When f is fixed, updating v_i and $\tilde{f}_i$ is done by calculating the metamorphosis between f and $\tilde{f}_i$. Either of the two algorithms discussed in the previous chapter will work here. Once the optimal v_i and f_i are found. We fix them in order to update f. The only terms in this sum that involve f are those at $s = 0$, since $f = \tilde{f}_i(\tau, 0)$. The resulting optimization problem is:

$$\tilde{f} = \arg\min_f \sum_i^P \frac{1}{\sigma^2}\|\tilde{f}_i(\tau + v_i(\tau, t_0), t_1) - f\|_{l^2}^2). \tag{4.3}$$

Taking the derivative with respect to f and setting it to 0, shows the update for f is the arithmetic mean of the undeformed first interpolants: $f^* = \frac{1}{P}\sum_i^P \tilde{f}_i(\tau + v(\tau, t_0), t_1)$. After updating f it is fixed in order to refine the estimates of v_i and $\tilde{f}_i$. The initial choice for f is the Euclidean (arithmetic) mean of $\{\tilde{f}_i\}$, and can be used if the variations in the population are small. Notice that the resulting algorithm for computing the Fréchet mean entails computing P metamorphoses at every iteration, which, depending on the number of iterations, could be computationally expensive, but in

general this algorithm requires fewer iterations, both internally with the metamorphosis updates, and externally over the update of f because the deformations are typically small.

Another way to pick a good atlas is by picking a member of the population more intelligently than by chance. Instead of calculating a mean, one could find a medoid (similar to a median) of the data. Like the Fréchet mean, a medoid minimizes the total distance to the entire population, but with the additional constraint that it has to be a member of the population. More specifically, the medoid M of a set of action potentials $\{\tilde{f}_i\}_i = 1^P$ is defined as:

$$M(\{\tilde{f}_i\}) = \underset{\{\tilde{f}_i\}_i^P}{\arg\min} \sum_j^P d^2(\tilde{f}_i, \tilde{f}_j). \tag{4.4}$$

Notice that the medioid can be a computed in a straightforward manner by first calculated the matrix of squared pairwise distance across the dataset (see Chapter 2) and then finding the row of the matrix with the smallest sum. While this can be computational intense when the distance d is the metamorphosis distance, this can be still more efficient that computing the Fréchet mean when the latter requires a large number of iterations. Another advantage of the medoid is that its shape corresponds to that of an action potential, as the medoid is equal to one of the action potentials in the dataset.

A comparison of the mean and medoid for a population of data in shown in Figure

4.2. The two references differ primarily in the phase 3 repolarization, as the medoid has a steeper repolarization phase. The flatter slope in the mean is likely due to the starting initialization being the Euclidean mean. In this case, since the medoid represents a member of the population, and the population doesn't vary wildly, it is a better reference to use, as the phase 3 repolarization in the medoid is more coherent with the rest of the data than the Fréchet mean.

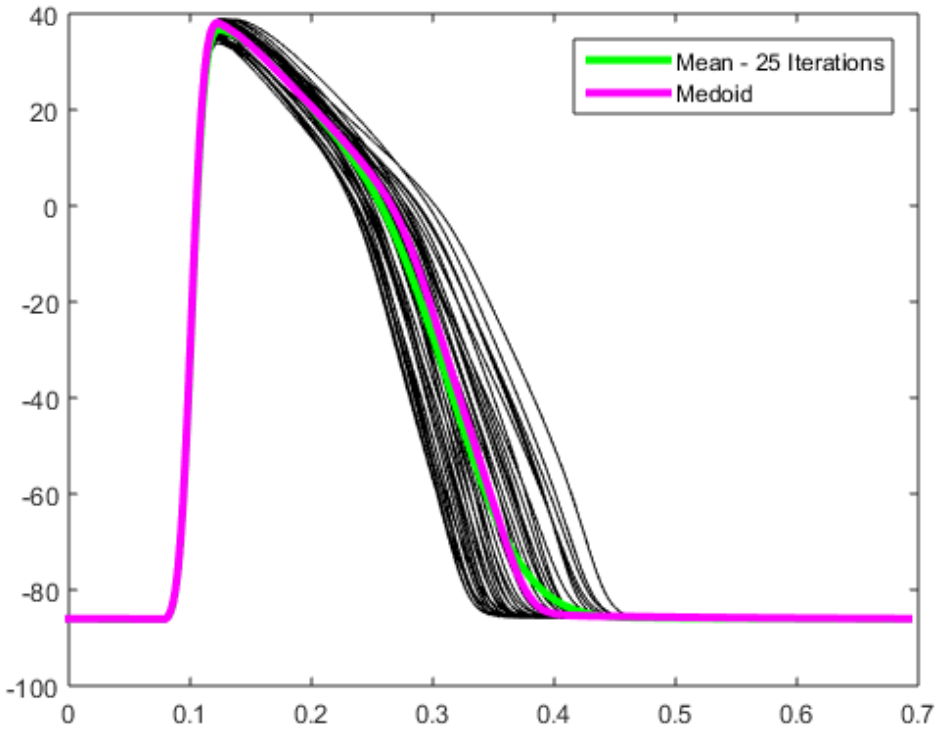

Figure 4.2: Mean Vs. Medoid of a population of 40 simulated ventricular cardiomyocyte action potentials using [45]

4.1.2 Schild's Ladder in the Metamorphosis domain

Now that we have appropriately chosen a frame of reference, we can discuss how to perform parallel transport to the chosen atlas. Like the circle, we have shown back in chapter 3 that the space of action potentials is not flat, since linear combinations of action potentials do not generate additional action potentials. However, unlike the sphere, the shape of the space of action potentials is unknown. This makes finding parallel transport in this space, like the case of the sphere, very difficult. If we, instead, take small enough steps towards the atlas, we can approximate parallel transport with parallelogrammoids, or parallelograms in curved spaces. This was the idea of Alfred Schild, who formalized this mathematically into a construction now appropriately called Schild's Ladder [66, 67].

We now formally discuss the construction of Schild's Ladder. Please refer to Figure 4.3 for additional visual aid. Suppose C_0 is original template point, and T_0 is original target such that the segment C_0T_0 is to be transported to a reference C_N. Let $C_1, \ldots, C_{N-1}$ be intermediates between C_0 and C_N. Given C_0, T_0, and C_1, first find the shortest distance (geodesic) path between T_0 and C_1. This is depicted in gray in Figure 4.3. This defines one diagonal of the parallelogram. Define the midpoint of this segment as $H_{0.5}$. In order build the other diagonal, calculate the geodesic between C_0 and $H_{0.5}$ and calculate the arc length l of the geodesic. Since one of the defining

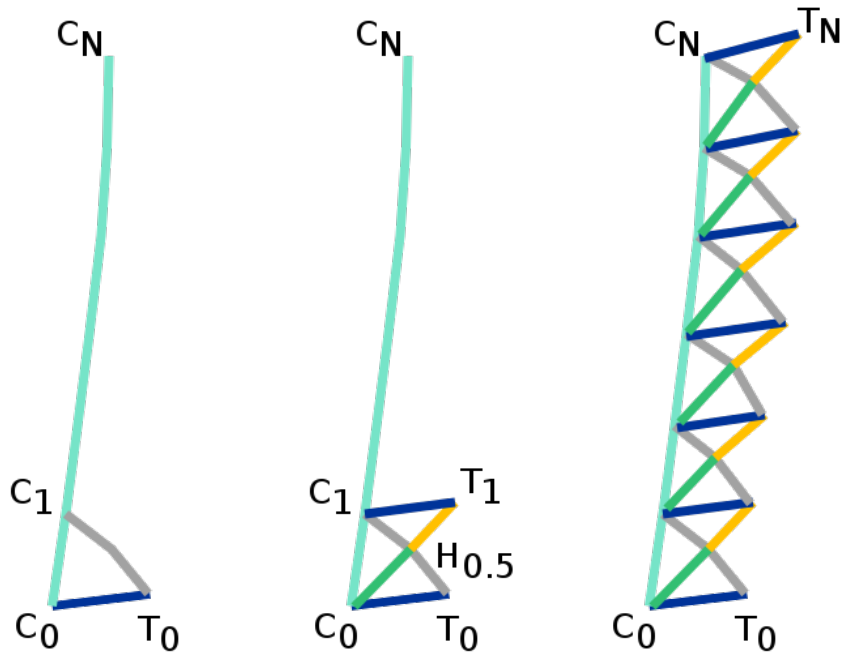

Figure 4.3: Schild's Ladder Schematic

properties of parallelograms is that the intersection of the diagonals, in this case $H_{0.5}$, bisects both diagonals, extend this new geodesic until it has length $2l$. The point where this geodesic ends is the approximation of the transported target T_1 at the new starting point C_1, e.g. C_1T_1 approximates the parallel transport of C_0T_0 at C_1. Repeat this process until you arrive at C_N, then T_N defines the transported target. The accuracy of the approximation is dependent on the curvature of the segments connecting C_i and C_{i+1}, so smaller steps are better as the curvature is closer to flat.

How does one apply Schild's Ladder to the space defined by the metamorphosis distance? First, we need to identify the geodesic in the metamorphosis metric between the original template C_0 and the point of reference C_N, as well as the geodesic

between C_0 and T_0. This can be done via the metamorphosis algorithms (Algorithms 2 and 3) presented in chapter 3. The intermediates in the C_0C_N metamorphosis can be used to define C_i. We can use the metamorphosis algorithms to find the geodesic between C_{i+1} and T_i, and from C_i to $H_{i+0.5}$. The only remaining step is to determine how to extend a previously defined geodesic path. This process is typically referred to as geodesic shooting. The general idea is that, by definition, minimizing the metamorphosis distance between two action potentials creates a geodesic. So by defining relationships achieved at the optimum of the metamorphosis algorithm, and following those relationships at subsequent, extrapolated points in the path, the resulting path is still a geodesic. So let us look at the governing equations of our boundary value metamorphosis problem (see equation (3.24)):

$$\begin{aligned}\nabla_{w(\tau,t_k)}E &= w(\tau,t_k) + \frac{1}{\sigma^2}K((f(\tau + Kw(\tau,t_k),t_{k+1}) - f(\tau,t_k))\frac{\partial f}{\partial \tau}(\tau + Kw(\tau,t_k),t_{k+1}) \\ &\approx w(\tau,t_k) + \frac{1}{\sigma^2}K(N_{v(\tau,t_k)}f(\tau,t_{k+1}) - f(\tau,t_k))\frac{\partial}{\partial \tau}N_{v(\tau,t_k)}f(\tau,t_{k+1}), \qquad (4.5)\end{aligned}$$

$$w(\tau,t_k) = Lv(\tau,t_k), \qquad (4.6)$$

$$\nabla_{f(\tau,t_k)}E \approx \frac{1}{\sigma^2}(N^T_{v(\tau,t_{k-1})}z(\tau,t_{k-1}) - z(\tau,t_k)), \qquad (4.7)$$

$$z(\tau,t_k) = N_{v(\tau,t_k)}f(\tau,t_{k+1}) - f(\tau,t_k). \qquad (4.8)$$

Recall that w and z are dummy variables defined to simplify the optimization (see 3.31 and 3.32) Let w^*, v^*, f^*, and z^* be the optimal parameters of any boundary value metamorphosis problem. These parameters then must set the gradients in the

previous set of equations to 0, since they are optimal. This leads to the following set of equations at the optimum:

$$0 \approx w^*(\tau, t_k) + \frac{1}{\sigma^2} K(N_{v^*(\tau,t_k)} f^*(\tau, t_{k+1}) - f^*(\tau, t_k)) \frac{\partial}{\partial \tau} N_{v^*(\tau,t_k)} f^*(\tau, t_{k+1}),$$

$$w^*(\tau, t_k) = L v^*(\tau, t_k),$$

$$0 \approx \frac{1}{\sigma^2}(N^T_{v^*(\tau,t_{k-1})} z^*(\tau, t_{k-1}) - z^*(\tau, t_k)),$$

$$z^*(\tau, t_k) = N_{v^*(\tau,t_k)} f^*(\tau, t_{k+1}) - f^*(\tau, t_k).$$

Rearranging the equations leads to the following series of updates, which allow the continuation of the geodesic:

$$w^*(\tau, t_k) \approx -\frac{1}{\sigma^2} K(z^*(\tau, t_k)) \frac{\partial}{\partial \tau}(z^*(\tau, t_k) + f^*(\tau, t_k)), \tag{4.9}$$

$$v^*(\tau, t_k) = K w^*(\tau, t_k), \tag{4.10}$$

$$z^*(\tau, t_k) \approx N^T_{v^*(\tau,t_{k-1})} z^*(\tau, t_{k-1}), \tag{4.11}$$

$$N^T_{v(\tau,t_k)} N_{v^*(\tau,t_k)} f^*(\tau, t_{k+1}) = N^T_{v(\tau,t_k)}(z^*(\tau, t_k) + f^*(\tau, t_k)). \tag{4.12}$$

There are two things of note from these equations. First, the template evolution $z(\tau, t_k) = N_{v(\tau,t_k)} f(\tau, t_{k+1}) - f(\tau, t_k)$ is the primary driver of the extension equations. Given $z(\tau, t_k)$ and $f(\tau, t_k)$ one can calculate $v(\tau, t_k)$ and $f(\tau, t_{k+1})$. This allows us to calculate $z(\tau, t_{k+1})$, continuing the geodesic shooting procedure. Thus, $z(\tau, t_k)$ is often called the momentum of the metamorphosis. Second, we have multiplied the

last equation by $N^T_{v(\tau,t_k)}$ to improve numerically solving for f_{k+1} via methods like preconditioned conjugate gradient. Even with the help of these numerical methods, the continuation of f runs a risk of instability, since it is not immediately clear that N_v is invertible, or if the inverse is well-conditioned if it does exist. Since N_v is computed via linear interpolation, invertibility can be ensured, at least, when the maximum absolute value, or the ∞-norm, of the velocity field is less than half of the discretized sampling of the action potential signal. Following this condition ensures that N_v is strictly diagonally dominant, a sufficient condition for invertibility. As we will show, however, invertibility is insufficient. Figure 4.4 compares 2 qualitative constant velocity fields, one "small" velocity field such that the interpolation matrix is close to the identity, and a larger velocity field that is still invertible. While N_v appears innocuous, the difference in the inversion of $N_v^T N_v$ is fairly significant, as we see oscillatory behavior as the velocity increases. This oscillation does not appear to manifest when the signal is smooth, but as soon as one deviates from smoothness, as in the triangle example in Figure 4.5, this behavior can cause a wildly unstable output. Thus when the velocity is small, the closer to the identity we are, and the more likely we are to avoid these unstable extensions, especially at aggressively varying portions of the action potential like phase 0 upstroke.

With the equations for geodesic shooting defined, we can construct the Schild's Ladder, using the procedure summarized in Algorithm 4. Since, as we have discussed, smaller geodesic extensions are typically more stable than larger geodesic extensions,

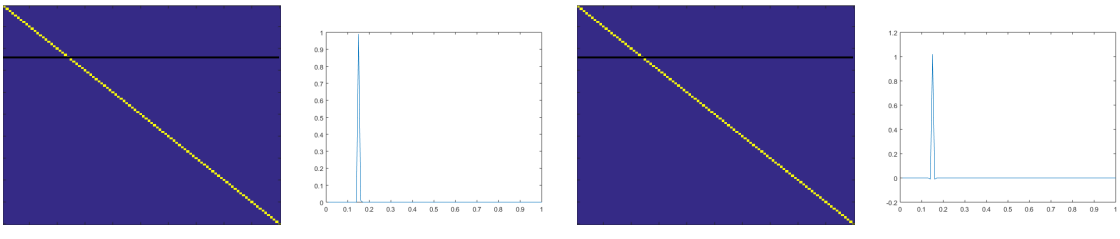

(a) $N_{v_{low}}$ (b) $(N_{v_{low}}^T N_{v_{low}})^{-1}$

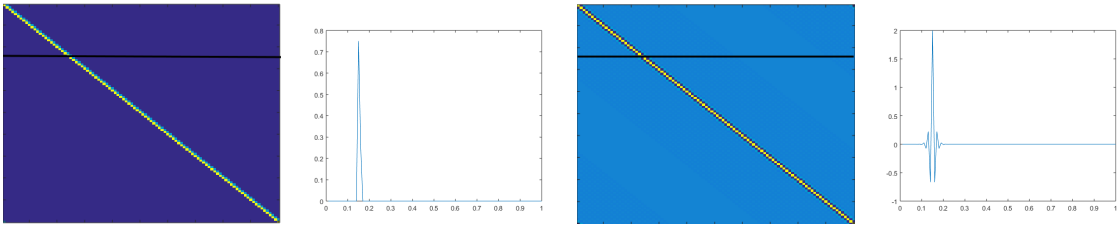

(c) $N_{v_{high}}$ (d) $(N_{v_{high}}^T N_{v_{high}})^{-1}$

Figure 4.4: Comparison of N_v and $(N_v^T N_v)^{-1}$ for two different constant velocities: a small velocity field, where N_v is close to the identity (top row), and a larger velocity field which deviates from the identity (bottom row). The plots show values for the specific row of the matrix on its left coded in black.

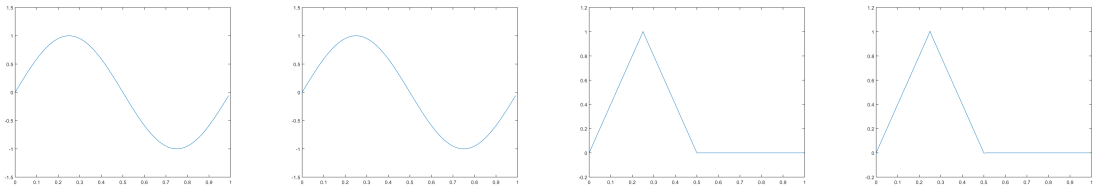

(a) $f_2^{v_{low}}$, sinusoid (b) $f_5^{v_{low}}$, sinusoid (c) $f_2^{v_{low}}$, triangle (d) $f_5^{v_{low}}$ triangle

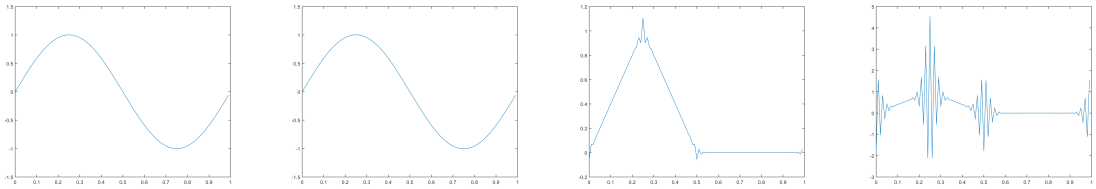

(e) $f_2^{v_{high}}$, sinusoid (f) $f_5^{v_{high}}$, sinusoid (g) $f_2^{v_{high}}$, triangle (h) $f_5^{v_{high}}$ triangle

Figure 4.5: Stability analysis: ladder steps of test signals with "low" (top) and "high" (bottom) velocity fields

Algorithm 4 Schild's Ladder Computation for Metamorphosis

Input: Original Template $C_0(\tau)$, Original Target $T_0(\tau)$, New Template $C_S(\tau)$, and metamorphosis parameters: Balance σ, Sobolev norm L_d, Number of Evolution Intermediates S, Convergence parameter δ, and Maximum number of Iterations M.

1: **Initialization**
2: Compute the metamorphosis (v_{cc}, f_{cc}) between $C_0(\tau)$ and $C_S(\tau)$. Let $C_i(\tau) = f_c(\tau, t_i), i = 1, \ldots, S-1$.
3: Compute the metamorphosis (v_{ct}, f_{ct}) between $C_0(\tau)$ and $T_0(\tau)$. Pick $\tilde{T}_0(\tau) = f_t(\tau, t_1)$.
4: **for** $i = 0, \ldots, S-1$ **do**
5: Compute the metamorphosis (v_i, f_i) between $C_{i+1}(\tau)$ and $\tilde{T}_i(\tau)$.
6: $H_{i+0.5}(\tau) = f_i(\tau, t_{S/2+1})$.
7: Compute the metamorphosis $(v_{i+0.5}, f_{i+0.5})$ between $C_i(\tau)$ and $H_{i+0.5}(\tau)$.
8: Use the geodesic extension equations (4.9) - (4.12) to extend the $(v_{i+0.5}, f_{i+0.5})$ metamorphosis from S intermediates to $2S$ intermediates.
9: $\tilde{T}_{i+1}(\tau) = f_{i+0.5}(\tau, t_{2S})$
10: **end for**
11: Compute the metamorphosis (v_s, f_s) between $C_S(\tau)$ and $\tilde{T}_S(\tau)$.
12: Use the geodesic extension equations 4.9 - 4.12 to extend the (v_s, f_s) metamorphosis from S intermediates to S^2 intermediates. $T_S(\tau) = f_s(\tau, t_{S^2})$.

Output: Transported Target $T_S(\tau)$.

included in the algorithm is the following step. Instead of transporting the entire geodesic from C_0 to T_0, let $\tilde{T}_0$ be the first intermediate along the geodesic connecting C_0 and T_0. Then transport the geodesic $C_0\tilde{T}_0$ via the ladder algorithm to obtain $\tilde{T}_N$. After calculating the transported geodesic, obtain T_N by extending $\tilde{T}_N$ back to the appropriate length. The benefit of this additional step is that instead of calculating N large geodesic extensions, this approach calculates N much smaller extensions and finishes with only 1 large extension, thus improving stability of the algorithm.

4.2 Experiments

In this section, we present several experiments to show how approximate parallel transport via Schild's Ladder can be applied to typical studies in the biology community, and how the resulting analyses are improved by this algorithm and the metamorphosis features.

4.2.1 Drug Sensitivity Analysis

In this first experiment, we use simulated action potential data to show the applicability of the metamorphosis features and our derived Schild's ladder algorithm to identify a drug sensitive population of cardiomyocytes. More specifically, we generate simulated 40 ventricular action potentials via the O'Hara ventricular model [45] by modulating the conductances of the ion channels between 80% and 120%. Half of the population, however, have an additional systemic 50% conductance block of the slow acting potassium rectifier (K_s) channel, simulating a loss in the repolarization reserve as described in [62]. The resulting population is then subjected to increasing conductance blocks in the fast action potassium rectifier (K_r), starting with at a 25% blockage and finishing at 45%. The population under the affect of the K_r blockage is shown in Figure 4.6. The action potentials are smoothed with a Gaussian filter to help the stability of the ladder algorithm.

Without knowing which action potentials belong to the "at-risk" group, we would

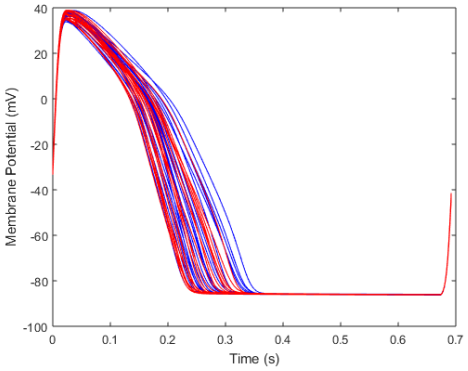

(a) 0% gKr Block

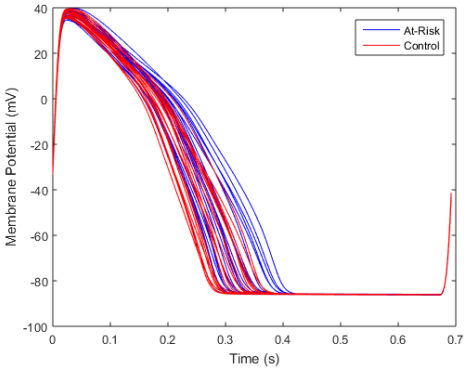

(b) 25% gKr Block

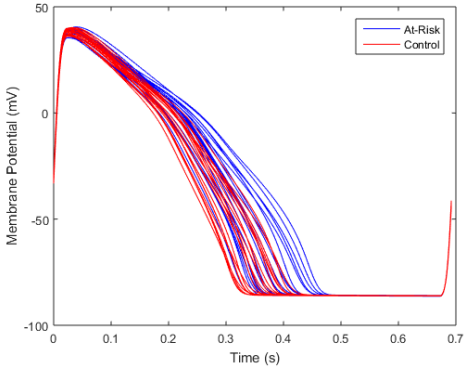

(c) 45% gKr Block

Figure 4.6: Population of 40 APs synthesized with the O'Hara model. The control population consists of 20 APs synthesized by modulating the conductance of the ion channels between 80% and 120%. The at-risk population consists of 20 APs with an additional 50% conductance block of the slow action potassium channel. The entire population is recorded at 0%, 25%, and 45% conductance blocks of the fast action potassium channel.

like to use the change in the action potential morphology to identify which cardiomyocytes have a reduced repolarization reserve (i.e. blockage in the K_s channel). Following the established standard of looking at the action potential duration features, we obtain the APD_{30}, APD_{50}, APD_{80} features of the signals at each stage of the gK_r blockage. To follow the conventional approach to identification, the percent change in these features between no block and 25% block and between 25% and 45% are shown in Figure 4.7. In order to see if they separate into the normal and at-risk groups automatically, we perform our spectral clustering algorithm into 2 groups with these difference features. Specifically, the percent changes are concatenated into a vector $y = [\%\Delta APD_{30}^{0\to25}, \%\Delta APD_{50}^{0\to25}, \%\Delta APD_{80}^{0\to25}, \%\Delta APD_{30}^{25\to45}, \%\Delta APD_{50}^{25\to45}, \%\Delta APD_{80}^{25\to45}]$, and the weight $w(i,j)$ is Gaussian kernel of the Euclidean distance between the vectors: $w(i,j) = e^{\frac{\|y_i - y_j\|^2}{\sigma^2}}$, with σ^2 being the average squared distance over all pairs. The ground truth labeling and the output of the spectral clustering result is shown in Figure 4.7. Notice that these features are not enough to resolve the difference between the two populations. In addition to looking at the accuracy of the labeling generated by spectral clustering, by building the graph for the spectral clustering process, we can look at the cost of cutting the graph according to the ground truth labeling. Recall that the cut cost ranges from 0 to 1 with lower meaning better separation between the defined groups. The cut cost is shown for the APD features in Table 4.1, where it does not suggest a particularly strong separation between the control and at-risk populations.

Table 4.1: Accuracy of predicted labels and the cut cost of spectral clustering separation based on APD and Metamorphosis feature sets

Feature Set	Predicted Label Result	SCC Result
%Δ APD Features	27/40	0.4664
Momentum $z(\tau, 0)$: Pre-Ladder	23/40	0.5065
Momentum $z^{SL}(\tau, 0)$: Post-Ladder	24/40	0.4862
Velocity $v(\tau, 0)$: Pre-Ladder	32/40	0.4180
Velocity $v^{SL}(\tau, 0)$: Post-Ladder	**39**/40	**0.3486**

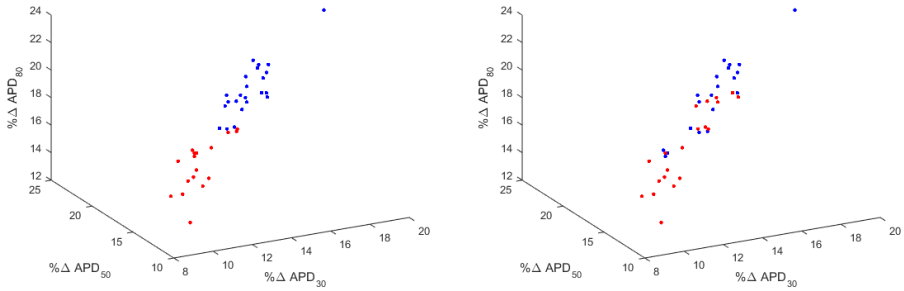

(a) 0% → 25% gKr Block

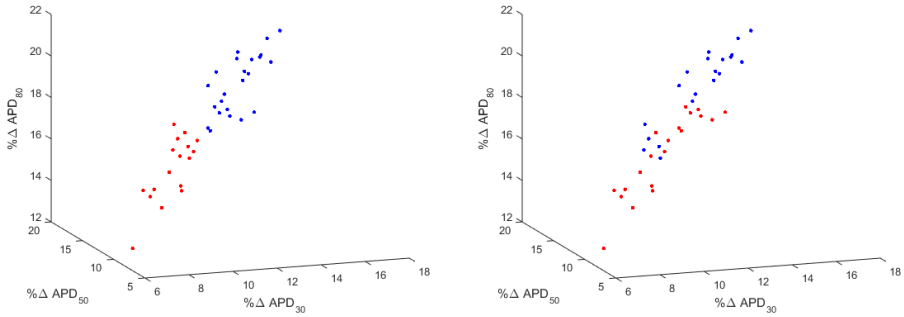

(b) 25% → 45% gKr Block

Figure 4.7: Plots of the percent change in APD_{30}, APD_{50}, and APD_{80} as the blockage in the conductance of the Kr channel increases from 0% to 25% (top) and 25% to 45% (bottom). The left figures show the assignments of the data according to the spectral clustering algorithm, while the right figures show the true assignments. Red indicates the control population, while blue indicates the at-risk population

In order to cluster based on the metamorphosis features, we first calculate a metamorphosis between the action potentials of no and 25% block and 25% and 45% block via Algorithm 3. After computing this metamorphosis, we look at the initial momentum ($z(\tau, t_0) = N_{v(\tau,t_0)} f(\tau, t_1) - f(\tau, t_0)$) and the initial velocity ($v(\tau, t_0)$) as two different features of the metamorphosis. Similar to the APD features, spectral clustering is performed on these metamorphosis features by concatenating the features and using the Gaussian kernel, but the Euclidean distance is used on the momentum, while the Sobolev norm is used on the velocity fields. The table with the resulting clustering accuracy and the cut cost of the ground truth labeling is provided in Table 4.1. Observe that the momentum does not improve the clustering result, but the velocity appears to the resolve the difference between the trajectories better than the APD features. However this result is not as helpful because the features are not aligned. We then transport the features to the medoid of the population via Schild's Ladder (Figure 4.2), and denote the transported momentum and velocity features as $z^{SL}(\tau, 0)$ and $v^{SL}(\tau, 0)$, respectively. Spectral clustering is performed on the transported features, and while this does not dramatically improve the momentum feature's ability to segregate the at-risk population from the normal population, the clustering based on the velocity feature greatly improves after parallel transport, performing the best of all features at self separating the two populations.

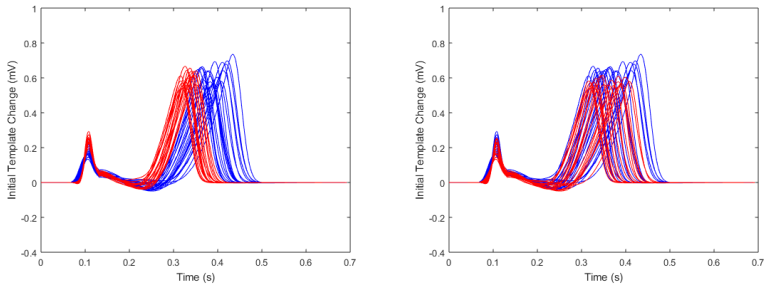

(a) 0% → 25% gKr Block: Pre-Ladder $z(\tau, 0)$

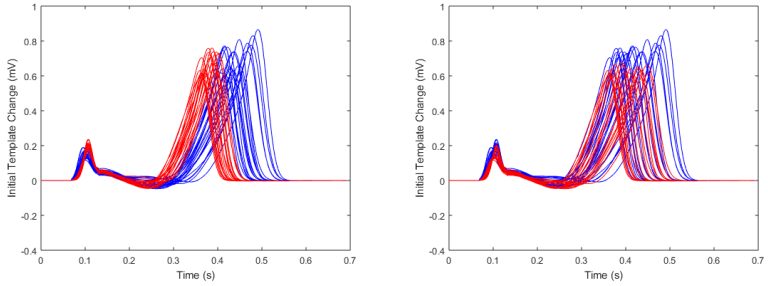

(b) 25% → 45% gKr Block: Pre-Ladder $z(\tau, 0)$

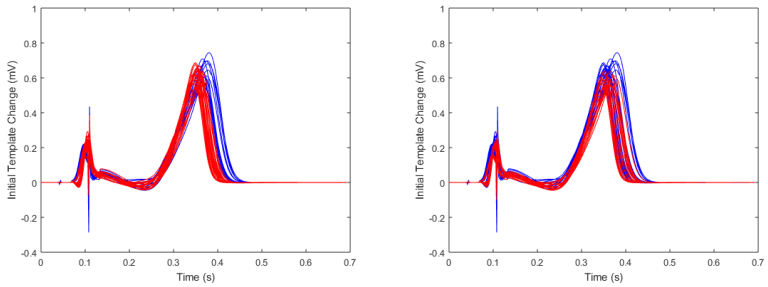

(c) 0% → 25% gKr Block: Post-Ladder $z^{SL}(\tau, 0)$

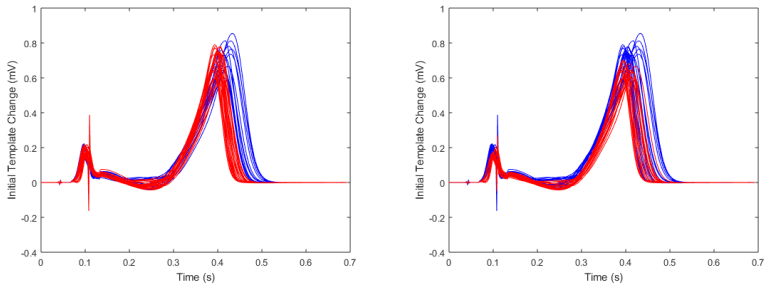

(d) 25% → 45% gKr Block: Post-Ladder $z^{SL}(\tau, 0)$

Figure 4.8: Metamorphosis momentum features before (a,b) and after (c,d) parallel transport. The hypothesized at-risk population as a result of the spectral clustering algorithm is presented in the left column, while the ground truth at-risk population is presented in the right column. Red indicates control population, while blue indicates at risk population.

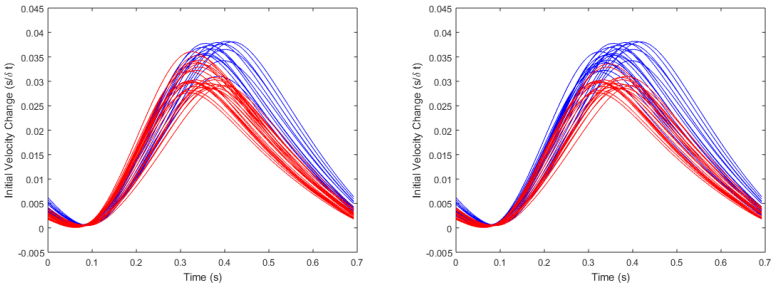

(a) 0% → 25% gKr Block: Pre-Ladder $v(\tau, 0)$

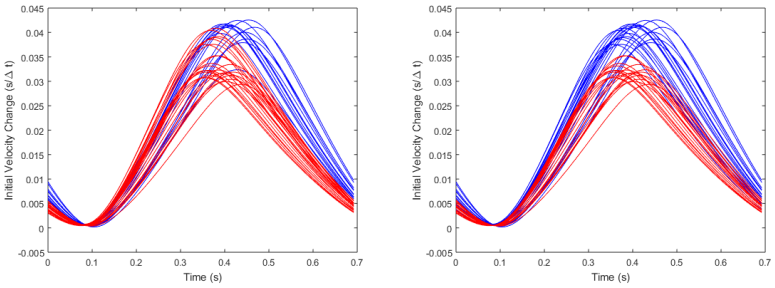

(b) 25% → 45% gKr Block: Pre-Ladder $v(\tau, 0)$

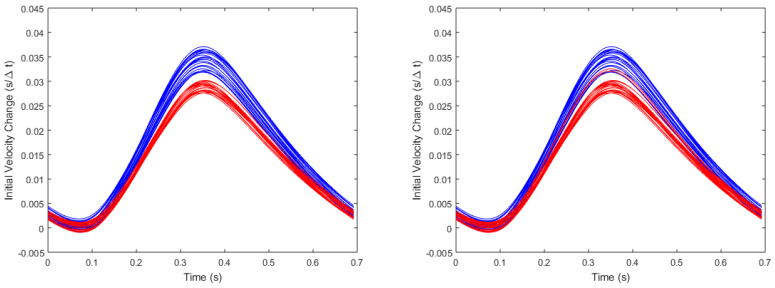

(c) 0% → 25% gKr Block: Post-Ladder $v^{SL}(\tau, 0)$

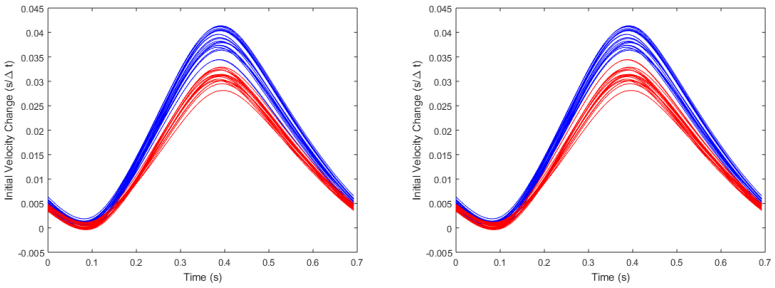

(d) 25% → 45% gKr Block: Post-Ladder $v^{SL}(\tau, 0)$

Figure 4.9: Metamorphosis velocity features before (a,b) and after (c,d) parallel transport. The hypothesized at-risk population as a result of the spectral clustering algorithm is presented in the left column, while the ground truth at-risk population is presented in the right column. Red indicates control population, while blue indicates at risk population.

4.2.2 Drug Comparison on a Single Cell Line

The purpose of this experiment is to define a drug response profile for a particular population of cardiomyocytes derived from a particular cell line/protocol. Populations of Cor.4U cardiomyocytes (Axiogenesis, Cologne, Germany) are isolated into monolayers and each of the monolayers and are subjected to a different concentration of various ion channel blockers and drug compounds, some of which have been previously implicated in arrhythmias [68–71]. Action potential recordings from the cells before and after their selected treatment is provided in Figure 4.10. The goal of the experiment is to potentially identify groups of drugs that effect the cardiomyocyte action potential in a similar way, and whether that is an indicator of the arrhythmic potential.

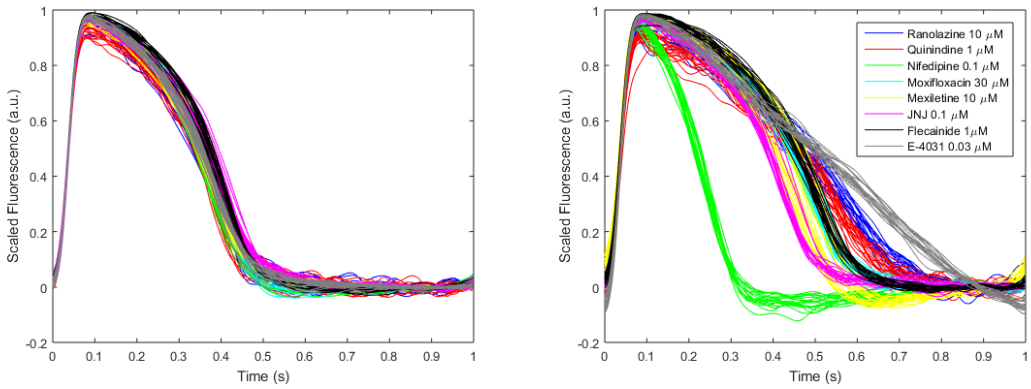

Figure 4.10: Cor.4U cardiomyocytes before (left) and after (right) prescribed drug treatment before transport.

Even without the ladder, there are some clear differences between some of the treatments. For example, it is clear to see that the Nifedipine treatment produces a

radically different response than the others, as it is the only treatment that reduces the action potential duration. The separation between the other treatments is less clear. The percent change in APD features in the remaining treatments (Figure 4.11) suggests at least 2, and perhaps as many as 4, distinct effects. The hope is that the metamorphosis features will be better equipped to discriminate between the drugs because of the additional information being encoded. However, it is hard to make appropriate comparisons before the ladder, as the features are not aligned due to the variability in the starting population.

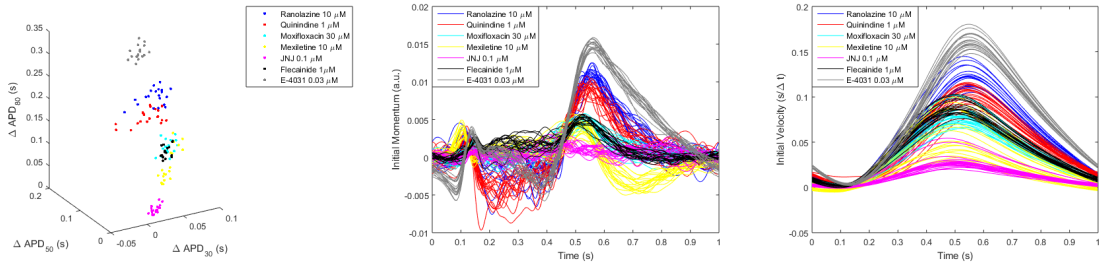

Figure 4.11: Change in APD_{30}, APD_{50} and APD_{80} (left), Metamorphosis Momentum (center), and Metamorphosis Velocity (right) as a result going from the pre-treatment to post-treatment state for the individual treatments on the Cor.4U Cardiomyocytes before parallel transport

To reduce this variability, we construct a ladder from each of the individuals to the mean of the population. While the effect on the action potential features is not clearly pronounced (Figure 4.12), the effect of transport to align the metamorphosis features is more significant, as now the velocity fields are coherent with each other. To see whether this effect is significant, we look at the clustering fitness of the drugs by their treatment before and after transport. Spectral clustering is performed with

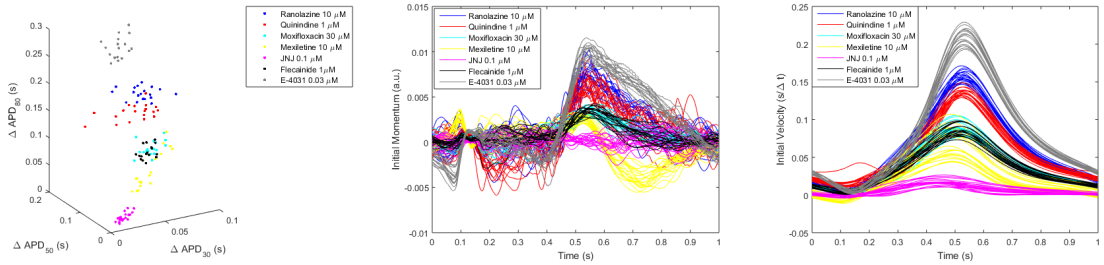

Figure 4.12: Change in APD_{30}, APD_{50} and APD_{80} (left), Metamorphosis Momentum (center), and Metamorphosis Velocity (right) as a result going from the pre-treatment to post-treatment state for the individual treatments on the Cor.4U Cardiomyocytes after parallel transport

a σ equal to the average distance of features between cells of the same group. Before the ladder, the APD features appear to separate the group the best, but once the features being compared are coherent with each other, the metamorphosis velocity has additional information to better separate the cardiomyocyte population (Table 4.2).

Table 4.2: Spectral clustering costs of separation schemes of the treatments on the Cor.4U cardiomyocyte dataset

Clustering Groups	Feature Set	Pre-Ladder Cost	Post-Ladder Cost
Each Treatment Separately	APD	0.4189	0.4353
	Momentum	0.5090	0.6108
	Velocity	0.5642	0.3784
Combining Moxifloxacin and Flecainide	APD	0.3219	0.3406
	Momentum	0.4316	0.5454
	Velocity	0.4949	0.2735

In particular, we want to note two things. First, Moxifloxacin and Flecainide appear to have the same action on the cell population. Regardless of feature, the two drugs aggregate together. To verify this, we compare the cut cost of considering them as one group instead of two separate groups. The cut cost (CC, (2.14)) reduces for

all of the groups by combining these treatments, suggesting that considering them one treatment is a better fit on this population than considering them separate. This result is interesting because the two drugs are not usually implicated in targeting the same channels, yet their multiple channel effects generate the same change in the action potential. Second, the effect of Mexiletine does not appear distinctly in the APD or metamorphosis velocity feature, but is more pronounced in the momentum feature. We verify this by looking at the cut cost of separating Mexiletine from the rest of the population using all three sets of features in Figure 4.13. The reduced cut cost further indicates that Mexiletine's effect on the action potential is best characterized by the momentum feature. This result, and how it manifests in the momentum feature makes sense given its function as a primary sodium channel blocker. It generally does not have an effect on APD, but the block in the influx of Na^+ makes it more difficult to balance the action of the potassium channels, leading to the increased APD and the temporary state of hyper-repolarization at the end of the action potential. Since this is a change in the lower bound of the action potential, and thus a non-diffeomorphic change in the action potential, it makes sense that this is captured in the momentum feature. This suggests that both metamorphosis features may be discriminative in different ways and certainly provides more information than the current standard of features.

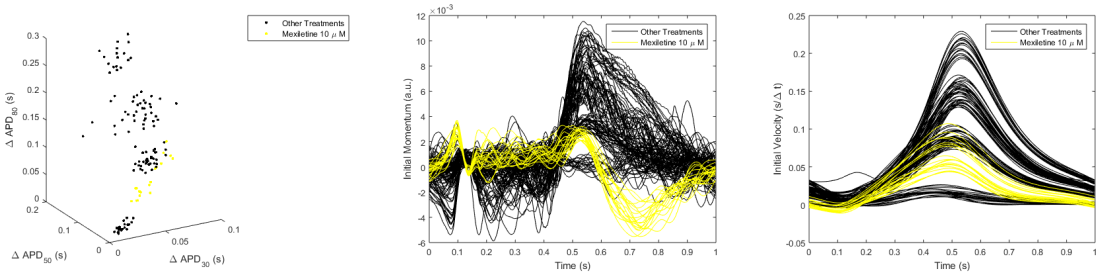

(a) Δ APD: CC = 0.3828 (b) Momentum: CC = 0.2461 (c) Velocity: CC = 0.3132

Figure 4.13: Mexiletine Separation Comparison

4.2.3 Cell Line Comparison of a Single Drug

In this test, we isolate the behavior of one drug across populations of cardiomyocytes derived from two different cell lines. The goal is to determine whether the cardiomyocytes from these two different isolation procedures react the same way to a treatment or where they differ if they do differ. We obtained a population of cardiomyocytes from a monolayer of iCell Cardiomyocytes (Cellular Dynamics International, Madison, WI, USA) and a monolayer of Cor.4U Cardiomyocytes (Axiogenesis, Cologne, Germany). Action potentials are recorded before and after the addition of $0.01\mu M$ E-4031 via optical mapping. Sample action potentials pre- and post-drug traces are shown in Figure 4.14.

First we look at the difference in APD_{30}, APD_{50}, APD_{80} before the ladder. A plot of the raw changes are shown in Figure 4.15. The plots suggest that there is separation between the cell lines in response to E-4031. To verify this, we perform two sample student's t-tests between the distributions of each of these features. The distributions of the individual features and the corresponding p-value of the t-tests are shown in

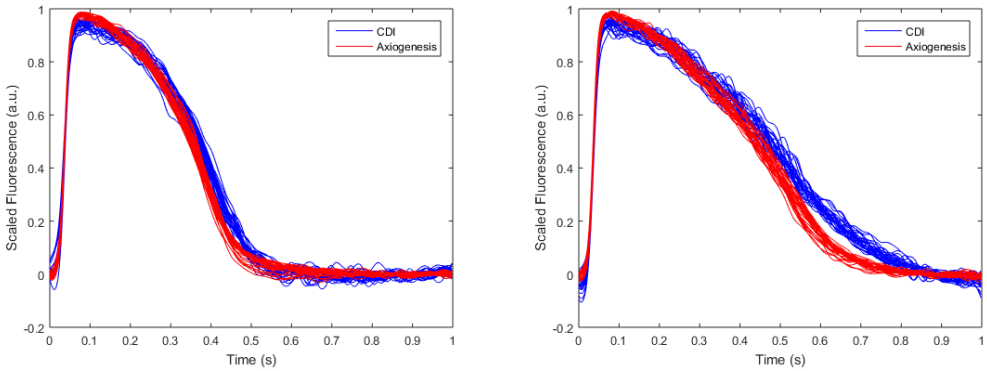

(a) Before Transport

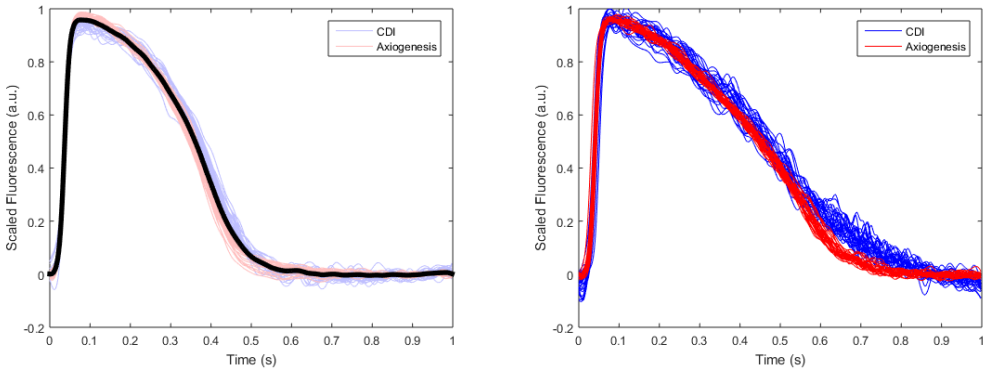

(b) After Transport

Figure 4.14: Action potential recordings from iCell cardiomyocytes (blue) and Cor.4U cardiomyocytes (red) before (left) and after (right) the addition of $0.01\mu M$ E-4031. The top row shows the recordings from the population before parallel transport. Parallel transport is performed to carry the changes to the Fréchet mean of the data (black, bottom row left), and the transported targets are shown on the bottom right

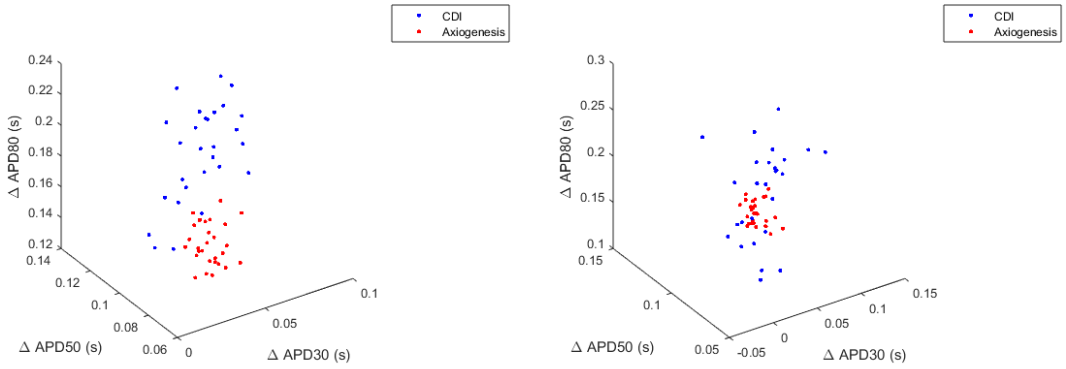

Figure 4.15: Plots of the change in APD_{30}, APD_{50}, and APD_{80} after the addition of E4031 to the CDI and Axiogenesis cardiomyocytes. Before parallel transport (left), the effect of the drug on APD appears to be different between the populations, but after parallel transport (right), there is more overlap in the effect of the drug on APD.

Figure 4.16. The p-values suggest that the two populations of cardiomyocytes are indeed different. But as discussed before, this is misleading because the differences in these features are heavily influenced by the individual members of the population generating these differences. The trajectories are then parallel transported to the mean of the population, and the features are re-evaluated. After parallel transport, the distributions of the APD features are much closer together, so much so that the t-tests for APD_{30} and APD_{50} no longer support the hypothesis that the distributions have different means (Figure 4.17). The only statistically significant separation occurs at APD_{80}, suggesting that the only difference between the cell lines occurs late in the repolarization phase.

We also investigate the metamorphosis features to see if we can be more exact about the differences between the two cell lines. The average action potential target, initial velocity field, and initial momentum of each cell line after parallel transport

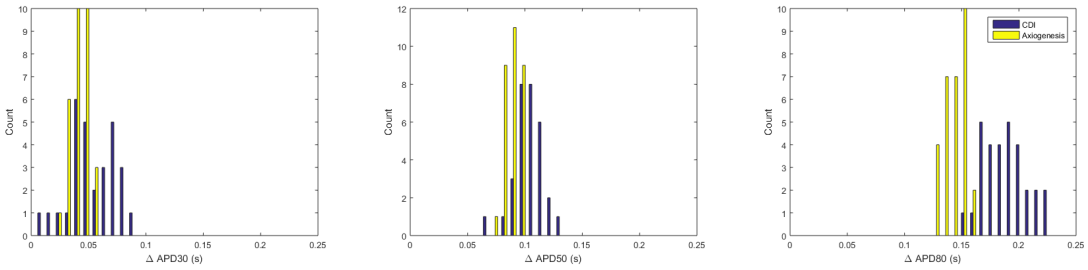

(a) ΔAPD_{30}: $p = 0.0277$ (b) ΔAPD_{50}: $p < 10^{-6}$ (c) ΔAPD_{80}: $p < 10^{-11}$

Figure 4.16: Histograms of change in APD before transport

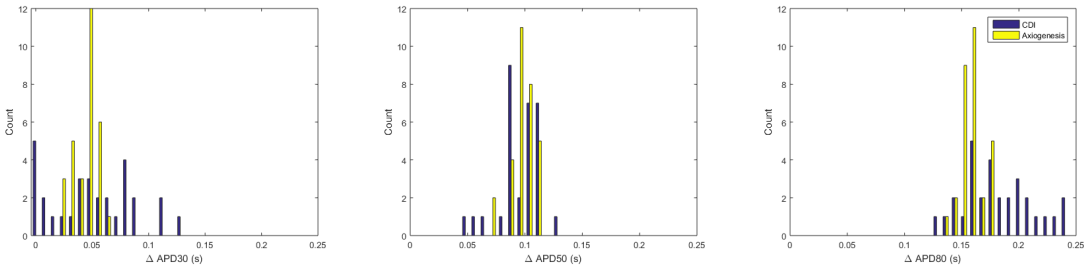

(a) ΔAPD_{30}: $p = 0.4024$ (b) ΔAPD_{50}: $p = 0.5675$ (c) ΔAPD_{80}: $p = 0.00011$

Figure 4.17: Histograms of change in APD after transport

is presented in Figure 4.18. In order to identify the locations where the two cell lines differ after transport, Two sample t-tests were performed at each timepoint of these features to identify where in the features the two cell lines differ. The areas coded in green indicate where the t-test discerns no difference ($p < 0.05$) between the populations. The result of the t-test on the action potential targets coincide with the APD result at around the 30% repolarization, the noisiness of the CDI data makes it difficult to corroborate the APD_{50} result. The momentum is coherent in the same way, but the velocity appears to encode more differences between the two datasets, and it appears that the CDI cardiomyocytes exhibit statistically more deformation

with respect to E4031 than the Axiogenesis cardiomyocytes.

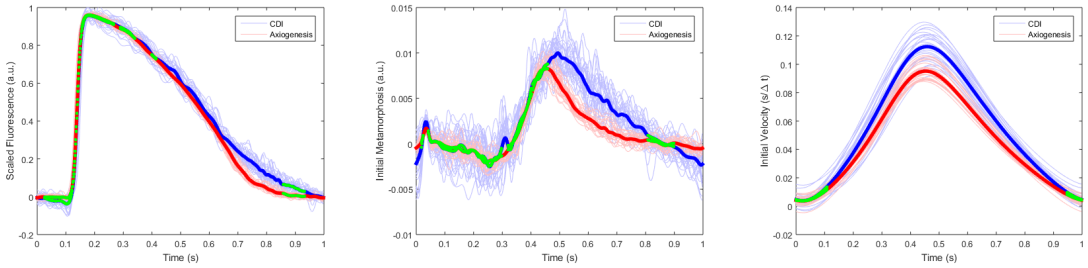

(a) Targets (b) Momenta (c) Velocity

Figure 4.18: Plots of the targets (left), and the metamorphosis momenta (center) and velocity fields (right) connecting the atlas to the targets after parallel transport. The populations specific means for each of the plots is overlaid on top in darker colors. In order to identify locations of statistically significant differences in these plots, two sample t-tests are performed at each timepoint in the plots. Timepoints with no statistical difference ($p < 0.05$) are highlighted in green.

4.2.4 Phenotype Classification Experiment

In this final experiment, we pilot a procedure to attempt to identify the phenotype of an embryonic action potential by its reaction to channel block. The hypothesis is that an embryonic cell will be affected by a change in the environment more similarly to the adult phenotype the cell will become than other phenotypes. To test this, embryonic ventricular-like and embryonic atrial-like action potentials were generated via the hiPSC-CM model of [48]. Adult atrial action potentials from the Nygren [41], Courtemanche [42], and Grandi Left and Right Atrial [43] models, and adult ventricular action potentials from the O'Hara endocardium and epicardium [45] and Ten Tusscher endocardium and epicardium [44] models were also utilized. Action

potentials were recorded from each of the models at normal parameters and after a 50% block in the conductance of the K_r channel. The population of mature and embryonic action potentials is shown in Figure 4.19. Unlike the other experiments, we do not smooth the signals with a Gaussian, as smoothing would lessen the effect of the much more predominant phase 1 notch out of the atrial signals. Instead, we mirror the signal, and combine the signals at the top of the upstroke to observe a more smoothly varying signal. This discards the information at phase 0, which we are willing to accept, as the fast acting potassium rectifier is not typically implicated during the rapid depolarization of the cardiomyocyte.

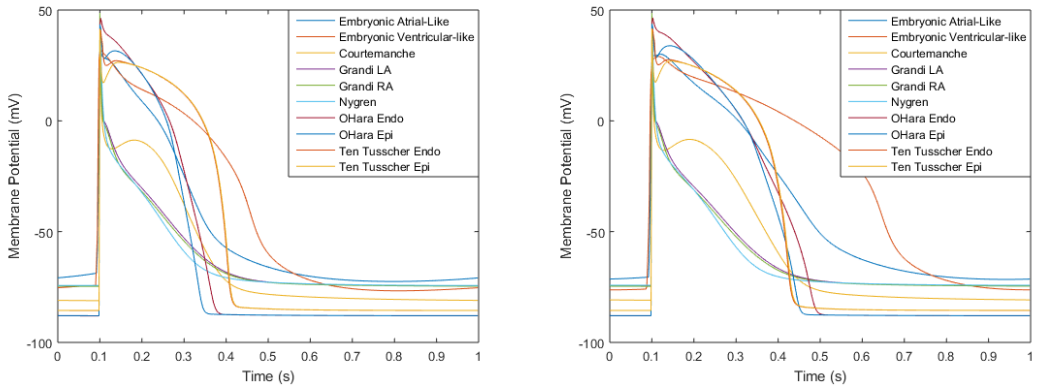

Figure 4.19: Computational Model AP Recordings before (left) and after (right) blocking conductance in K_r channel

Given the different action potential morphologies it is does not make sense to compare the features from the previously discussed experiments before transport. So in order to test the hypothesis, we aim to transport the metamorphosis trajectories of the effects of the channel block from adult models to the embryonic models. Unfortunately, it appears we have approached a limitation in the presented algorithm, as the

transport process appears to fail at different levels for each of the trajectories being transported. Some examples of the failure modes are presented in Figure 4.20. The potentially un-smooth changes in morphology being introduced by the momentum combined with the larger deformation required to arrive at the embryonic action potentials cause the transport process to break down, either during the ladder portion or the final extension.

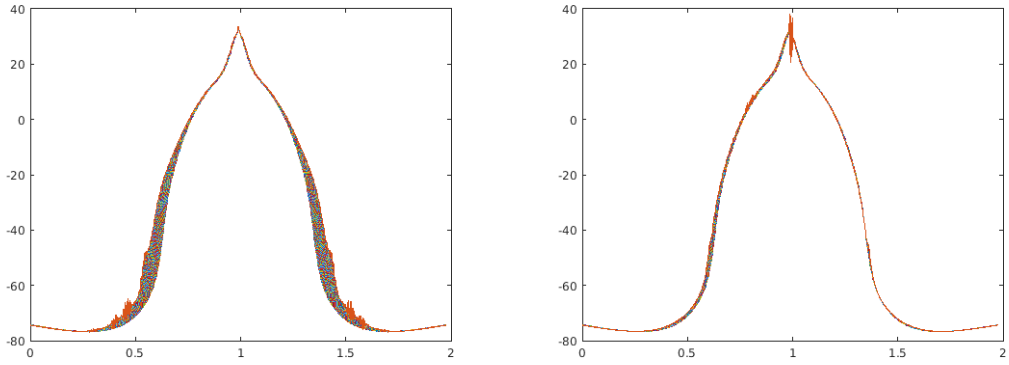

Figure 4.20: Examples of failures when trying to perform parallel transport to the embryonic ventricular-like model. Either the extensions are large and cause instabilities (left) or they are small but located at highly sensitive regions of the action potential like at the peak (right)

Chapter 5

Conclusion

In this thesis, I have presented a shape theoretic model and various machine learning methods to investigate embryonic and mature cardiomyocytes based on action potential morphology. These methods can be used to determine heterogeneity of cardiomyocytes in populations, providing estimates of the eventual chamber specific fate of an embryonic cardiomyocyte, and comparisons of multiple treatments or multiple cardiomyocyte preparations. As the methodology for obtaining cardiomyocytes and recording their action potential with high fidelity and signal-to-noise ratio continues to evolve, the work presented in this thesis will prove ever more important in providing high-throughput analysis of populations of cardiomyocytes.

5.1 Future Directions

While this thesis introduces multiple methods to analyze action potentials for various applications, there are still plenty of unexplored extensions of these methods that are potential future directions of this project. In particular, while I have presented two different algorithms for finding the metamorphosis between two action potentials in chapter 3, improving the computational efficiency of these algorithms could be vital in higher throughput applications. One future direction is to identify a closed form solution for the velocity fields (v) with a fixed metamorphosis family (f). Another potential direction, as discussed in 3.3.3, learning the relationship between incremental changes in v and the resulting interpolation matrix (N_v) and aggregates ($R_{i,j}, A_{l,m}$) could prove particularly helpful increasing the computational efficiency of the metamorphosis optimization, as having to recompute each of these matrices at every iteration when the change in the velocity should get smaller as we get closer to the optimum seems unnecessarily computationally burdensome. To see how one could potentially do this, let us look at N_v.

$$(N_v)_{ij} = \begin{cases} \frac{\tau_i+v}{\Delta\tau} - (j-1) & j-1 <= \frac{\tau_i+v}{\Delta\tau} < j \\ j+1-\frac{\tau_i+v}{\Delta\tau} & j <= \frac{\tau_i+v}{\Delta\tau} < j+1 \\ 0 & \text{otherwise} \end{cases} . \tag{5.1}$$

If we now modulate v by a small change δv:

$$(N_{v+\delta v})_{ij} = \begin{cases} \frac{\tau_i+v+\delta v}{\Delta\tau} - (j-1) & j-1 <= \frac{\tau_i+v+\delta v}{\Delta\tau} < j \\ j+1-\frac{\tau_i+v+\delta v}{\Delta\tau} & j <= \frac{\tau_i+v+\delta v}{\Delta\tau} < j+1 \\ 0 & \text{otherwise} \end{cases}$$

$$= \begin{cases} (N_v)_{ij} + \frac{\delta v}{\Delta\tau} & j-1 <= \frac{\tau_i+v}{\Delta\tau}, \frac{\tau_i+v+\delta v}{\Delta\tau} < j \\ (N_v)_{ij} - \frac{\delta v}{\Delta\tau} & j <= \frac{\tau_i+v}{\Delta\tau}, \frac{\tau_i+v+\delta v}{\Delta\tau} < j+1 \\ (N_v)_{ij} = 0 & j-1 > \frac{\tau_i+v}{\Delta\tau}, \frac{\tau_i+v+\delta v}{\Delta\tau} \text{ or } \frac{\tau_i+v}{\Delta\tau}, \frac{\tau_i+v+\delta v}{\Delta\tau} >= j+1 \end{cases}$$

The open question is how do we define this relationship for the other cases, when the change in v takes us between segments of the piecewise function, and subsequent linear operators such as $R_{i,j}, A_{l,m}$. Computing these incremental changes would greatly cut down the amount computations required to calculate a metamorphosis, speeding up the optimization.

Additionally, while the metamorphosis algorithm is quantitatively coherent with those obtained by biologists on a small dataset, and qualitatively coherent on a large dataset, there is still a gap in systematically showing quantitative coherence between metamorphosis and the established biology at a large scale. Thus a future direction is to obtain a large dataset annotated either manually or via a phenotype specific biomarker such as myosin heavy chain for ventricular cells [11]. Such a dataset could be utilized to validate the metamorphosis metric as a viable predictor of eventual em-

bryonic cardiomyocyte phenotype. Transferability across various validated datasets could prove integral to the metamorphosis algorithm's general acceptance and applicability.

One major limitation of the clustering algorithm presented in Chapter 2 is that the output is, by problem definition, a partition of the data. This presents issues when the data does not necessarily like in clusters, but along a continuum. Clustering algorithms typically do not have criterion to conclude whether the data lies in one large continuum vs multiple clusters. There are heuristic approaches to provide insight into this question. For example, one can observe the projected data points from the spectral clustering algorithm before k-means, and such a projection is provided for the optical mapping dataset in Figure 5.1. But this projection discards some of the higher dimensional information of the raw AP signal that the result may not be completely accurate. Finding such criterion is important to the cardiomyocyte problem because cardiomyocytes of all chamber phenotypes derive from the pluripotent stem cells, so identifying when in the maturation process they differentiate from a continuous population of embryonic heart muscle cells into chamber specific cardiomyocytes is of particular importance for developing tissue models for scientific and therapeutic applications. A future direction of this task is modify the objective to allow for the potential for a continuum. One way of doing this is by introducing a prior on the number of clusters in the dataset, and then performing clustering based on the fixed number of clusters and identifying the likelihood of the data given the clusters and

the prior on the number of clusters.

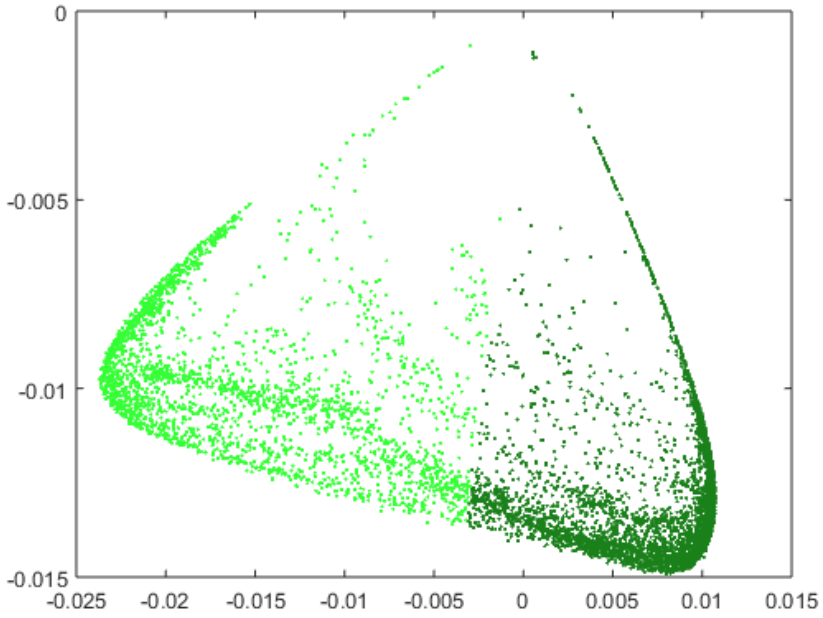

Figure 5.1: Data projection of optical mapping dataset from section 2.4, with the clustering labels from the two group spectral clustering algorithm

Perhaps the biggest room for improvement is in the parallel transport process. As was shown in Chapter 4, there is still a problem with utilizing the ladder as a tool for embryonic classification. One of the suspected reasons for this is the large discontinuity at the upstroke. As discussed previously, phase 0 depolarization is a very integral part of the action potential morphology, and so we would like to preserve the information there as much as possible. While smoothing allows us to perform parallel transport with reduced risks of instabilities, it is not a panacea.

There are a couple of different approaches to addressing this concern in the action potential signals. One approach is to separate the upstroke from the rest of the

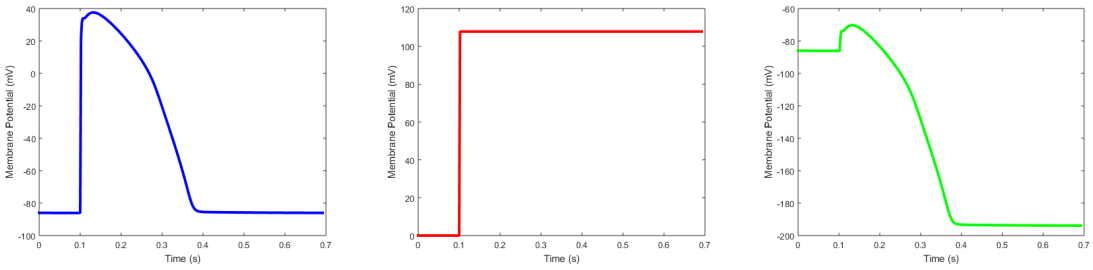

(a) $f(\tau)$ (b) $u_{\tau_0}(\tau)$ (c) $\hat{f}(\tau)$

Figure 5.2: Action potential decomposition with a step function

action potential by trying to model it via a step function. To observe this visually, an example action potential $f(\tau)$ and its decomposition into a step function $u_{\tau_0}(\tau)$, i.e., $u_{\tau_0} = 1$ when $\tau > \tau_0$ and 0 otherwise, and the smoothly decreasing signal $\hat{f}(\tau)$ is shown in Figure 5.2. In order to see how this affects the metamorphosis, let us start with the discrete metamorphosis interpolation from equation (3.24).

$$d^2_{\mathcal{M}}(f_0, f_1) \approx \min_{\substack{v,f \\ f(\tau,0)=f_0(\tau) \\ f(\tau,1)=f_1(\tau)}} \sum_{s=0}^{S-1} (\|v(\tau, t_s)\|^2_V + \frac{1}{\sigma^2}\|f(\tau + v(\tau, t_s), t_{s+1}) - f(\tau, t_s)\|^2_{l^2}). \quad (5.2)$$

Now allowing $f(\tau, \cdot) = c(\cdot)u_{\tau_0(\cdot)}(\tau) + \hat{f}(\tau, \cdot)$:

$$d^2_{\mathcal{M}}(f_0, f_1) \approx \min_{\substack{v,\hat{f},c,\tau_0 \\ f(\tau,t_s)=c(t_s)u_{\tau_0(t_s)}(\tau)+\hat{f}(\tau) \\ f(\tau,0)=f_0(\tau) \\ f(\tau,1)=f_1(\tau)}} \sum_{s=0}^{S-1} (\|v(\tau, t_s)\|^2_V \quad (5.3)$$

$$+ \frac{1}{\sigma^2}\|\hat{f}(\tau + v(\tau, t_s), t_{s+1}) - \hat{f}(\tau, t_s) + c(t_{s+1})u_{\tau_0(t_{s+1})}(\tau + v(\tau, t_s)) - c(t_s)u_{\tau_0(t_s)}(\tau)\|^2_{l^2}),$$

$$d^2_{\mathcal{M}}(f_0, f_1) \approx \min_{\substack{v,\hat{f},c,\tau_0 \\ f(\tau,t_s)=c(t_s)u_{\tau_0(t_s)}(\tau)+\hat{f}(\tau) \\ f(\tau,0)=f_0(\tau) \\ f(\tau,1)=f_1(\tau)}} \sum_{s=0}^{S-1} (\|v(\tau,t_s)\|^2_V \tag{5.4}$$

$$+ \frac{1}{\sigma^2}\|\hat{f}(\tau + v(\tau,t_s), t_{s+1}) - \hat{f}(\tau,t_s) + c(t_{s+1})u_{\tau_0(t_{s+1}-v(\tau,t_s))}(\tau) - c(t_s)u_{\tau_0(t_s)}(\tau)\|^2_{l^2}). \tag{5.5}$$

Under this construction, the remaining signal $\hat{f}$ undergoes metamorphosis as we have previously defined, while u remains a step function and moves with the velocity field of the metamorphosis on $\hat{f}$. By modeling the action potential in this way, we allow the metamorphosis to act on a much smoother signal, reducing the risk of instabilities during the transport process. Identifying how to minimize this energy, whether it can be manipulated into a form amenable to the Schild's ladder process, and its continuous analog are all interesting scientific questions worth investigating.

At the heart of the previously discussed approach is that there are deformation dynamics occurring at different scales throughout action potential; dynamics around the upstroke occur over the scale of milliseconds while the morphology changes in the action potential duration occur at the scale of tens to hundreds of milliseconds. As such, constraining the Sobolev norm to one scale of smoothness may not be appropriate. The step function addresses this by accounting for most of the activity of phase 0, leaving only the APD changes. Instead of splitting the action potential in this

way, another approach is to try to the perform the metamorphosis in a coarse-to-fine manner. In this multi-scale approach, a metamorphosis at the tens to hundreds of milliseconds scale, to attempt to best align the larger morphology changes, then gradually reducing the smoothness constraint over the signal to allow the deformation to adjust for fine changes at the smaller scales, such as at depolarization. The primary concern with this approach is marrying different scales together, as each choice of α induces a new Sobolev norm, which induces a new metamorphosis metric, making for a more complex minimization.

Improving the methods for obtaining and transporting metamorphosis trajectories is just one part of the picture. There is still a lot that can be expanded upon in the analysis of these trajectories. One such problem that we only briefly touched upon in Chapter 4 is the question of similarity between drug actions, where the same general trajectory may occur at different dosages for two drugs. A similar phenomenon can be found in the activity recognition literature, when an action, such as walking, occurs at different rates due to gait, body type, etc. Similar to the effect of drugs on action potential morphology, actions also traverse Riemannian trajectories. Thus we should be able to transfer some of the existing techniques for classifying actions to the space of action potentials. One potential avenue worth investigating is based on the work of [72], where actions are not only normalized to remove temporal variability, but also use common dimensionality reduction techniques on the tangent space to help visualize the trajectories and improve classification. They not only

showed state-of-the-art classification results on activity classification datasets, but also expanded their scope to include stroke rehabilitation classification and action reconstruction. Pursuing an analog for the action potential drug data to improve drug action clustering and classification should prove increasingly insightful to clinical cardiology and anti-arrhythmic studies.

Bibliography

[1] I. Kehat, D. Kenyagin-Karsenti, M. Snir, H. Segev, M. Amit, A. Gepstein, E. Livne, O. Binah, J. Itskovitz-Eldor, and L. Gepstein, "Human embryonic stem cells can differentiate into myocytes with structural and functional properties of cardiomyocytes," *Journal of Clinical Investigation*, vol. 108, no. 3, pp. 407–414, 2001.

[2] M. A. Laflamme, K. Y. Chen, A. V. Naumova, V. Muskheli, J. A. Fugate, S. K. Dupras, H. Reinecke, C. Xu, M. Hassanipour, S. Police, C. O'Sullivan, L. Collins, Y. Chen, E. Minami, E. A. Gill, S. Ueno, C. Yuan, J. Gold, and C. E. Murry, "Cardiomyocytes derived from human embryonic stem cells in pro-survival factors enhance function of infarcted rat hearts." *Nature Biotech.*, vol. 25, no. 9, pp. 1015–24, 2007.

[3] N. M. Mordwinkin, P. W. Burridge, and J. C. Wu, "A review of human pluripotent stem cell-derived cardiomyocytes for high-throughput drug discovery, cardiotoxicity screening, and publication standards." *Journal of Cardiovascular*

Translational Research, vol. 6, no. 1, pp. 22–30, Nov. 2012.

[4] K. Blinova, J. Stohlman, J. Vicente, D. Chan, L. Johannesen, M. P. Hortigon-Vinagre, V. Zamora, G. Smith, W. J. Crumb, L. Pang, B. Lyn-Cook, J. Ross, M. Brock, S. Chvatal, D. Millard, L. Galeotti, N. Stockbridge, and D. G. Strauss, "Comprehensive translational assessment of human-induced pluripotent stem cell derived cardiomyocytes for evaluating drug-induced arrhythmias," *Toxicological Sciences*, vol. 155, no. 1, pp. 234–247, 2017.

[5] C. L. Mummery, J. Zhang, E. S. Ng, D. a. Elliott, A. G. Elefanty, and T. J. Kamp, "Differentiation of human embryonic stem cells and induced pluripotent stem cells to cardiomyocytes: a methods overview." *Circulation Research*, vol. 111, no. 3, pp. 344–58, Jul. 2012.

[6] W.-Z. Zhu, Y. Xie, K. W. Moyes, J. D. Gold, B. Askari, and M. A. Laflamme, "Neuregulin/ErbB signaling regulates cardiac subtype specification in differentiating human embryonic stem cells." *Circulation Research*, vol. 107, no. 6, pp. 776–86, Sep. 2010.

[7] M. Gaur, C. Ritner, and R. Sievers, "Timed inhibition of p38MAPK directs accelerated differentiation of human embryonic stem cells into cardiomyocytes," *Cytotherapy*, vol. 12, no. 6, pp. 807–817, 2010.

[8] M. Habib, O. Caspi, and L. Gepstein, "Human embryonic stem cells for car-

diomyogenesis." *Journal of Molecular and Cellular Cardiology*, vol. 45, no. 4, pp. 462–74, 2008.

[9] X. Zhang, J.-P. Guo, Y.-L. Chi, Y.-C. Liu, C.-S. Zhang, X.-Q. Yang, H.-Y. Lin, E.-P. Jiang, S.-H. Xiong, Z.-Y. Zhang, and B.-H. Liu, "Endothelin-induced differentiation of Nkx2.5 cardiac progenitor cells into pacemaking cells." *Molecular and Cellular Biochemistry*, vol. 366, no. 1-2, pp. 309–18, 2012.

[10] J.-D. Fu, S. Rushing, D. Lieu, C. Chan, C.-W. Kong, L. Geng, K. Wilson, N. Chiamvimonvat, K. Boheler, J. Wu, G. Keller, R. Hajjar, and R. Li, "Distinct roles of microRNA-1 and -499 in ventricular specification and functional maturation of human embryonic stem cell-derived cardiomyocytes." *PloS One*, vol. 6, no. 11, 2011.

[11] J. Zhang, G. F. Wilson, A. G. Soerens, C. H. Koonce, J. Yu, S. P. Palecek, J. A. Thomson, and T. J. Kamp, "Functional cardiomyocytes derived from human induced pluripotent stem cells." *Circulation Research*, vol. 104, no. 4, pp. e30–41, Feb. 2009.

[12] A. G. Kléber and Y. Rudy, "Basic mechanisms of cardiac impulse propagation and associated arrhythmias," *Physiological reviews*, vol. 84, no. 2, pp. 431–488, 2004.

[13] S. Nattel, A. Maguy, S. Le Bouter, and Y.-H. Yeh, "Arrhythmogenic ion-channel

remodeling in the heart: heart failure, myocardial infarction, and atrial fibrillation," *Physiological reviews*, vol. 87, no. 2, pp. 425–456, 2007.

[14] M. Saad, A. Mahmoud, I. Y. Elgendy, and C. Richard Conti, "Ranolazine in cardiac arrhythmia," *Clinical Cardiology*, vol. 39, no. 3, pp. 170–178, 2016.

[15] V. A. Maltsev, J. Rohwedel, J. Hescheler, and A. M. Wobus, "Embryonic stem cells differentiate in vitro into cardiomyocytes representing sinusnodal, atrial and ventricular cell types," *Mechanisms of development*, vol. 44, no. 1, pp. 41–50, 1993.

[16] J.-Q. He, Y. Ma, Y. Lee, J. A. Thomson, and T. J. Kamp, "Human embryonic stem cells develop into multiple types of cardiac myocytes: action potential characterization." *Circulation Research*, vol. 93, no. 1, pp. 32–9, 2003.

[17] S. Weinberg, E. A. Lipke, and L. Tung, "In Vitro Electrophysiological Mapping of Stem Cells," *Methods in Molecular Biology*, vol. 660, pp. 215–237, 2010.

[18] A. Petrie and X. Zhao, "Estimating eigenvalues of dynamical systems from time series with applications to predicting cardiac alternans," *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 2012.

[19] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2006.

[20] A. C. Tan, D. Q. Naiman, L. Xu, R. L. Winslow, and D. Geman, "Simple decision

rules for classifying human cancers from gene expression profiles," *Bioinformatics*, vol. 21, no. 20, pp. 3896–3904, 2005.

[21] D. L. Pham, C. Xu, and J. L. Prince, "Current methods in medical image segmentation," *Annual review of biomedical engineering*, vol. 2, no. 1, pp. 315–337, 2000.

[22] A. Dasgupta, Y. V. Sun, I. R. Konig, J. E. Bailey-Wilson, and J. D. Malley, "Brief review of regression-based and machine learning methods in genetic epidemiology: the Genetic Analysis Workshop 17 experience." *Genetic Epidemiology*, vol. 35 Suppl 1, pp. S5–11, 2011.

[23] L. Deng and X. Li, "Machine learning paradigms for speech recognition: An overview," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 21, no. 5, pp. 1060–1089, 2013.

[24] S. Lloyd, "Least squares quantization in PCM," *Technical Report, Bell Laboratories, Published in 1982 in IEEE Trans. Inf. Theory 28: 128-137*, 1957.

[25] J. Shi and J. Malik, "Normalized cuts and image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, pp. 888–905, 2000.

[26] A. Ng, Y. Weiss, and M. Jordan, "On spectral clustering: analysis and an algorithm," in *Neural Information Processing Systems*, 2001, pp. 849–856.

[27] U. von Luxburg, "A tutorial on spectral clustering," *Statistics and Computing*, vol. 17, no. 4, pp. 395–416, 2007.

[28] N. Ahmad, D. Alahakoon, and R. Chau, "Cluster identification and separation in the growing self-organizing map: application in protein sequence classification," *Neural Computing & Applications*, pp. 531–542, 2010.

[29] D. Davies and D. Bouldin, "A cluster separation measure," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 1, no. 2, pp. 224–227, 1979.

[30] H. Saitoh, J. C. Bailey, and B. Surawicz, "Alternans of action potential duration after abrupt shortening of cycle length: differences between dog purkinje and ventricular muscle fibers." *Circulation Research*, vol. 62, no. 5, pp. 1027–1040, 1988.

[31] J. C. Moore, J. Fu, Y.-C. Chan, D. Lin, H. Tran, H.-F. Tse, and R. A. Li, "Distinct cardiogenic preferences of two human embryonic stem cell (hESC) lines are imprinted in their proteomes in the pluripotent state." *Biochemical and biophysical research communications*, vol. 372, no. 4, pp. 553–8, 2008.

[32] S. Peng, A. E. Lacerda, G. E. Kirsch, A. M. Brown, and A. Bruening-Wright, "The action potential and comparative pharmacology of stem cell-derived human cardiomyocytes." *Journal of Pharmacological and Toxicological Methods*, vol. 61, no. 3, pp. 277–86, 2010.

[33] A. Blazeski, R. Zhu, D. W. Hunter, S. H. Weinberg, K. R. Boheler, E. T. Zambidis, and L. Tung, "Electrophysiological and contractile function of cardiomyocytes derived from human embryonic stem cells," *Progress in Biophysics and Mol. Bio.*, vol. 110, no. 2-3, pp. 178–195, 2012.

[34] Q. Zhang, J. Jiang, P. Han, Q. Yuan, J. Zhang, X. Zhang, Y. Xu, H. Cao, Q. Meng, L. Chen, T. Tian, X. Wang, P. Li, J. Hescheler, G. Ji, and Y. Ma, "Direct differentiation of atrial and ventricular myocytes from human embryonic stem cells by alternating retinoid signals." *Cell Research*, vol. 21, no. 4, pp. 579–87, 2011.

[35] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, Sep 1995.

[36] N. Cristianini and J. Shawe-Taylor, *An Introduction to Support Vector Machines*. Cambridge, UK: Cambridge University Press, 2000.

[37] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.

[38] M. Minsky and S. Papert, "Perceptrons," 1969.

[39] C. M. Bishop, *Neural networks for pattern recognition*. Oxford University press, 1995.

[40] A. L. Hodgkin and A. F. Huxley, "A quantitative description of membrane cur-

rent and its application to conduction and excitation in nerve," *The Journal of physiology*, vol. 117, no. 4, pp. 500–544, 1952.

[41] A. Nygren, C. Fiset, L. Firek, J. Clark, D. Lindblad, R. Clark, and W. Giles, "Mathematical model of an adult human atrial cell: The role of K+ currents in repolarization," *Circulation Research*, vol. 82, no. 1, pp. 63–81, 1998.

[42] M. Courtemanche, R. J. Ramirez, and S. Nattel, "Ionic mechanisms underlying human atrial action potential properties: insights from a mathematical model," *American Journal of Physiology - Heart and Circulatory Physiology*, vol. 275, no. 1, pp. H301–H321, 1998.

[43] E. Grandi, S. V. Pandit, N. Voigt, A. J. Workman, D. Dobrev, J. Jalife, and D. M. Bers, "Human atrial action potential and ca2+ model," *Circulation Research*, 2011.

[44] K. H. W. J. ten Tusscher, D. Noble, P. J. Noble, and A. V. Panfilov, "A model for human ventricular tissue," *American Journal of Physiology - Heart and Circulatory Physiology*, vol. 286, no. 4, pp. H1573–H1589, 2004.

[45] T. O'Hara, L. Virág, A. Varró, and Y. Rudy, "Simulation of the undiseased human cardiac ventricular action potential: model formulation and experimental validation." *PLoS computational biology*, vol. 7, no. 5, 2011.

[46] M. Boyett, H. Zhang, A. Garny, and A. Holden, "Control of the pacemaker ac-

tivity of the sinoatrial node by intracellular ca2+. experiments and modelling," *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 359, no. 1783, pp. 1091–1110, 2001.

[47] Y. Kurata, I. Hisatome, S. Imanishi, and T. Shibamoto, "Roles of l-type ca 2+ and delayed-rectifier k+ currents in sinoatrial node pacemaking: insights from stability and bifurcation analyses of a mathematical model," *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 285, no. 6, pp. H2804–H2819, 2003.

[48] M. Paci, J. Hyttinen, K. Aalto-Setälä, and S. Severi, "Computational models of ventricular- and atrial-like human induced pluripotent stem cell derived cardiomyocytes," *Annals of Biomedical Engineering*, vol. 41, no. 11, pp. 2334–2348, Nov 2013.

[49] D. G. Kendall, "Shape manifolds, procrustean metrics, and complex projective spaces," *Bulletin of the London Mathematical Society*, vol. 16, no. 2, pp. 81–121, 1984.

[50] D. W. Thompson, *On Growth and Form*, J. T. Bonner, Ed. Cambridge University Press, 1992.

[51] U. Grenander, *General Pattern Theory: A Mathematical Study of Regular Structures*. New York, NY, USA: Oxford University Press, Inc., 1994.

[52] U. Grenander and M. Miller, *Pattern Theory: From Representation to Inference*. New York, NY, USA: Oxford University Press, Inc., 2007.

[53] M. F. Beg and A. Khan, "Computing an average anatomical atlas using LDDMM and geodesic shooting," in *3rd IEEE International Symposium on Biomedical Imaging: Nano to Macro, 2006.*, pp. 1116–1119.

[54] C. Ceritoglu, X. Tang, M. Chow, D. Hadjiabadi, D. Shah, T. Brown, M. Burhanullah, H. Trinh, J. Hsu, K. Ament, D. Crocetti, S. Mori, S. Mostofsky, S. Yantis, M. Miller, and J. T. Ratnanather, "Computational analysis of LDDMM for brain mapping," *Frontiers in Neuroscience*, vol. 7, p. 151, 2013.

[55] L. Wang, F. Beg, T. Ratnanather, C. Ceritoglu, L. Younes, J. C. Morris, J. G. Csernansky, and M. I. Miller, "Large deformation diffeomorphism and momentum based hippocampal shape discrimination in dementia of the Alzheimer type," *IEEE Transactions on Medical Imaging*, vol. 26, no. 4, pp. 462–470, April 2007.

[56] L. Younes, *Shapes and Diffeomorphisms*, ser. Applied Mathematical Sciences, S. Antman, J. Marsden, and L. Sirovich, Eds. Springer, 2010, vol. 171.

[57] L. Garcin and L. Younes, "Geodesic image matching: A wavelet based energy minimization scheme," *Energy Minimization Methods in Computer Vision and Pattern Recognition*, pp. 349–364, 2005.

[58] A. Trouvé and L. Younes, "Local geometry of deformable templates," *SIAM Journal on Mathematical Analysis*, vol. 37, no. 1, pp. 17–59, 2005.

[59] S. Iravanian and L. Tung, "A novel algorithm for cardiac biosignal filtering based on filtered residue method." *IEEE Transactions on Biomedical Engineering*, vol. 49, no. 11, pp. 1310–7, 2002.

[60] D. T. M. Du, N. Hellen, C. Kane, and C. M. Terracciano, "Action Potential Morphology of Human Induced Pluripotent Stem Cell-Derived Cardiomyocytes Does Not Predict Cardiac Chamber Specificity and Is Dependent on Cell Density," *Biophysical Journal*, vol. 108, no. 1, pp. 1–4, 2015.

[61] E. Wettwer, T. Christ, S. Endig, N. Rozmaritsa, K. Matschke, J. J. Lynch, M. Pourrier, J. K. Gibson, D. Fedida, M. Knaut, and U. Ravens, "The new antiarrhythmic drug vernakalant: ex vivo study of human atrial tissue from sinus rhythm and chronic atrial fibrillation," *Cardiovascular Research*, vol. 98, no. 1, pp. 145–154, 2013.

[62] S. Braam, L. Tertoolen, S. Casini, E. Matsa, H. Lu, A. Teisman, R. Passier, C. Denning, D. Gallacher, R. Towart, and C. Mummery, "Repolarization reserve determines drug responses in human pluripotent stem cell derived cardiomyocytes," *Stem Cell Research*, vol. 10, no. 1, pp. 48 – 56, 2013.

[63] J. K. Gibson, Y. Yue, J. Bronson, C. Palmer, and R. Numann, "Human stem cell-derived cardiomyocytes detect drug-mediated changes in action potentials and

ion currents," *Journal of Pharmacological and Toxicological Methods*, vol. 70, no. 3, pp. 255 – 267, 2014, 11th annual focused issue on methods in safety pharmacology.

[64] C. Heylman, R. Datta, A. Sobrino, S. George, and E. Gratton, "Supervised machine learning for classification of the electrophysiological effects of chronotropic drugs on human induced pluripotent stem cell-derived cardiomyocytes," *PLOS ONE*, vol. 10, no. 12, pp. 1–15, 12 2015.

[65] J. Stoker, *Differential Geometry*. Wiley, 1989.

[66] J. Ehlers, F. A. E. Pirani, and A. Schild, "Republication of: The geometry of free fall and light propagation," *General Relativity and Gravitation*, vol. 44, no. 6, pp. 1587–1609, Jun 2012.

[67] M. Lorenzi, N. Ayache, and X. Pennec, "Schild's ladder for the parallel transport of deformations in time series of images," in *Proceedings of Information Processing in Medical Imaging: 22nd International Conference, IPMI 2011.*, G. Székely and H. K. Hahn, Eds. Springer Berlin Heidelberg, pp. 463–474.

[68] X. Chen, J. D. Cass, J. A. Bradley, C. M. Dahm, Z. Sun, E. Kadyszewski, M. J. Engwall, and J. Zhou, "QT prolongation and proarrhythmia by moxifloxacin: concordance of preclinical models in relation to clinical outcome," *British Journal of Pharmacology*, vol. 146, no. 6, pp. 792–799, 2005.

[69] S. Sossalla, B. Kallmeyer, S. Wagner, M. Mazur, U. Maurer, K. Toischer, J. D. Schmitto, R. Seipelt, F. A. Schndube, G. Hasenfuss, L. Belardinelli, and L. S. Maier, "Altered Na+ currents in atrial fibrillation," *Journal of the American College of Cardiology*, vol. 55, no. 21, pp. 2330 – 2342, 2010.

[70] R. H. Falk, "Flecainide-induced ventricular tachycardia and fibrillation in patients treated for atrial fibrillation," *Annals of Internal Medicine*, vol. 111, no. 2, pp. 107–111, 1989.

[71] J. Morganroth and J. E. Goin, "Quinidine-related mortality in the short-to-medium-term treatment of ventricular arrhythmias. a meta-analysis." *Circulation*, vol. 84, no. 5, pp. 1977–1983, 1991.

[72] R. Anirudh, P. Turaga, J. Su, and A. Srivastava, "Elastic functional coding of riemannian trajectories," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 5, pp. 922–936, 2017.

Vita

Giann Gorospe received his B.S.E in Bioengineering and B.A. in Mathematics from the University of Pennsylvania in spring of 2009, and enrolled in the Biomedical Engineering Ph.D. program at Johns Hopkins University in the following fall. Research interests include machine learning applied to biomedical systems, in particular cardiac electrophysiology, as well as deformable shape analysis. His current research project involves developing automated methods for cardiomyocyte classification and clustering based on properties of the action potential. His work entitled, "Automated Grouping of Action Potentials of Human Embryonic Stem Cell-Derived Cardiomyocytes" was featured on the cover of the September 2014 issue of IEEE Transactions in Biomedical Engineering.