

Machine Learning Assisted Non-Rigid Surface Tracking in Biological Systems

by William Henry Hempstead, M.S.B.E.

July 2020

Directed by: Zhen Zhu, Ph.D.

Department of Engineering

Occlusions, obstructions, and lighting changes that occur in a camera's field-of-view (FOV) during a medical procedure can cause tracking algorithms to lose track of a particular region-of-interest (ROI). Various approaches to reacquire the tracking of rigid objects have been developed, however, the non-rigid nature of biological structures requires more complicated approaches. The purpose of this research is to improve the performance of an existing non-rigid tracking algorithm under these types of adverse conditions. This existing algorithm was previously shown to be accurate and efficient under ideal conditions but exhibited a high rate of tracking failures due to the aforementioned occlusions, obstructions, and lighting changes. To improve tracking under these conditions, a tissue motion machine learning model was developed to provide predictions of future ROI grid motion. The combination of this machine learning technique along with various improvements to the base algorithm was shown to greatly reduce the number of tracking resets and allow the tracking grid to briefly follow an expected motion pattern during a simulated occlusion.

Machine Learning Assisted Non-Rigid Surface Tracking in Biological Systems

A THESIS

Presented to the Faculty of the Department of Engineering

College of Engineering and Technology

East Carolina University

In Partial Fulfillment of the Requirements For the Degree

Master of Science in Biomedical Engineering

by

William Henry Hempstead

July, 2020

©Copyright 2020

William Henry Hempstead

Machine Learning Assisted Non-Rigid Surface Tracking in Biological Systems

By

William Henry Hempstead

APPROVED BY:

DIRECTOR OF
THESIS: _____

Zhen Zhu, Ph.D.

COMMITTEE MEMBER: _____

Sunghan Kim, Ph.D.

COMMITTEE MEMBER: _____

Rui Wu, Ph.D.

GRADUATE PROGRAM DIRECTOR: _____

Sunghan Kim, Ph.D.

CHAIR OF THE
DEPARTMENT OF ENGINEERING: _____

Barbara J. Muller-Borer, Ph.D.

DEAN OF THE
GRADUATE SCHOOL: _____

Paul J. Gemperline, Ph.D.

"Those who have knowledge, don't predict. Those who predict, don't have knowledge."

-Lao Tzu

ACKNOWLEDGMENTS

I would like to give my highest thanks to my thesis advisor, Dr. Zhen Zhu, of the Department of Engineering at East Carolina University for his insight and guidance during the development of this research.

In addition, I would like to thank the other members of my thesis committee, Dr. Sunghan Kim of the Department of Engineering and Dr. Rui Wu of the Department of Computer Science, for their continued support.

Thank you to Dr. Barbara Muller-Borer for your encouragement and support in my return to study in the field of biomedical engineering.

Thanks to Dr. Cheng Chen and Dr. Bruce Ferguson of RFPi, Inc. for providing the data set used in this research.

Finally, thank you to all of my friends, new and old, near and far, who have been supportive of me during the course of this research.

TABLE OF CONTENTS

| | Page |
|---|-------------|
| ACKNOWLEDGMENTS | v |
| LIST OF TABLES | x |
| LIST OF FIGURES | xi |
| | |
| CHAPTER | |
| 1. INTRODUCTION | 1 |
| 1.1 Problem Statement and Proposed Work | 1 |
| 1.2 Organization of Thesis | 3 |
| 2. BACKGROUND INFORMATION | 5 |
| 2.1 Previous Work | 5 |
| 2.2 Time Series Forecasting | 8 |
| 2.2.1 Characteristics of Time Series Data | 8 |
| 2.2.2 Deterministic Models | 10 |
| 2.2.3 Machine Learning Methods | 10 |
| 3. METHODS | 15 |
| 3.1 Programming Framework | 15 |
| 3.2 Software Packages and Development Environment | 16 |
| 3.2.1 MATLAB | 17 |
| 3.2.2 Linux | 17 |
| 3.2.3 Python | 18 |

| | | |
|-----------|--|-----------|
| 3.2.4 | OpenCV | 18 |
| 3.2.5 | TensorFlow | 19 |
| 3.2.6 | Hardware | 19 |
| 3.3 | Improvements \ Changes to the Original Algorithm | 20 |
| 3.3.1 | Frame Pre-processing | 20 |
| 3.3.2 | Feature Detection Algorithm/Parameters | 24 |
| 3.3.3 | Tracking Algorithm/Parameters | 25 |
| 3.3.4 | Initial Frame Selection | 25 |
| 3.3.5 | Pulse Detection | 26 |
| 3.3.6 | Determination of Number of Patches | 27 |
| 3.3.7 | Neighboring Patch Smoothness Constraint | 27 |
| 3.3.8 | Tracking Re-initialization | 29 |
| 3.3.9 | Parameter Values | 30 |
| 3.3.10 | Occlusion Detection | 31 |
| 3.4 | Feature Extraction | 32 |
| 3.4.1 | Data Set | 32 |
| 3.4.2 | Extraction Method / Algorithm | 34 |
| 3.4.3 | Sequence Length | 35 |
| 3.4.4 | Scaling | 36 |
| 3.5 | Neural Network | 36 |
| 3.5.1 | Architecture | 36 |
| 3.5.2 | Training | 38 |
| 3.6 | Testing Scenarios | 39 |
| 4. | RESULTS | 41 |
| 4.1 | Results for 20160725T124323.avi | 43 |
| 4.2 | Results for 20160725T124609.avi | 46 |
| 4.3 | Results for 20160725T124909.avi | 49 |
| 4.4 | Results for 20160725T132840.avi | 52 |
| 4.5 | Results for 20160725T133113.avi | 55 |
| 4.6 | Neural Network Training | 57 |
| 4.7 | Speed of Execution | 59 |

| | |
|---------------------------------------|-----------|
| 5. DISCUSSION | 61 |
| 5.1 Tracking Performance | 61 |
| 5.2 Prediction Performance | 62 |
| 5.3 Speed of Execution | 63 |
| 5.4 Limitations / Future Improvements | 65 |
| 5.4.1 Tracking Algorithm | 65 |
| 5.4.2 Data Collection | 66 |
| 5.4.3 Neural Network | 67 |
| 5.4.4 Speed of Execution | 68 |
| 6. CONCLUSIONS | 70 |
| REFERENCES | 71 |
| APPENDICES | |
| APPENDICES | 77 |
| A. SOURCE CODE | 78 |
| A.1 Main Algorithm | 79 |
| A.1.1 Overview | 80 |
| A.1.2 Code | 81 |
| A.2 Feature Point Sequence Extraction | 92 |
| A.2.1 Overview | 92 |
| A.2.2 Code | 93 |
| A.3 Data Set Creation | 99 |
| A.3.1 Overview | 99 |
| A.3.2 Code | 100 |
| A.4 Neural Network Training | 105 |
| A.4.1 Overview | 105 |

| | | |
|-------|--------------------------------------|-----|
| A.4.2 | Code | 106 |
| A.5 | Supporting Classes / Functions | 109 |
| A.5.1 | Classes | 109 |
| A.5.2 | Functions | 112 |
| A.6 | MAT to AVI Conversion | 151 |

LIST OF TABLES

| Table | Page |
|--|-------------|
| 3.1 Extracted Sequence Example. | 35 |
| 3.2 RNN Model Layers. | 37 |
| 4.1 Tracking Performance Summary (20160725T124323.avi). | 44 |
| 4.2 Tracking Performance Summary (20160725T124609.avi). | 47 |
| 4.3 Tracking Performance Summary (20160725T124909.avi). | 50 |
| 4.4 Tracking Performance Summary (20160725T132840.avi). | 53 |
| 4.5 Tracking Performance Summary (20160725T133113.avi). | 56 |
| 4.6 Execution Speed Summary. | 60 |
| 5.1 Tracking Failure Summary and Comparison. | 61 |
| 5.2 Grid Prediction Error Summary. | 62 |
| 5.3 Profiler Summary for 20160725T124323.avi. | 64 |

LIST OF FIGURES

| Figure | Page |
|--|-------------|
| 1.1 Example of Occlusion from Data Set | 3 |
| 2.1 Example of Region-of-Interest (ROI) | 6 |
| 2.2 Simplified Block Diagram of Original Algorithm | 7 |
| 2.3 Components of a Time Series | 9 |
| 2.4 Multi-Layer Perceptron (MLP) Diagram | 11 |
| 2.5 Recurrent Neural Network (RNN) Diagram | 12 |
| 2.6 Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) Nodes | 13 |
| 2.7 Convolutional Neural Network (CNN) Diagram | 14 |
| 3.1 Block Diagram of Improved Algorithm | 16 |
| 3.2 Comparison of Histogram Equalization Methods | 22 |
| 3.3 Comparison of Denoising Methods | 24 |
| 3.4 Optical Flow from Feature Point Movement | 27 |
| 3.5 Smoothness Criteria Constraint | 28 |
| 3.6 Sample Frame From Data Set | 33 |
| 3.7 Sample of Extracted Feature Points | 34 |

| | | |
|------|--|----|
| 3.8 | ROI Selection Landmarks | 39 |
| 4.1 | Representative Screenshots (20160725T124323.avi)..... | 43 |
| 4.2 | Analysis Timeline (20160725T124323.avi) | 44 |
| 4.3 | Average Grid Prediction Error (x, y, and x-y (absolute)) - (20160725T124323.avi) | 45 |
| 4.4 | Representative Screenshots (20160725T124609.avi)..... | 46 |
| 4.5 | Analysis Timeline (20160725T124609.avi) | 47 |
| 4.6 | Average Grid Prediction Error (x, y, and x-y (absolute)) - (20160725T124609.avi) | 48 |
| 4.7 | Representative Screenshots (20160725T124909.avi)..... | 49 |
| 4.8 | Analysis Timeline (20160725T124909.avi) | 50 |
| 4.9 | Average Grid Prediction Error (x, y, and x-y (absolute)) - (20160725T124909.avi) | 51 |
| 4.10 | Representative Screenshots (20160725T132840.avi) | 52 |
| 4.11 | Analysis Timeline (20160725T132840.avi) | 53 |
| 4.12 | Average Grid Prediction Error (x, y, and x-y (absolute)) - (20160725T132840.avi) | 54 |
| 4.13 | Representative Screenshots (20160725T133113.avi). | 55 |
| 4.14 | Analysis Timeline (20160725T133113.avi) | 56 |
| 4.15 | Average Grid Prediction Error (x, y, and x-y (absolute)) - (20160725T133113.avi) | 57 |
| 4.16 | Neural Network Training Loss | 58 |
| 4.17 | Neural Network Training Accuracy | 59 |

CHAPTER 1

INTRODUCTION

1.1 Problem Statement and Proposed Work

The increase in machine-assisted and robotic-assisted medical procedures over the past several years has driven a need for more robust tracking techniques in computer vision. One of the most well-known robotic surgical systems, the da Vinci Surgical System, has grown to 5,582 total installations since entering the market in 2000 with approximately 1,229,000 procedures performed in 2019 [1]. These robotic, as well as, traditional laparoscopic surgeries are increasingly dependent on computer vision technologies for procedures such as tumor or cardiac ablations [2, 3] and tumor resections [4] where the surgeon's visibility and frame of reference are limited.

Other areas that can benefit from image-based tracking of organ surfaces are: fusion of various diagnostic images and endoscopic camera feeds [4, 5], instrument guidance and motion compensation [6, 7], and deformation/strain assessment of cardiac structures [8].

Fusion of different types of medical imagery can be useful in various types of robotic and endoscopic surgery. For example, by projecting an ultrasound image of a liver tumor onto a live camera feed of the moving organ surface, the surgeon can make better decisions about instrument positioning and tissue removal [5]. These types of augmented reality (AR) overlays have shown great promise in assisting with tasks such as tumor location, but accurate, precise tracking to the sub-millimeter scale continues to be a challenge [4, 9].

Many endoscopic procedures, such as ablation of specific tissue regions to destroy tumors or reduce internal hemorrhaging have become commonplace [10], while more complicated procedures such minimally-invasive Coronary Artery Bypass Grafting (CABG) continue to develop and

improve [11]. These procedures require precise instrument movements synchronized with moving tissue surfaces or normal breathing motions. A robust tissue tracking algorithm could be used in such procedures to either mechanically compensate for organ movements [11, 12], or to trigger timing of ultrasound therapies to direct energy to a specific target [6].

Tracking the deformation and strain of cardiac tissues has become an important monitoring tool during complex heart surgeries. Techniques have been developed to calculate tissue strain values via the displacements derived from optical flow in cardiac videos [8, 13]. These types of measurements can serve as a "virtual" strain gauge indicating a relative load on the heart during surgery.

The original algorithm that provides the basis for the research presented in this thesis has demonstrated that accurate tracking of these non-rigid organ surfaces in 2-D imagery without *a priori* understanding of the underlying motion is possible [14]. However, obscured surface features cannot be tracked without some understanding of the underlying motion during extended periods of surface occlusion.

Throughout any typical surgical procedure, numerous occlusions and obstructions can enter the camera's field-of-view (FOV) including those caused by instruments, surgical materials, glare, or, in the case of open surgery, the surgeon's hands. Figure 1.1 shows an example of a clamping instrument and surgical thread in the FOV that could block a region-of-interest (ROI).




Figure 1.1: Example of Occlusions from Data Set. The left arrow points to a surgical thread and the right arrow points to a clamping instrument.

This research aims to improve upon the original work through various refinements to the algorithm and through implementation of a more optimized computer vision software library. A machine learning model will be used to model and predict the motion of the ROI.

The proposed algorithm framework will:

- Modify the image pre-processing in the original tracking algorithm to improve identification of image features.
- Improve the ROI re-acquisition strategy based on the cyclic motion of the target.
- Predict the future motion of an occluded region based on a window of previous motion.

1.2 Organization of Thesis

This research is presented as follows: Chapter 2 will review background relevant to the original algorithm and sequence prediction. The underlying methods used for algorithm development and

the experimental design will be presented in Chapter 3. Lastly, results, discussion, and conclusions will be presented in Chapters 4, 5, and 6 respectively. All of the source code used for this project is listed in the Appendices.

CHAPTER 2

BACKGROUND INFORMATION

This chapter provides a review of topics relevant to the proposed work. The first section will highlight some key points from the original research [14]. The second will discuss point sequence prediction as a time series forecasting problem.

2.1 Previous Work

The research that established the original algorithm developed a method for tracking non-rigid surfaces based on the establishment of constraints (statistical tests) for the optical flow of distinct feature points within small, piecewise areas (patches) of a particular region (See Figure 2.1). Feature points are distinctive locations in an image that are detected through a particular detection algorithm (in this case Shi-Tomasi [15]). Figure 2.2 shows a simplified block diagram of the original algorithm for reference.




Figure 2.1: Example of Region-of-Interest (ROI). The grid area is divided into smaller, piecewise areas (patches).

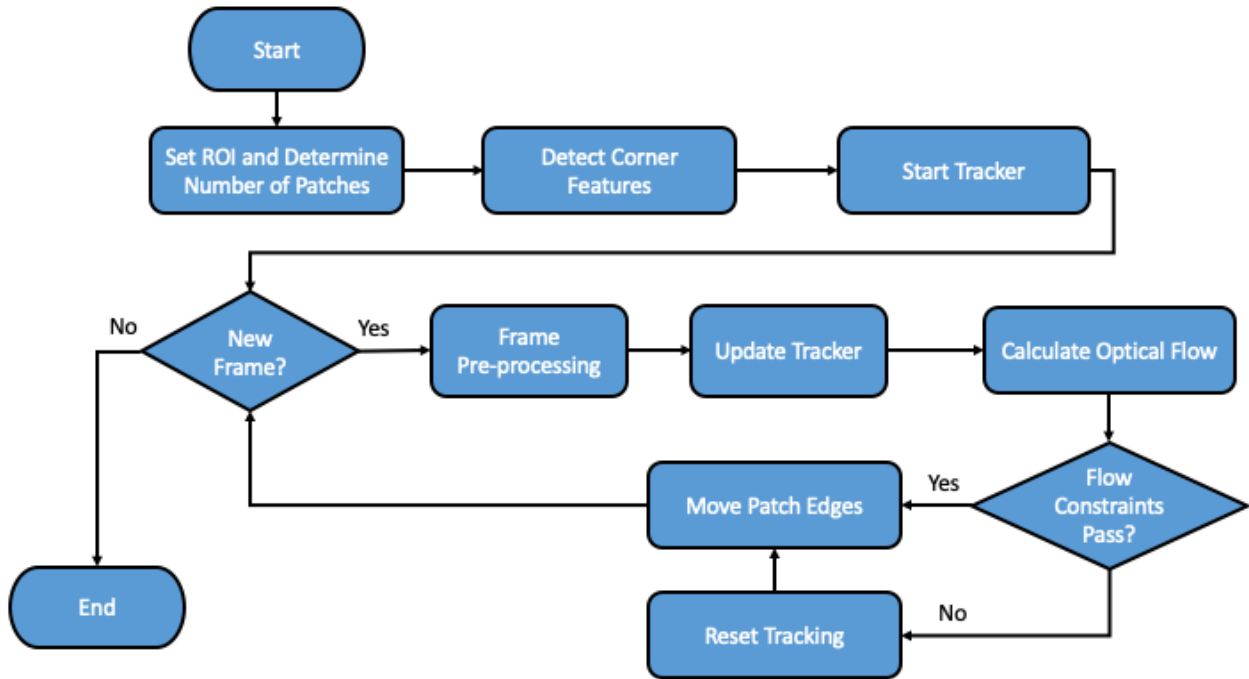


Figure 2.2: Simplified Block Diagram of Original Algorithm [14].

These tracking constraints were: 1) the motion of feature points must follow a normal distribution within a patch and 2) the motion of a particular patch must follow a normal distribution relative to its neighboring patches. If these constraints are not met, tracking is re-initialized using a more computationally expensive homography transform which warps the perspective of the grid points to match a reference frame [16].

The original algorithm was able to continue tracking under some cases of partial occlusion, but failed when the ROI was completely obscured. To improve the performance during a total occlusion, the original research proposed the addition of machine learning-based ROI prediction. By using some knowledge of the previous motion history of the ROI grid points before occlusion, a prediction of future grid points could be made until the occlusion is cleared. This translates into a multiple point sequence prediction problem.

2.2 Time Series Forecasting

The topic of point sequence prediction can be grouped under the larger category of time series forecasting. Research has been done on many different types of data to predict future values based on past history. Much of this research focuses on financial data analysis [17] and scientific sensor data [18].

Most of these prediction approaches fall into two categories: deterministic models or machine learning methods. The following sections will discuss these methods in greater detail along with some of the characteristics of time series data important to forecasting.

2.2.1 Characteristics of Time Series Data

One of the fundamental assumptions for forecasting a time series with traditional techniques is stationarity. If a data set is said to have stationarity, one or more of the statistical distribution measures (mean, variance, and co-variance) are constant. Without at least one constant statistical property, a signal would be considered noise and a reliable prediction would not be possible [19].

A typical time series can be separated into three components: a trend, seasonal (cyclic) variation, and random variation. Figure 2.3 demonstrates how each of these components contribute to the overall series.

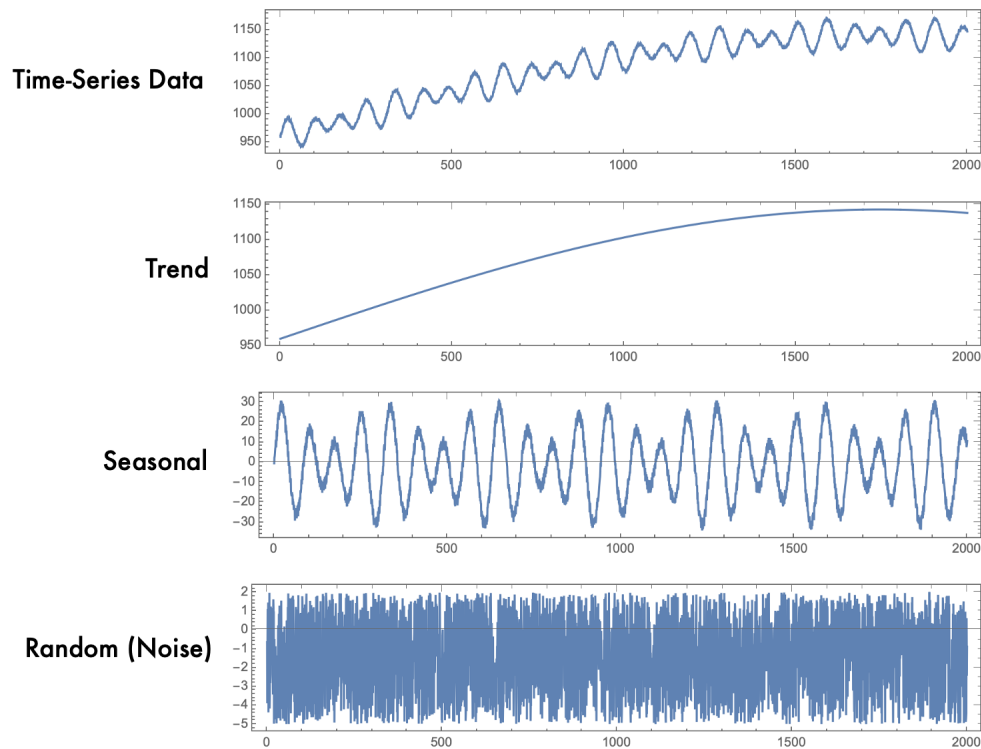


Figure 2.3: Components of a Time Series. The trend, seasonal, and random parts of the series sum to form the composite signal.

A trend is a long term directional change in the data that can be either linear or nonlinear. Seasonal (periodic cyclic) components are repeated, long or short-term patterns that occur throughout the data series. On top of this trend/seasonal structure, a degree of randomness or noise from error and unpredictability in measurements is typically present [19].

To achieve stationarity in a time series, mathematical transformations are performed to remove the trend and seasonal components leaving an aperiodic data series. This can be achieved with various filtering and differencing methods [19].

Time series data can also be classified as univariate or multivariate. Univariate time series data consists of the values of a single variable over a period of time such as the daily temperature over a

year. Multivariate time series data consists of multiple variables such as temperature, rainfall, and atmospheric pressure over a concurrent period of time.

2.2.2 Deterministic Models

A deterministic model would attempt to use mathematical formulas to strictly define the motion based on expected patterns. The development of deterministic models is difficult in problems dealing with non-rigid tissue motion because of the complex spatial dynamics of the surface, thus this research will not focus on the use of deterministic models.

2.2.3 Machine Learning Methods

Various machine learning techniques have been developed to perform time series forecasting on collected data. The two main methods found in the literature are Autoregression and Neural Networks.

Autoregression

Autoregression analyses are a form of state-space models that use linear equations to predict future values based on differences with previous values at various lag times. The state-space model consists of two equations: an observation equation describing the output of a system and a state equation describing a system's current state based on a previous observation [19].

The Autoregressive Moving Average (ARMA) model is one of the standard analyses for stationary data and the Autoregressive Integrated Moving Average (ARIMA) model is used for non-stationary data [19]. Univariate and multivariate versions of both of these techniques have been developed.

A multivariate autoregression could be used to analyze several time sequences at once and make a prediction for each sequence. An advantage of this approach would be that spatial relationships between the time sequences could be preserved if the spatial locations were fixed.

Neural Networks

MLP Multi-layer perceptrons (MLPs) are one of the most basic types of neural networks consisting of a series of input nodes connected to one or more layers of hidden nodes and output nodes (Figure 2.4). MLPs are classified as feed-forward networks because the flow of information is always from the input nodes through the hidden layers to the output nodes. Therefore, the MLP network is said to have no "memory" or feedback of previous states or inputs [20].

The MLP network is initialized with a rigid input/output structure of a fixed size. While these sizes are fixed, variable size inputs can be used with an MLP structure if the input size is chosen to be greater than the longest individual input and the inputs are "zero-padded" [21], however, the output size will remain a fixed length.




Figure 2.4: Multi-Layer Perceptron (MLP) Diagram.

RNN Recurrent Neural Networks (RNNs) have a structure similar to MLPs with the addition of feedback connections to add a form of "memory" that can be useful in sequence analysis (See

Figure 2.5) [20]. One advantage of this type of network over a standard MLP is that it can fed with variable length inputs to generate variable length outputs.



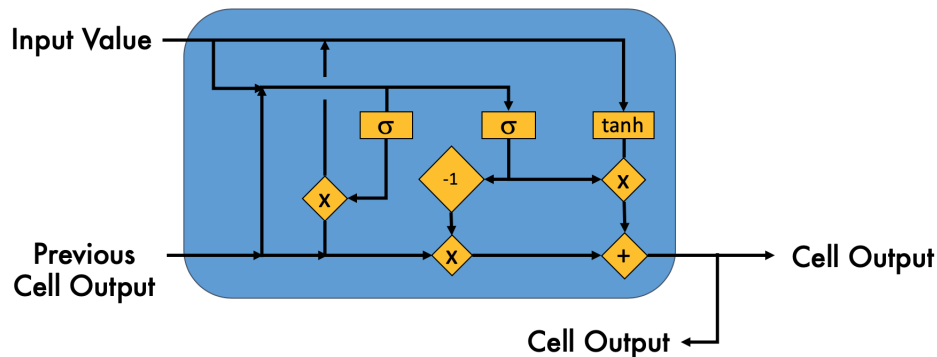


Figure 2.5: Recurrent Neural Network (RNN) Diagram.

One large disadvantage of RNNs is that the memory feedback loop can lead to the problem of exploding or vanishing gradients during backpropagation in the training phase. The errors in the derivative calculations which determine the weighting of the nodes can, in some cases, drive the weight gradients to very small values (vanishing) or to very large values (exploding) [21]. Either case results in a network of limited value that cannot be effectively trained.



(a) LSTM Cell



(b) GRU Cell

Figure 2.6: Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) Nodes.

The addition of alternate information pathways, gates, within the network nodes can help to better control the flow of the information through the network and eliminate the exploding/vanishing gradient issue. The two most well-known modifications to the RNN node are the Long Short-Term Memory (LSTM) cell and the Gated Recurrent Unit (GRU) cell (Figure 2.6) [22]. The LSTM cell has the addition of three gates: the forget gate, the input gate, and the output gate. The GRU is a more recent invention intended to improve processing speed by reducing the number of gates and, therefore, the number of operations. It only has two gates: the update gate and the reset gate. These two network types can have similar performance in many situations, but research to establish the best case uses is ongoing.

CNN A Convolutional Neural Network (CNN) is optimized for learning patterns in two-dimensional data occurring in a regular spatial pattern [20]. This data could take the form of image data or any other 2-D array of values in a grid-like structure. Data flows forward through the network (See Figure 2.7) and is downsampled through a series of convolutions to extract essential features and perform tasks such as classification and segmentation [23].

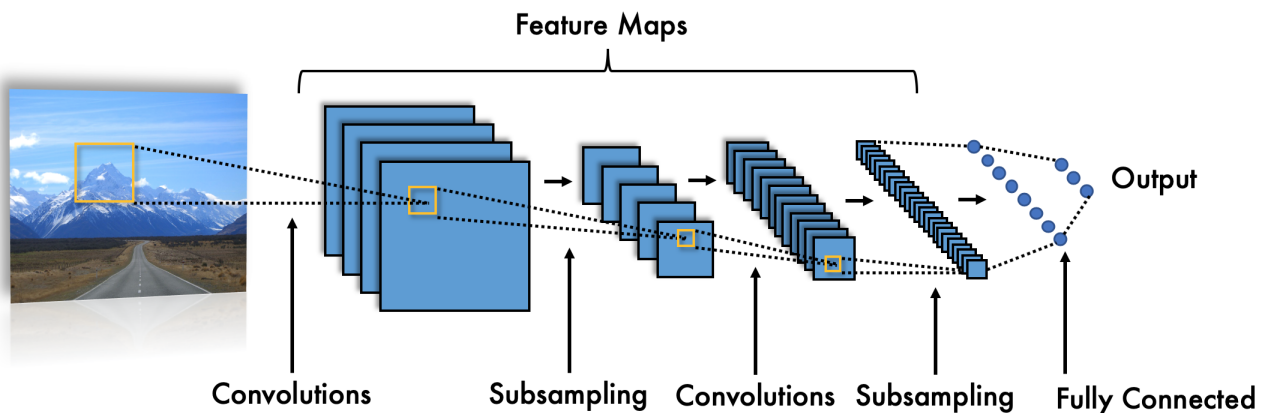


Figure 2.7: Convolutional Neural Network (CNN) Diagram.

CNNs have come into greater prominence after the success of the Krizhevsky et al. team in the 2012 ImageNet competition where image classification was improved by over 10% compared to the previous year's winner [24]. Recurrent CNNs (R-CNNs) have also been successfully applied to many types of video analyses such as video captioning [25], object tracking [26], and frame prediction [27]. The hardware requirements for network training are high [27], but inference operations have been able to generate new complete frames in real-time (20-30 ms per image [~ 30 fps]) on a consumer level GPU [28].

CHAPTER 3

METHODS

The following sections describe the modifications to the original MATLAB code, the overall framework for incorporating the occlusion prediction into the original algorithm, the rationale for the selected software packages and neural network, and the structure of the experimental tests of the framework.

3.1 Programming Framework

The overall structure of the research will be arranged as follows:

1. Rewrite the original algorithm [14] to improve tracking performance and support a future real-time implementation.
2. Develop a feature extraction method to create a point motion sequence data set.
3. Train the selected machine learning model using the point motion sequence data set.
4. Integrate the machine learning model into the main algorithm for prediction of the ROI.
5. Continue ROI tracking using the machine learning prediction.
6. Assess the performance of the tracking prediction.

Figure 3.1 shows a simplified block diagram of the improved algorithm.

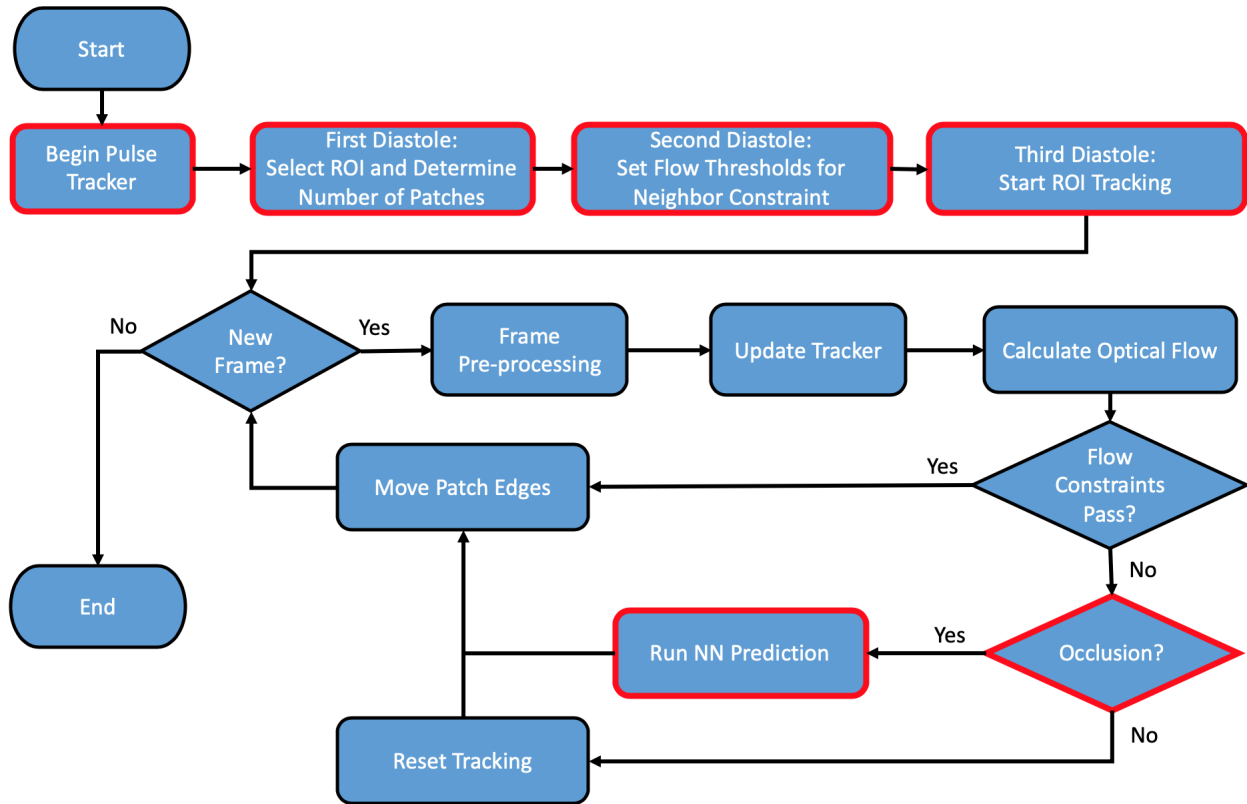


Figure 3.1: Block Diagram of Improved Algorithm. The red outlines indicate added routines.

3.2 Software Packages and Development Environment

The following is a list of the main software packages and versions used in this research:

- MATLAB R2019a
- Linux - Ubuntu 18.04 LTS
- Python 3.6.9
- OpenCV 4.2.0-dev
- NumPy 1.18.1
- TensorFlow 2.0.0

The next sections give some background on each package and rationale for its use in this research.

3.2.1 MATLAB

MATLAB is a programming environment/language developed by Mathworks that provides the ability to create programs and run computational code for an array of scientific analyses. It has the capability of running various image processing and computer vision algorithms and was used in the original algorithm research [14].

In recent years, MATLAB has added more capabilities to interface with C++, NVIDIA Graphics Processing Units (GPUs), and Python. Mathworks continues to improve the speed of its algorithms by providing access to C++ functions through MEX [29], GPU-enabled algorithms, and machine learning models that are claimed to be up to 7x faster than TensorFlow [30]. These approaches may be viable, but the practical implementation of combining and optimizing all of these methods was beyond the scope of this research.

Therefore, a more streamlined approach was taken to keep everything open-source (free) and cohesive by using the Python language combined with NumPy, SciPy, OpenCV Python, and TensorFlow. In addition, many of the MATLAB image processing functions are not optimized for a GPU and the transfer of data between the CPU and GPU requires extra programming attention to optimize speed. Mismanagement of the CPU-GPU data transfer can make the code run slower than without a GPU [31].

3.2.2 Linux

Development for this project was done under the Linux operating system (Ubuntu 18.04 LTS [Long Term Stable]). Linux is a free, open-source operating system and its ability to scale to many different types of hardware could be useful for future implementation in a real-time system or actual product.

3.2.3 Python

Python has become more popular as a programming language in recent years for scientific computing due to its open-source nature, clear syntax, and the proliferation of add-on packages that provide access to many essential programming functions [32].

C and C++ are traditionally the most popular languages for performance-driven computing applications especially in embedded systems [33]. However, increasingly Python is being used to create bindings to faster C++ functions to achieve similar performance with more intuitive, understandable code. The Cython language, a superset of Python, can also be used to generate C or C++ code from regular Python code for faster performance [34]. These are additional options that may improve the performance of the code but were not implemented in this research.

Packages such as NumPy [35], SciPy [36], Scikit-learn [37], and Matplotlib [38] have given Python expanded functionality for various mathematical and scientific computing operations. NumPy adds support for array generation and manipulation similar to that found in MATLAB. SciPy adds many scientific computation modules including modules for signal and image processing and statistics. Matplotlib adds graphical plotting capabilities similar to those available with MATLAB.

3.2.4 OpenCV

OpenCV (Open-source Computer Vision) was developed as a cross-platform real-time computer vision library and has both Python and C++ versions [39].

The conversion of the original MATLAB code to OpenCV provides several advantages. OpenCV Python allows easy integration with other Python-based packages such as TensorFlow and NumPy. OpenCV was developed for real-time systems and runs on Linux allowing applications to be more easily scaled and ported to various types of hardware. Finally, OpenCV is open-source and free-to-use compared to the cost of a non-educational MATLAB license.

At this time, the Python version of OpenCV is not optimized for the CUDA language. Compute Unified Device Architecture (CUDA) is a software interface for graphical processing units (GPUs)

created by NVIDIA, the leading manufacturer of GPUs [40]. CUDA allows many operations to be performed in parallel on NVIDIA GPUs including the types of matrix operations common in image processing and machine learning.

OpenCV Python can, however, make use of OpenCL acceleration via the Transparent API (T-API) added in OpenCV Version 3 [41]. OpenCL is a more general GPU language that can run as an alternative to CUDA on non-Nvidia GPUs. It can also run on Nvidia GPUs, but has much slower performance than CUDA. It is possible to create Python bindings for the C++ functions and achieve increased speed using OpenCV with CUDA instead of OpenCL.

Due to the frame-by-frame nature of the tracking algorithm, the GPU was not used for processing during tracking as these image processing tasks are better performed in large batches that minimize the transfer of data between the system and GPU memory.

3.2.5 TensorFlow

TensorFlow is an open-source machine learning library developed and maintained by Google that is compatible with Python, C++, and CUDA. Version 2.0 included the popular high-level machine learning interface Keras as part of the core TensorFlow package. TensorFlow Lite allows the flexibility to run models on less powerful devices [42].

3.2.6 Hardware

Image processing, neural network training, and final algorithm testing were performed on a custom-built workstation containing the following components:

- Intel Core i9-9900X 3.5 GHz 10-Core Processor
- Zotac NVIDIA GeForce RTX 2080 Ti 11 GB Video Card
- Asus WS X299 SAGE SSI CEB LGA2066 Motherboard
- Corsair Vengeance LPX 32 GB (2 x 16 GB) DDR4-2133 CL13 Memory
- Samsung 960 EVO 500 GB M.2-2280 NVME Solid State Drive

These components are near the top of the performance ladder available for workstation desktop PCs at the time of this research.

3.3 Improvements \ Changes to the Original Algorithm

The following sections describe the improvements and modifications that were made to the original algorithm including changes to frame pre-processing, a new method to determine the initial feature detection frame and detect the pulse cycle, a modification to the neighboring patch constraint, and modifications to the tracking reset routine.

3.3.1 Frame Pre-processing


To improve feature point detection and tracking in the original algorithm, two standard imaging processing techniques were applied to each frame: a histogram equalization and denoising (blur) filter. These techniques are still used in the new framework, but they have been modified to give better results as described in the following paragraphs.

Equalization


Histogram equalization enhances the contrast of the image by spreading the pixel intensity levels across the full spectrum of grayscale levels. By improving contrast, features on the surface will be more distinct and more easily found by the feature detection and tracking algorithms.

The histogram equalization algorithm used in the original research was the Matlab function "histeq" with its default settings [43]. For this research, improved performance was found by using the Contrast Limited Adaptive Histogram Equalization (CLAHE) algorithm of OpenCV [44]. This algorithm performs a similar function as the "histeq" algorithm multiple times on smaller tiled sections of the image with an additional option to limit the amount of level changes. A 20x20 pixel tile with a level clip threshold of 30 was found to give good performance on this data set.

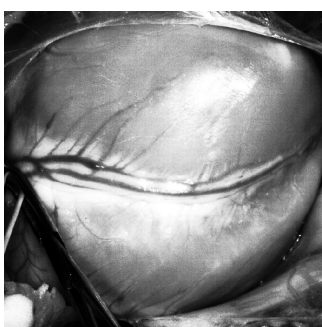
Figure 3.2 shows an example of the difference in enhancement between these two equalization algorithms. The original images were very dark near the low end of the intensity scale. The basic Matlab and OpenCV equalization algorithms give different results but both spread the intensity levels across the full grayscale range. The CLAHE algorithm gives much greater contrast across the image than the basic equalization algorithm with a more dense histogram. The main coronary artery is much more defined in the equalized image using this method. It is expected to provide better point feature tracking performance. The highlights in the areas of glare (upper right center) are also greatly reduced helping to improve tracking performance.




(a) Unprocessed Image




(b) Unprocessed Image - Histogram




(c) MATLAB Basic Equalization [14]




(d) MATLAB Basic Equalization - Histogram [14]




(e) OpenCV Basic Equalization



(f) OpenCV Basic Equalization - Histogram



(g) OpenCV CLAHE Equalization



(h) OpenCV CLAHE Equalization - Histogram

Figure 3.2: Comparison of Histogram Equalization Methods.

Denoising

Denoising, or blurring, can remove or reduce small, speckle-like artifacts in the image by applying a smoothing filter. This helps the feature detector or tracker to focus on the strongest image features instead of the undesired noise artifacts.

OpenCV has multiple denoising algorithms: the box filter, gaussian filter, median filter, and bilateral filter. The first three of these algorithms are effective in reducing noise but can also smooth over the edges in the image [45]. The bilateral filter takes the pixel intensity of neighboring pixels into account to reduce noise while better preserving edges [45].

The original research used the median filter in MATLAB. After some experimentation using the different blurring filters, this research chose to use the bilateral filter for its slight improvement in retaining edges which are important for optimal feature detection. Figure 3.3 shows a comparison of the unblurred image, median filter, and bilateral filter. The difference is subtle, but there is a slight increase in edge contrast with the bilateral filter (more apparent in the magnified images).




Figure 3.3: Comparison of Denoising Methods. The Bilateral filter reduces speckling while better preserving edges.

3.3.2 Feature Detection Algorithm/Parameters

The original algorithm uses the MATLAB function “detectMinEigenFeatures” to detect the initial feature points [14]. This function incorporates the Shi-Tomasi minimum eigenvalue algorithm to find feature points [46].

The equivalent OpenCV function is “goodFeaturesToTrack” which uses the Shi-Tomasi algorithm [47]. By setting the “useHarrisDetector” parameter to “False” or “0,” this function will use the minimum eigenvalue calculation and operate similar to the MATLAB function.

The “qualityLevel” parameter was decreased to 0.001 from the 0.005 value used in the original research which gave a better distribution of points within the ROI.

3.3.3 Tracking Algorithm/Parameters

The original algorithm uses the MATLAB function “vision.PointTracker” to track the set of detected points throughout the image sequence [48]. This function incorporates the Kanade-Lucas-Tomasi (KLT) feature-tracking algorithm [49].

The equivalent OpenCV function is “calcOpticalFlowPyrLK” which uses the iterative Lucas-Kanade method with pyramids [50].

The implementation of the OpenCV KLT tracker is different from the MATLAB KLT tracker and does not include the error parameter, “MaxBidirectionalError,” that is set to “1” in the original research. This parameter does a calculation forward and backward between frames to double-check the tracking calculation and uses this difference as an error criterion at the cost of additional computational overhead [48]. With a setting of “1”, the tracking would be considered invalid if the difference in the forward and backward tracking location was greater than 1. According to the OpenCV documentation, this backward check is implemented into “calcOpticalFlowPyrLK” but control of this error criterion is not directly accessible [50].

3.3.4 Initial Frame Selection

A recommendation was made in the previous research to use machine learning to assist in the selection of the initial feature point detection frame. It was also found that the diastole point of the cycle was the best frame to use for initialization of the feature points [14]. This finding makes intuitive sense because the tissue surface is most relaxed and stretched during diastole and thus covers more surface area (pixels) providing more information to the feature detector.

While a machine learning approach may be viable, it would require an effort to collect and label many examples of heart images at the diastole phase and training of another classification algorithm. A more simple approach can make use of the pulse detection algorithm described in the next section to determine the approximate frame corresponding to the diastole phase.

3.3.5 Pulse Detection

A method for detecting the pulse cycle was implemented in order to assist in the initialization of feature detection described in Section 3.3.4 and the re-initialization of tracking. An additional benefit of pulse cycle detection is the calculation of an approximate heart rate based only on the video imagery.

To detect the cycle, a sparse KLT feature point tracker (using the 10 strongest features across the entire frame) is started at the beginning of the video sequence. The average of the absolute Euclidean distance of all of the tracked points from their initial location is calculated for each frame. This distance should have a rough correlation to the overall motion in the frame.

The peaks of this average distance should correspond to the approximate systole and diastole points of the cardiac cycle where motion is at a minimum and changing direction. By tracking the direction of the distance changes, the type of peak can be determined (i.e. increasing distance before a peak would correspond to a diastole [relaxation expansion] and decreasing distance would correspond to a systole [contraction]).

This cycle detection runs continuously in the background independently of the main feature tracker for the selected ROI. An example of this overlay of the optical-flow based cycle is shown in the red graph at the bottom of Figure 3.4.

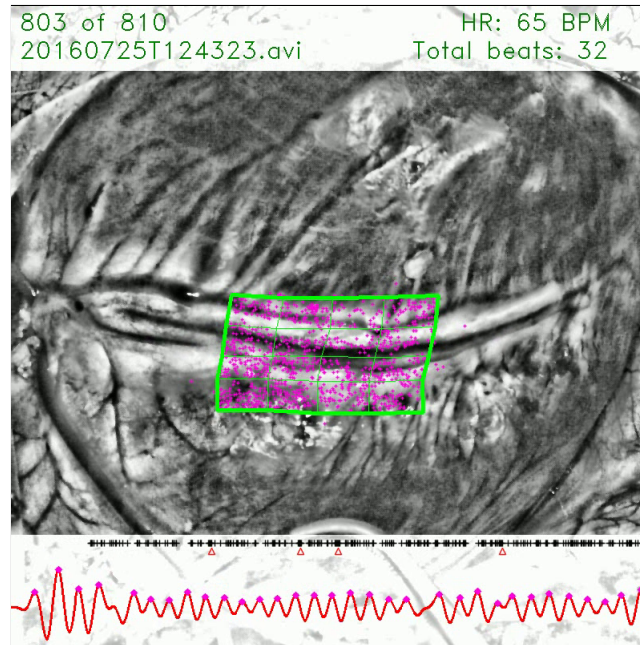


Figure 3.4: Optical Flow from Feature Point Movement.

3.3.6 Determination of Number of Patches

The function to determine the number of patches in the original algorithm relied on incrementally increasing the number of patches until the two constraint conditions described in Section 2.1 were met. The patches are set to increase as a square matrix (i.e. 2x2, 3x3) to a maximum of 36 patches (6x6). This is similar to the method in the original algorithm.

If the maximum number of patches is reached without satisfying the constraint criteria, the calculation will proceed, but a warning will be displayed.

3.3.7 Neighboring Patch Smoothness Constraint

Criteria

The neighboring patch smoothness constraint (Section 2.1, Constraint 2) from the original algorithm tested the normality of the optical flow difference values of each patch's corresponding

row and column and reinitialized tracking if this condition was not met (See Examples in Fig. 3.5 a, b).

Patch

| Numbers | 1 | 2 | 3 | 4 | 5 |
|---------|---|---|---|---|---|
| 1 | 4 | 4 | 1 | 3 | 2 |
| 2 | 5 | 9 | 2 | 6 | 0 |
| 3 | 1 | 9 | 2 | 4 | 1 |
| 4 | 2 | 6 | 4 | 3 | 5 |
| 5 | 5 | 2 | 4 | 2 | 4 |

(a) Original Constraint 2 Comparison - Example 1

Patch

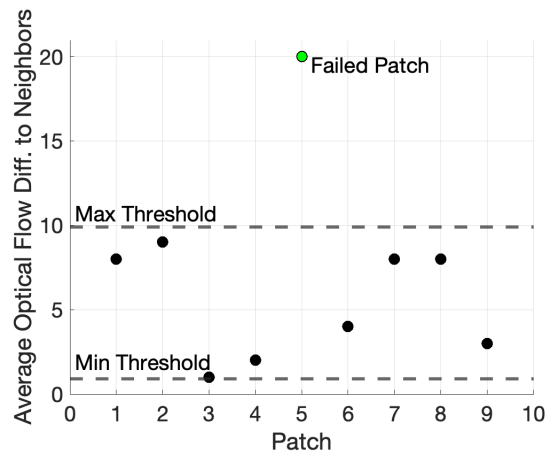
| Numbers | 1 | 2 | 3 | 4 | 5 |
|---------|---|---|---|---|---|
| 1 | 2 | 3 | 2 | 0 | 4 |
| 2 | 1 | 1 | 9 | 4 | 1 |
| 3 | 1 | 3 | 4 | 2 | 4 |
| 4 | 3 | 3 | 4 | 3 | 3 |
| 5 | 1 | 3 | 2 | 2 | 4 |

(b) Original Constraint 2 Comparison - Example 2

Neighbor Average: 5.1

| | 1 | 2 | 3 | 4 | 5 |
|---|---|---|----|---|---|
| 1 | 3 | 2 | 4 | 4 | 2 |
| 2 | 2 | 8 | 9 | 1 | 1 |
| 3 | 4 | 2 | 20 | 4 | 4 |
| 4 | 1 | 8 | 8 | 3 | 1 |
| 5 | 2 | 2 | 2 | 2 | 4 |

(c) New Constraint 2 Comparison to Surrounding Neighbors



(d) Outside Threshold Range -> Failed Constraint Check

Figure 3.5: Smoothness Criteria Constraint.

Because this algorithm consisted of several nested loops that were slow, unclear, and difficult to modify, a different measure of smoothness was implemented in this research. Instead of calculating multiple U (horizontal) and V (vertical) flow differences between each patch in the entire row and column and then checking normality, the average U and V flow for each patch will be calculated and compared to the average of the neighboring patch flows (See Figure 3.5 c).

This allows for a simpler constraint check as a comparison of the individual patch flow to the expected range of neighbor flow. In addition, the influence of the immediate diagonal neighboring patches is now considered without the influence of distant patches. The next paragraph describes how the flow threshold range was determined.

Threshold Range

After the number of patches are determined during the first full heart cycle, the difference between the average optical flow of each patch and the average optical flow of its neighbors is monitored during the second heart cycle. Each patch is assigned a maximum and minimum threshold range (plus margin) based on the observed difference values for later comparison during tracking. Any patch with an extreme difference from its neighboring patches would violate these thresholds and the patch would be flagged as failing the constraint check (See Fig. 3.5 d).

The original algorithm included analysis of several frames in order to determine a maximum threshold for the smoothness across a cycle. This was performed in the same routine used to determine the number of patches and slowed program execution during calculation. In this research, this threshold determination procedure was moved into a separate routine to allow for more options to calculate the threshold while the video sequence is moving forward.

3.3.8 Tracking Re-initialization

The original algorithm re-initialized tracking after the constraint check failed on three consecutive frames. Once three failures were detected, another routine began to search backward from the current frame through previous frames to find a best homography match. This homography was used to transform the region of interest boundaries and reacquire tracking points.

This algorithm has a long execution time as implemented and could benefit from several improvements. The backward search through multiple frames performs the same calculations multiple times which is inefficient for use in a real-time application. Performing operations in a forward

direction would be more desirable in a live video stream application. Thus, in this research, the original ROI selection frame and ROI grid are saved and used to obtain a homography matrix when tracking is lost. The homography matrix is used to warp the original ROI grid to the current frame.

In scenarios where frames with occluded ROIs could be compared to the non-occluded selection frame, re-initialization error is likely to be high, triggering more constraint failures and attempts to re-initialize tracking, but tracking should return to normal when the occlusion is cleared. Future work could add changes to limit this undesired behavior.

A homography re-initialization is also done at each diastole cycle whether or not there has been a tracking failure. This modification should help to reduce tracking failures in general and periodic tracking resets are recommended by the OpenCV development team for more robust tracking [51]. While each added re-initialization slows performance, it was anecdotally observed that less total re-initializations were needed with periodic resets compared to waiting for tracking failure.

3.3.9 Parameter Values

Optimization of the tracking parameters is not the focus of this research, but several observations were made while re-factoring the algorithm. The quality of unobstructed tracking depends on a balance of the number of initial points tracked, the tracking failure threshold number, size of the ROI, and the presence of glare.

Ensuring that there are a minimum number of points within each patch to meet the statistical normality tests is a good guideline for setting the number of corner feature points to detect. If the maximum number of patches is set to 36 (6 x 6) and the minimum number of points to meet the normality test is 20, a minimum of $36 \times 20 = 720$ points would be required. More points than the minimum required should lead to better tracking that is less sensitive to outliers. The detection threshold was set to 1000 features in this research to give 720 points plus a margin. Care should be taken to avoid using an excessive number of points which would slow down the overall calculation

time. A dynamic patch point threshold of 20 points per patch (plus a margin) could be set to reduce excess feature points and improve run times.

The threshold for the number of tracking failures before a reset can be increased if the tracking error is relatively low. Setting the failure limit too low can lead to frequent, perhaps unnecessary, re-initializations that can cause slow-downs and potential re-initialization errors. The failure limit setting is also closely related to the frame rate of the video under analysis. If the frame rate is very high, multiple failures may occur over several frames that represent a very short length of time. Setting the failure limit high could be warranted in this case to prevent resets from occurring too often. Likewise, if the frame rate is very low, movement (and error) could be higher from frame to frame and a lower failure limit would be justified. The tracking failure limit in this research was able to be increased to five consecutive failures without excessive tracking resets.

In general, larger ROIs with larger patch sizes were tracked better than smaller overall ROIs or ROIs with smaller patches. This is likely due to a larger ROI containing a greater number of distinctive features overall and more detected features within each patch. This may vary with the resolution of the frame as well.

ROI areas with glare seemed to generate more overall tracking resets. Glare is similar to an occlusion in the sense that it hides the tissue surface information from the view of the camera. However, some amount of glare is likely always visible in biological images due to moisture and ambient lighting. This may be reduced using the CLAHE equalization described in Section 3.3.1, however, it will not be completely eliminated because the pixel intensity has been saturated to a maximum value and texture detail will be lost in the overexposed areas.

3.3.10 Occlusion Detection

No changes were made in this research to improve the detection of occlusions in the original algorithm. The original research assumed that a partial occlusion would trigger constraint failure and an eventual re-initialization. Failure of constraints on all of the patches would be considered a

full, non-recoverable occlusion and may require a pause of tracking until the occlusion is removed [14].

In this research, a full occlusion will be simulated by switching the ROI grid from the tracking algorithm to a prediction from the machine learning model for a specified selection of frames. A shaded rectangle will be displayed to indicate an occlusion, but this is cosmetic and does not trigger the prediction or affect the tracking algorithm.

In a final implementation, additional work may be needed to define criteria for switching to a neural network prediction or triggering a tracking re-initialization. Actual occlusions would have several variables such as the speed across the ROI, the size relative to the ROI, and the period of occlusion which could trigger unwanted and inaccurate re-initialization behavior.

Expected types of real occlusions fall into three categories: glare, shadows, and objects. Routines could be created to deal with the unique properties of each of these separately. Glare causes a complete saturation of intensity levels, shadows cause a partial decrease in intensity levels, and an object introduces new features and edges that can confuse the feature detector/tracker.

3.4 Feature Extraction

3.4.1 Data Set

Five data files containing image sequences of cardiac motion were provided by Dr. Bruce Ferguson and Dr. Cheng Chen for the previous research [14]. These images were obtained during an unrelated, previously conducted porcine (pig) animal study conducted at ECU facilities under Institutional Animal Care and Use Committee (IACUC) supervision and approval.

The image sequences contain an upright, anterior view of the heart with the coronary artery in a horizontal orientation within the image frame as shown in Figure 3.6. The heart rate for all of the sequences is approximately 60-70 beats per minute (bpm). The image intensity levels are very low (dark) without any level adjustments.



Figure 3.6: Sample Data Set Image Frame. Raw image levels adjusted for clarity.

For the majority of the heart cycles, most of the heart surface is in clear view with minimal occlusion of the moving surface. Occasionally, a pair of forceps or surgical threading material moves across the heart surface temporarily obstructing a part of the surface. This occurs mainly in the lower left area of the frame.

The files were provided in a MATLAB data file format (.mat). Each file contains approximately 2,430 grayscale, 8-bit image frames at 800 x 800 resolution captured at 81 frames per second (fps). This corresponds to a total image sequence time of 30 seconds.

The original filenames are: 20160725T124323.mat, 20160725T124609.mat, 20160725T124909.mat, 20160725T132840.mat, 20160725T133113.mat. The filenames correspond to the date and time of image capture.

A MATLAB script (Appendix A.6) was created to convert the .mat image sequence files into uncompressed audio-video interleave (.avi) video files that can be read by the OpenCV file input algorithms.

To increase the speed of processing, each file was downsampled by a factor of 3 to 27 fps. The original files were captured at a high camera frame rate that is not necessary for accurate tracking and 27 fps is closer to the 30 fps framerate that is typical of many commonly available cameras.

3.4.2 Extraction Method / Algorithm

An OpenCV Python script (Appendix A.2) was created to perform feature point detection and tracking on each video and save each point sequence to a NumPy data file (.npy). This tracking script uses the pulse detection method described in Section 3.3.5 to locate the peak values of the flow corresponding the systole and diastole phases of the cycle.

Initial feature detection will begin at the first diastole frame and tracking will continue through the remainder of the video as shown in Figure 3.7. Point sequences that the tracking algorithm designates with a "lost" status will not be saved.




Figure 3.7: Sample of Extracted Feature Points Across Tracking Period.

The points are detected in a circular region in the center of the image to avoid collecting unwanted data from tracking points on the surgical instruments or from stationary tissues. This will limit predictions to the central area of the image frame.

Each frame is pre-processed before feature detection and tracking using the equalization and denoising techniques described in Section 3.3.

3.4.3 Sequence Length

The continuous data sequence for each point can be divided up into sub-sequences of arbitrary length. These smaller sub-sequences will be used as training data for the neural network.

A sub-sequence length of 10 time steps (frames) was chosen. At a heart rate of approximately 70 bpm and frame rate of 27 fps, each cardiac cycle would cover approximately 32 frames and each 10 step sub-sequence would contain information for approximately 1/3 of a cardiac cycle.

With approximately 7000 points detected in each video and 810 frames per video, this would generate over 5.6 million training sequences per video and over 28 million sequences for the entire data set.

Table 3.1 shows a typical example of an extracted point sequence that will be used to train the neural network. Out of the 10 steps captured, the first nine points will be the training features (used to set the weights of the network), and the last point will be a training label (target predicted output of the network). Note that even though pixel image coordinates are normally expressed in integer values, the tracking algorithm uses subpixel (decimal) values to estimate pixel locations. Training and tracking did not perform well when these values were rounded to integers likely due to accumulation of rounding errors.

Table 3.1: Extracted Sequence Example.

| Data Type | Sequence Step | Coordinate Values (x,y) |
|-----------|---------------|------------------------------|
| Feature | 1 | [381.06005859, 410.21243286] |
| | 2 | [380.56082153, 415.64233398] |
| | 3 | [380.24166870, 424.26119995] |
| | 4 | [387.43267822, 427.93359375] |
| | 5 | [396.51394653, 432.11978149] |
| | 6 | [400.73373413, 434.64465332] |
| | 7 | [389.73266602, 428.98709106] |
| | 8 | [369.53717041, 414.84133911] |
| | 9 | [349.72045898, 402.50061035] |
| Label | 10 | [342.07440000, 400.10553000] |

3.4.4 Scaling

While not always required, neural networks typically perform better with normalized data (scaled to -1 to 1 or 0 to 1) or standardized data with a zero mean and unit variance [52].

The entire data set will be standardized using the StandardScaler function of Scikit-learn before it is used in the neural network. Outputs of the neural network will thus be standardized and require inverse scaling to return to image coordinates. The scaler is fit only on the training data set to maintain data independence.

Data should also normally be transformed to remove stationarity [52], but, in this case since the time sequences cover a short period without significant trends no stationarity transform is added.

3.5 Neural Network

This section describes the selected neural network architecture and training parameters.

3.5.1 Architecture

As described in Section 2.2.3, MLPs, RNNs, and CNNs can all be applied to learn patterns in sequential data. This research will focus on the use of the RNN due to its strengths in dealing with variable length inputs and the capability to generate variable outputs. While we will only predict one future sequence point at a time with each new frame, multiple-step predictions could be useful in further experiments that expand upon this research.

The multivariate autoregression and CNN networks should have an advantage in their ability to better capture spatial dependencies between features compared to a single sequence prediction model. However, because of the variable ROI selection and non-ordered nature of feature point detection used in the baseline algorithm, these networks are not practical. The autoregression and CNN networks assume a fixed number of inputs or grid points, but the baseline algorithm can have a variable number of patches and grid points determined at the time of ROI selection. Collection of training data would be complicated by the large number of grid sequences that would need

to be collected in order to cover the entire range of possible grid permutations. In addition, if the prediction of a series of tracked point features is desired rather than ROI grid points, there would be no *a priori* knowledge of the number or order of features nor any consistent structure to the feature locations since any ROI location could be selected. Individual grid point prediction could be desirable in a future iteration where the ROI is partially obscured and additional statistical checks are required for a single patch.

The RNN will be structured with GRU cells to improve speed due to the reduced number of operations (See Section 2.2.3). The model will consist of two GRU layers and one dense layer as shown in Table 3.2.

Table 3.2: RNN Model Layers.

| Number | Layer Type | Activation | Output Shape | Parameter Number |
|---------------------------|------------|------------|------------------|------------------|
| 1 | GRU | ReLU | (None, None, 64) | 13,056 |
| 2 | GRU | ReLU | (None, 32) | 9,408 |
| 3 | Dense | Linear | (None, 2) | 66 |
| Total Parameters: | | | | 22,530 |
| Trainable Parameters: | | | | 22,530 |
| Non-trainable Parameters: | | | | 0 |

The input GRU layer will have 64 units (nodes), the second GRU layer will have 32 units, and the final Dense layer has 2 units to match the size of the desired output (x and y point coordinates).

Each GRU layer uses the default "ReLU" activation function. ReLU stands for Rectified Linear Unit and is a type of function applied to transform each node summation value to an output value in a way that minimizes computations and allows for larger networks that do not suffer from vanishing

gradients [20]. The final layer uses a "Linear" activation function which is recommended for regression outputs [53].

Initial experiments added dropout layers after each GRU layer which can aid in reducing overfitting via a sampling process that simulates an average of many possible network configurations [54]. However, prediction performance for this data set was improved by removing these dropout layers. Predictions using the dropout layer model had large deformation and location errors.

3.5.2 Training

The network will be trained using a split of 80% of the data set as training data and 20% of the data set as testing (validation) data. The testing data will be held back during each training iteration (epoch) and used as a test of how well the network was trained on the other 80%. One measure of optimal training for a neural network is when the overall network losses on the training data are approximately equal to the losses on the testing data [55]. Before the data split, the point sequences are shuffled so that the training and test data are randomly distributed.

The losses will be calculated using the Mean Squared Error (MSE) metric which emphasizes large errors in the predictions and results in worse loss values for these cases [56]. Because the network is trained on individual point sequences, these losses would be related to an individual point prediction and not the overall grid.

The accuracy of the model is a percentage value calculated based on a comparison of how well the model has predicted the last sequence point during inference compared with the training label.

The Adaptive Momentum (AdaM) optimization algorithm was chosen for training as this is becoming more of an industry standard [57]. The optimization algorithm is the particular convergence method used to set the network weights while minimizing the overall loss through the network. AdaM adds some enhancements to the traditional gradient descent approaches to dampen the convergence to a system minimum and thereby greatly speed up training [58].




Figure 3.8: ROI Selection Landmarks. Repeatable points were chosen for use in the manual ROI selection step for each video. The left arrow points to the intersection of the larger artery with a smaller branch. The right arrow points to the middle of a large, circular dark region.

3.6 Testing Scenarios

Because each video will have some variation in the heart surface and orientation at diastole, each tracking sequence will be initiated with manual selection of a ROI based on repeatable landmarks. Figure 3.8 shows the landmarks chosen for the experiments presented in this research. These landmarks were chosen to be near the center of the frame (within the feature extraction ROI - See Section 3.4) and contain a region that may be more clinically relevant (coronary artery). The upper left landmark is at the intersection of the coronary artery and the indicated vessel branch and the lower right landmark is near the center of the dark circular region indicated.

The performance of the grid prediction will be measured over 30 frames. At each frame, the predicted grid points will be added to the grid point sequence history and used as the next input sequence for prediction. Frames 201 to 230 will be selected for the prediction period in each video. The heart cycles are not synchronized across the videos so this frame range will begin at different

points in the cardiac cycle. After the prediction period, normal tracking will continue until the end of the video.

CHAPTER 4

RESULTS

The following sections provide a summary of the tracking and grid prediction results for each data set.

A series of representative screenshots from normal tracking (green grid) and predictive tracking (yellow grid) are shown for each data set. At the top of each screenshot is an overlay that displays the current frame number, the data set name, the estimated heart rate, and the total number of beats counted (diastole peaks). At the bottom of each screenshot is an overlay showing a plot of the pulse cycle (based on overall optical flow) with frame markers above. Each black (+) marker on the bottom overlay represents a frame with failed constraints. Each red (Δ) marker represents a tracking re-initialization due to multiple consecutive constraint failures.

A summary of tracking results is presented in a table and in a timeline graph. Each of these figures show: the total number of frames in the video, the number of frames where the tracking algorithm was active, the number of diastole peaks detected (and corresponding homography re-initialization), the number tracking resets (undesired homography re-initialization due to consecutive failure of the tracking constraints), the number of frames that had a failure of one or more constraints (flag raised), and the number of frames where the grid prediction was active. The tracking resets and failed constraints percentages are relative to the number of tracked frames.

The prediction performance is presented in a graph that shows the distance between the predicted grid points and the actively tracked grid points (ground truth) versus the predicted frame number. These values are shown as differences in the x-direction only, the y-direction only, and the absolute Euclidean distance in the x-y plane. The average of all the differences across the entire

grid is shown by the main marker. Error bars have been added to each marker to show the maximum and minimum grid differences for each ROI prediction. This allows the graph to show the overall spread of differences for the entire ROI without showing each individual grid difference.

4.1 Results for 20160725T124323.avi

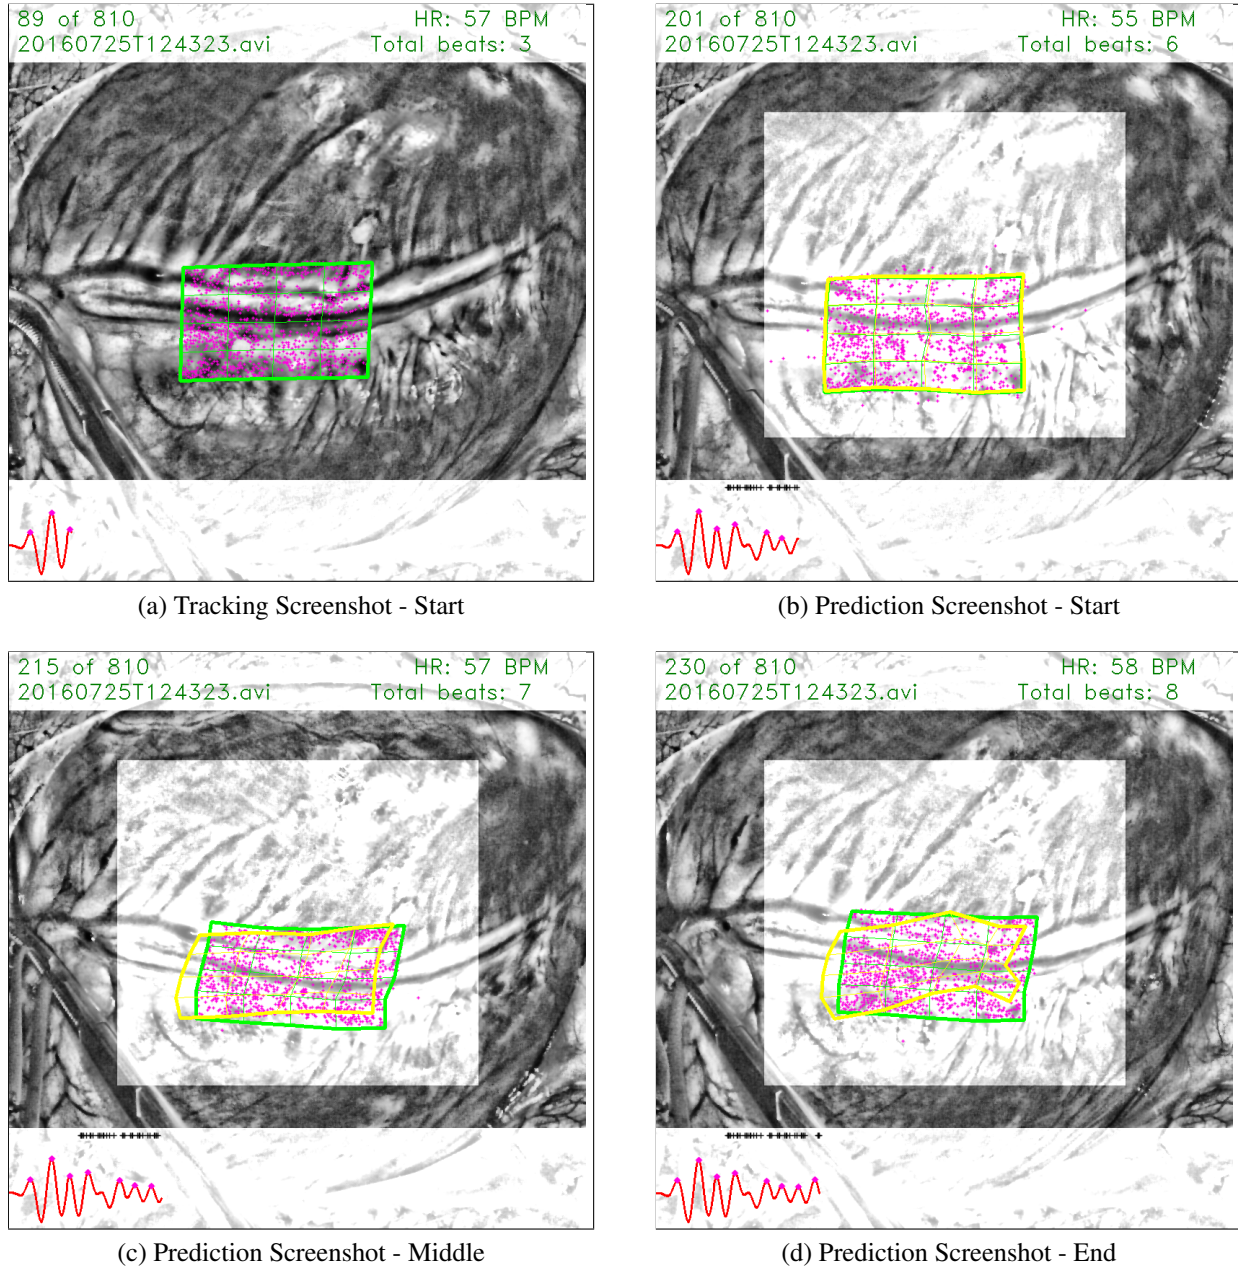


Figure 4.1: Representative Screenshots (20160725T124323.avi).

Figure 4.1 shows that the prediction begins close to the tracked ROI and continues to deviate as the prediction continues. In the middle of the prediction, the predicted grid has a similar shape

as the ground truth. By the last prediction frame, the predicted grid has remained centered in the same area, but it is further deformed compared to the ground truth.

Table 4.1: Tracking Performance Summary (20160725T124323.avi).

| Parameter | Total | % of Total |
|------------------|-------|------------|
| Total Frames | 810 | - |
| Tracked Frames | 721 | 89.0 % |
| Diastole Peaks | 32 | 4.0 % |
| Tracking Resets | 4 | 0.6 % |
| Flag Raised | 234 | 32.5 % |
| Predicted Frames | 30 | 4.0 % |

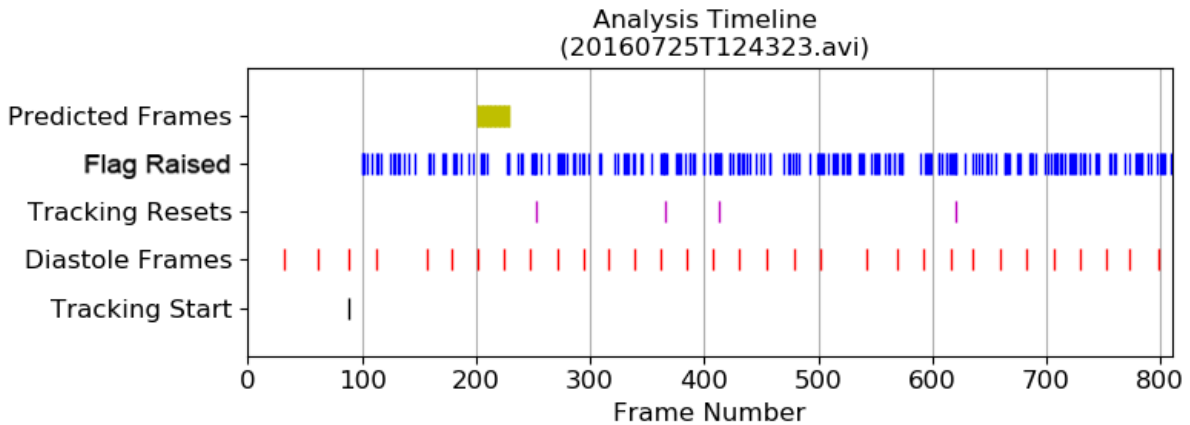


Figure 4.2: Analysis Timeline (20160725T124323.avi).

Figure 4.2 shows that there are a large number of frames where the optical flow has not passed both of the tracking constraint criteria (flag raised) but only 4 frames where the tracking was forced to reset.

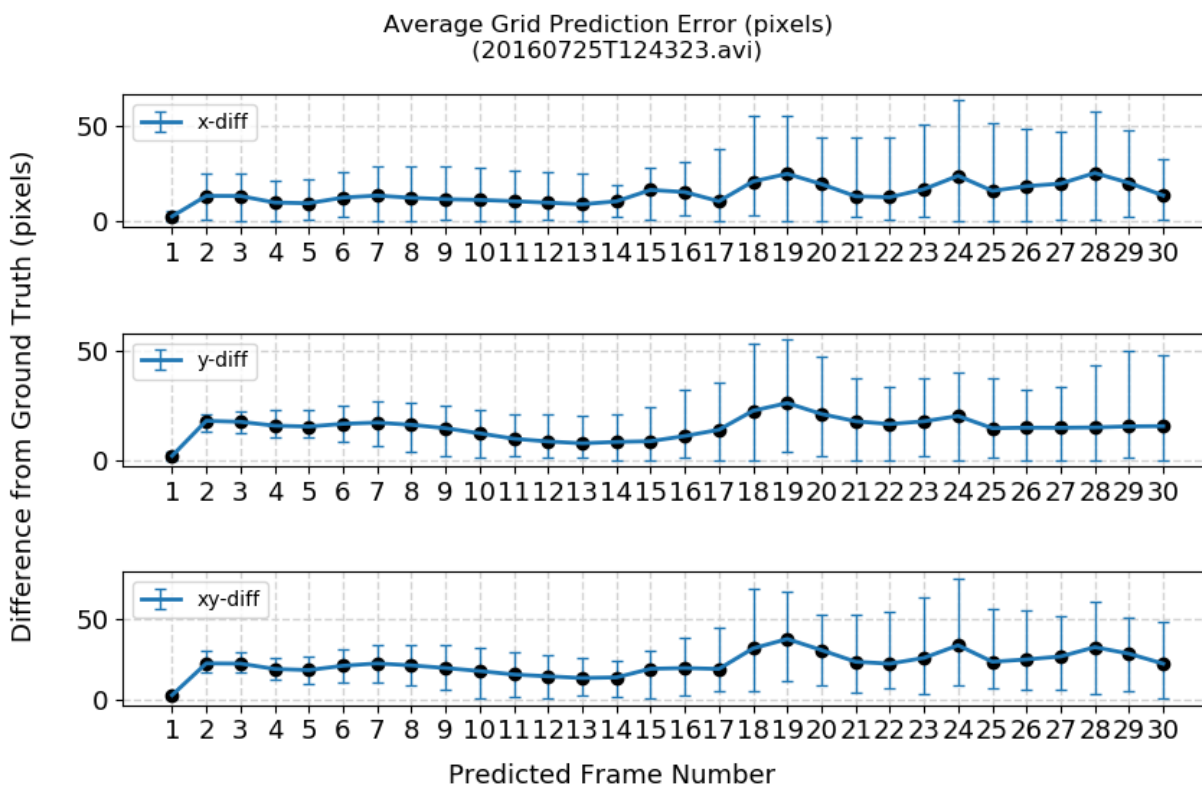


Figure 4.3: Average Grid Prediction Error (x, y, and x-y [absolute]). Error bars show max and min error values. Image is 800 x 800 pixels. Initial ROI is ~160 x 280 pixels. (20160725T124323.avi).

Figure 4.3 shows that the average difference, as well as the spread of the differences, generally increases as the prediction time increases. However, the graph shows that there are periods during which the errors decrease as the cyclic motion changes direction. This is more evident in the video output files. The average error remains below 50 pixels throughout the prediction and remains consistent in the x- and y-directions.

4.2 Results for 20160725T124609.avi

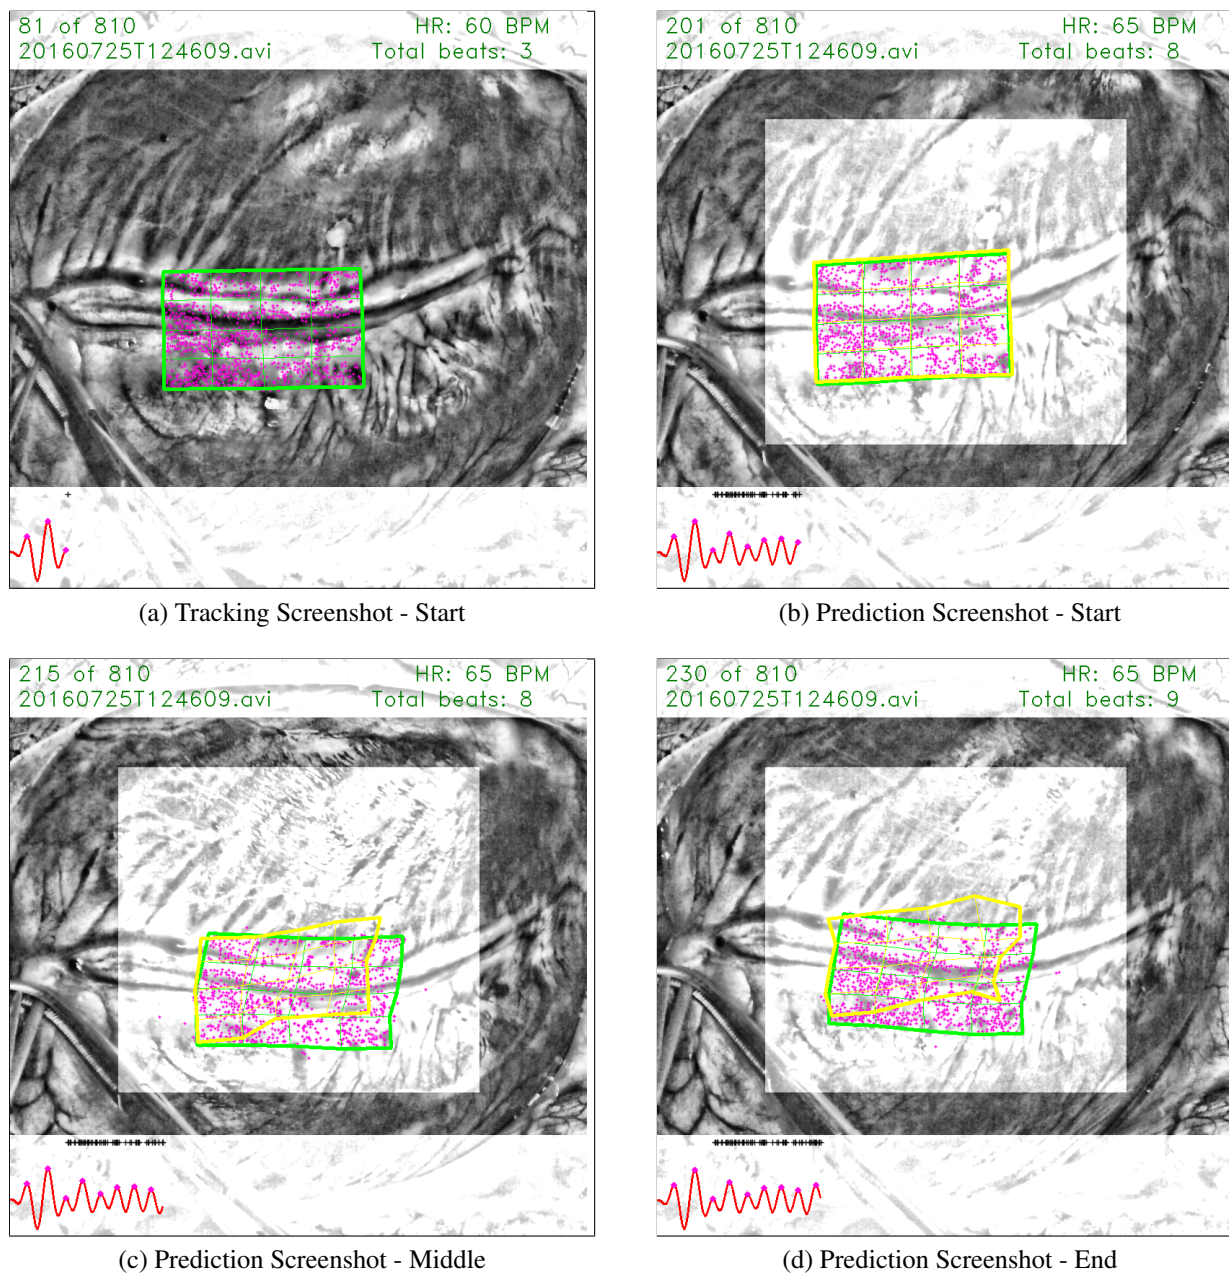


Figure 4.4: Representative Screenshots (20160725T124609.avi).

Figure 4.4 shows that the prediction begins close to the tracked ROI and continues to deviate as the prediction continues. In the middle of the prediction, the predicted grid has moved in a similar

direction as the ground truth, but it shows a large amount of deformation. By the last prediction frame, the predicted grid has a similar amount of deformation as the middle frame.

Table 4.2: Tracking Performance Summary (20160725T124609.avi).

| Parameter | Total | % of Total |
|------------------|-------|------------|
| Total Frames | 810 | - |
| Tracked Frames | 729 | 90.0 % |
| Diastole Peaks | 34 | 4.2 % |
| Tracking Resets | 1 | 0.1 % |
| Flag Raised | 255 | 35.0 % |
| Predicted Frames | 30 | 4.0 % |




Figure 4.5: Analysis Timeline (20160725T124609.avi).

Figure 4.5 shows that there are a large number of frames where the optical flow has not passed both of the tracking constraint criteria (flag raised) but only 1 frame where the tracking was forced to reset.

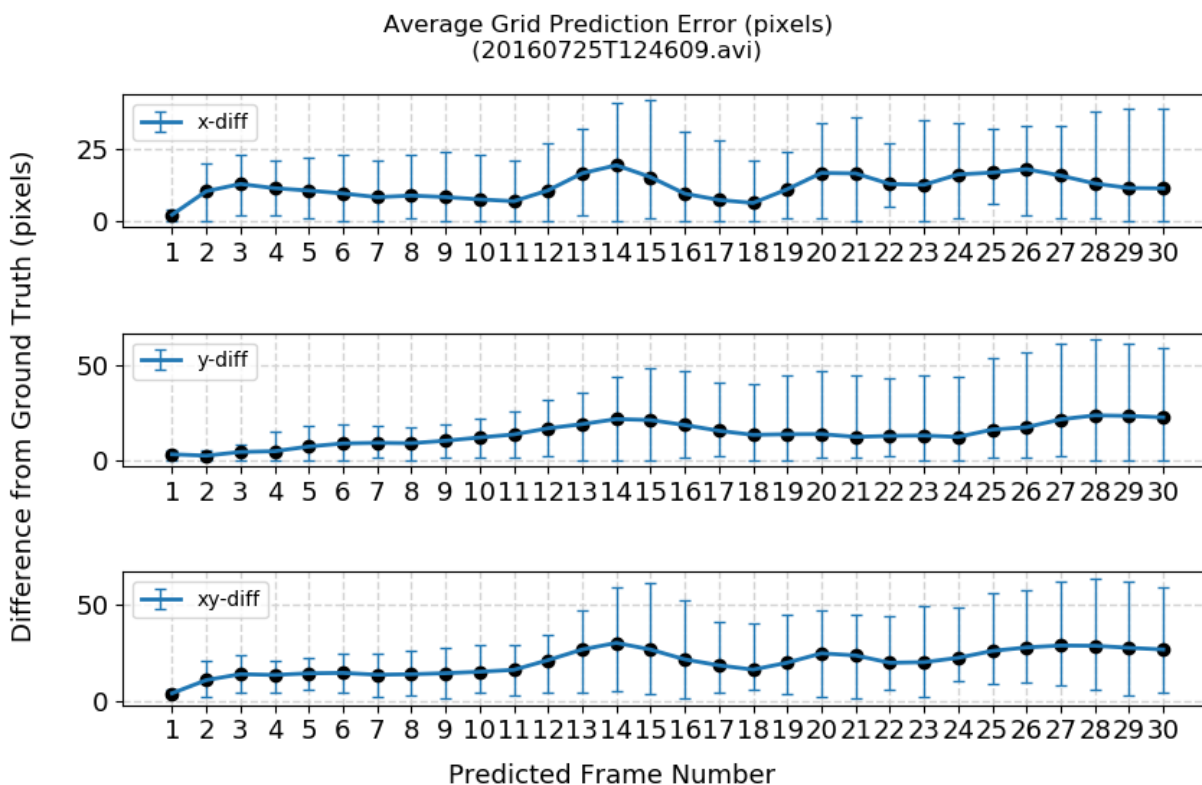


Figure 4.6: Average Grid Prediction Error (x, y, and x-y [absolute]). Error bars show max and min error values. Image is 800 x 800 pixels. Initial ROI is ~160 x 280 pixels. (20160725T124609.avi).

Figure 4.6 shows that the average difference, as well as the spread of the differences, generally increases as the prediction time increases. However, the graph shows that there are periods during which the errors decrease as the cyclic motion changes direction. This is more evident in the video output files. The average error remains below 50 pixels throughout the prediction. The errors are lower in the x-direction than in the y-direction.

4.3 Results for 20160725T124909.avi

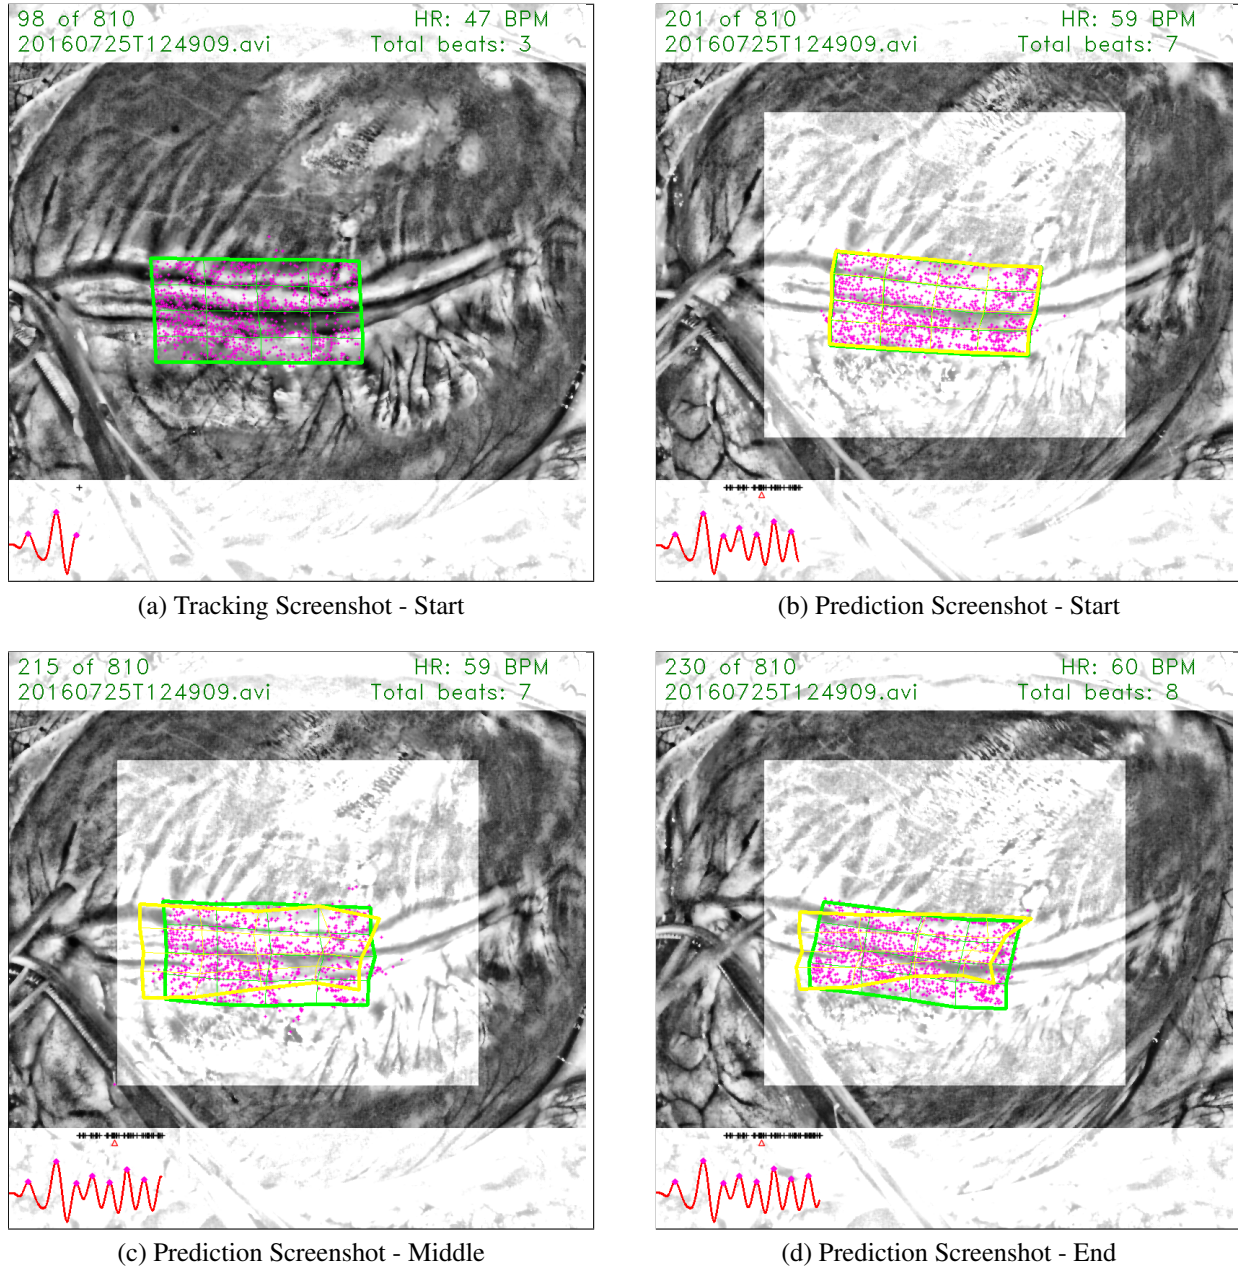


Figure 4.7: Representative Screenshots (20160725T124909.avi).

Figure 4.7 shows that the prediction begins close to the tracked ROI and continues to deviate as the prediction continues. In the middle of the prediction, the predicted grid has moved in a similar

direction as the ground truth with some deformation. By the last prediction frame, the predicted grid error and deformation has further increased.

Table 4.3: Tracking Performance Summary (20160725T124909.avi).

| Parameter | Total | % of Total |
|------------------|-------|------------|
| Total Frames | 810 | - |
| Tracked Frames | 712 | 87.9 % |
| Diastole Peaks | 35 | 4.3 % |
| Tracking Resets | 6 | 0.8 % |
| Flag Raised | 304 | 42.7 % |
| Predicted Frames | 30 | 4.1 % |




Figure 4.8: Analysis Timeline (20160725T124909.avi).

Figure 4.8 shows that there are a large number of frames where the optical flow has not passed both of the tracking constraint criteria (flag raised) but only 6 frames where the tracking was forced to reset.

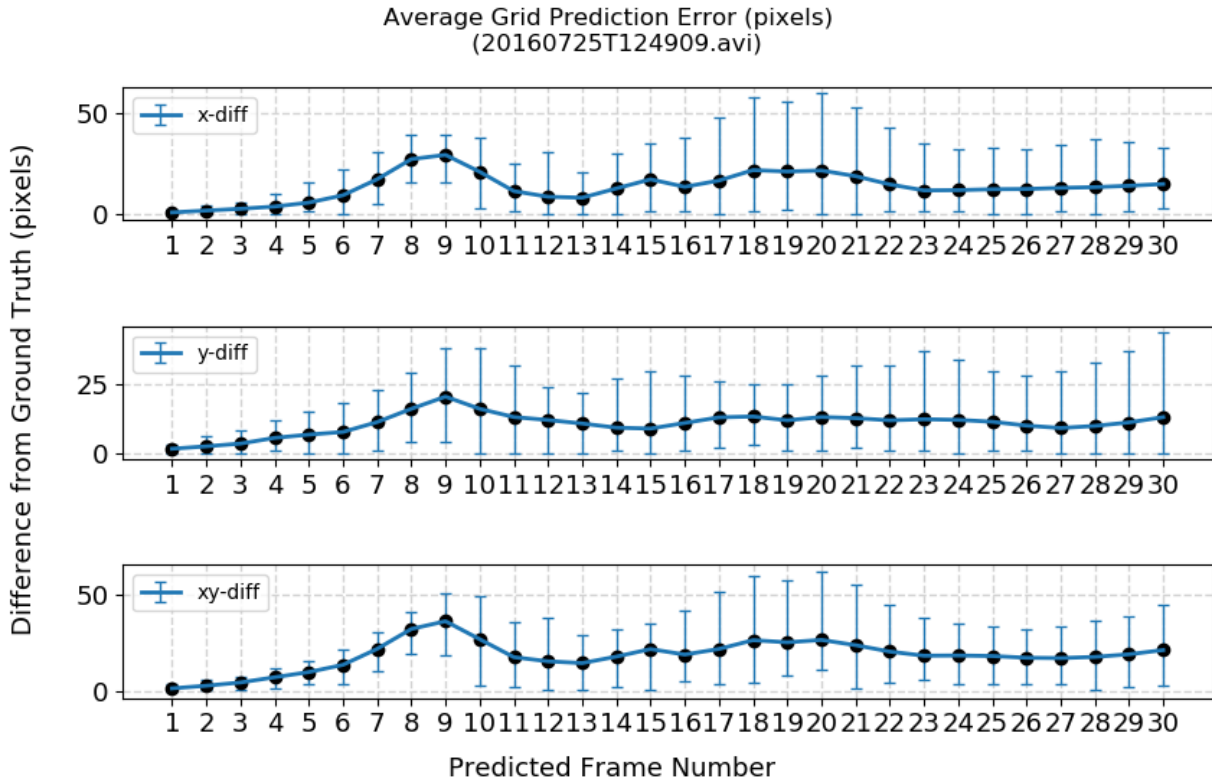


Figure 4.9: Average Grid Prediction Error (x, y, and x-y [absolute]). Error bars show max and min error values. Image is 800 x 800 pixels. Initial ROI is ~160 x 280 pixels. (20160725T124909.avi).

Figure 4.9 shows that the average difference, as well as the spread of the differences, generally increases as the prediction time increases. However, the graph shows that there are periods during which the errors decrease as the cyclic motion changes direction. This is more evident in the video output files. The average error remains below 50 pixels throughout the prediction. The errors are lower in the y-direction than in the x-direction.

4.4 Results for 20160725T132840.avi

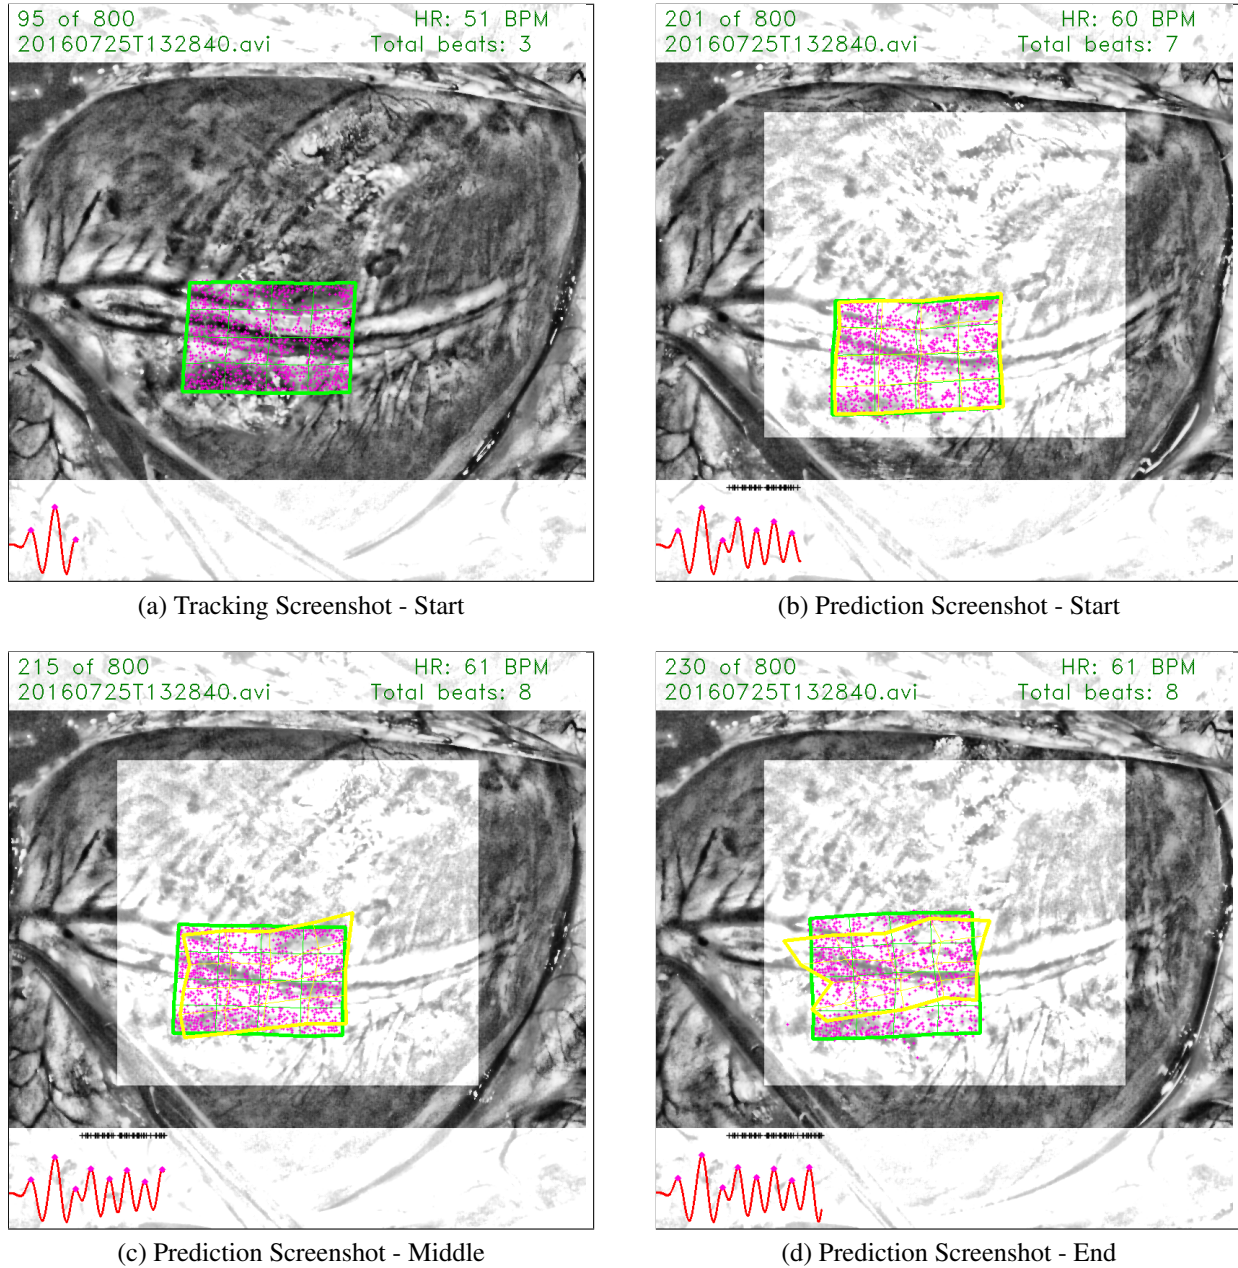


Figure 4.10: Representative Screenshots (20160725T132840.avi)

Figure 4.10 shows that the prediction begins close to the tracked ROI and continues to deviate as the prediction continues. In the middle of the prediction, the predicted grid has moved in a

similar direction as the ground truth and has a very similar shape. By the last prediction frame, the predicted grid error has increased, and the grid shape shows larger deformation.

Table 4.4: Tracking Performance Summary (20160725T132840.avi).

| Parameter | Total | % of Total |
|------------------|-------|------------|
| Total Frames | 810 | - |
| Tracked Frames | 705 | 88.1 % |
| Diastole Peaks | 40 | 5.0 % |
| Tracking Resets | 4 | 0.6 % |
| Flag Raised | 257 | 36.5 % |
| Predicted Frames | 30 | 4.1 % |




Figure 4.11: Analysis Timeline (20160725T132840.avi).

Figure 4.11 shows that there are a large number of frames where the optical flow has not passed both of the tracking constraint criteria (flag raised) but only 4 frames where the tracking was forced to reset.




Figure 4.12: Average Grid Prediction Error (x, y, and x-y [absolute]). Error bars show max and min error values. Image is 800 x 800 pixels. Initial ROI is ~160 x 280 pixels. (20160725T132840.avi).

Figure 4.12 shows that the average difference, as well as the spread of the differences, generally increases as the prediction time increases. However, the graph shows that there are periods during which the errors decrease as the cyclic motion changes direction. This is more evident in the video output files. The average error remains below 50 pixels throughout the prediction. The errors are lower in the y-direction than in the x-direction.

4.5 Results for 20160725T133113.avi

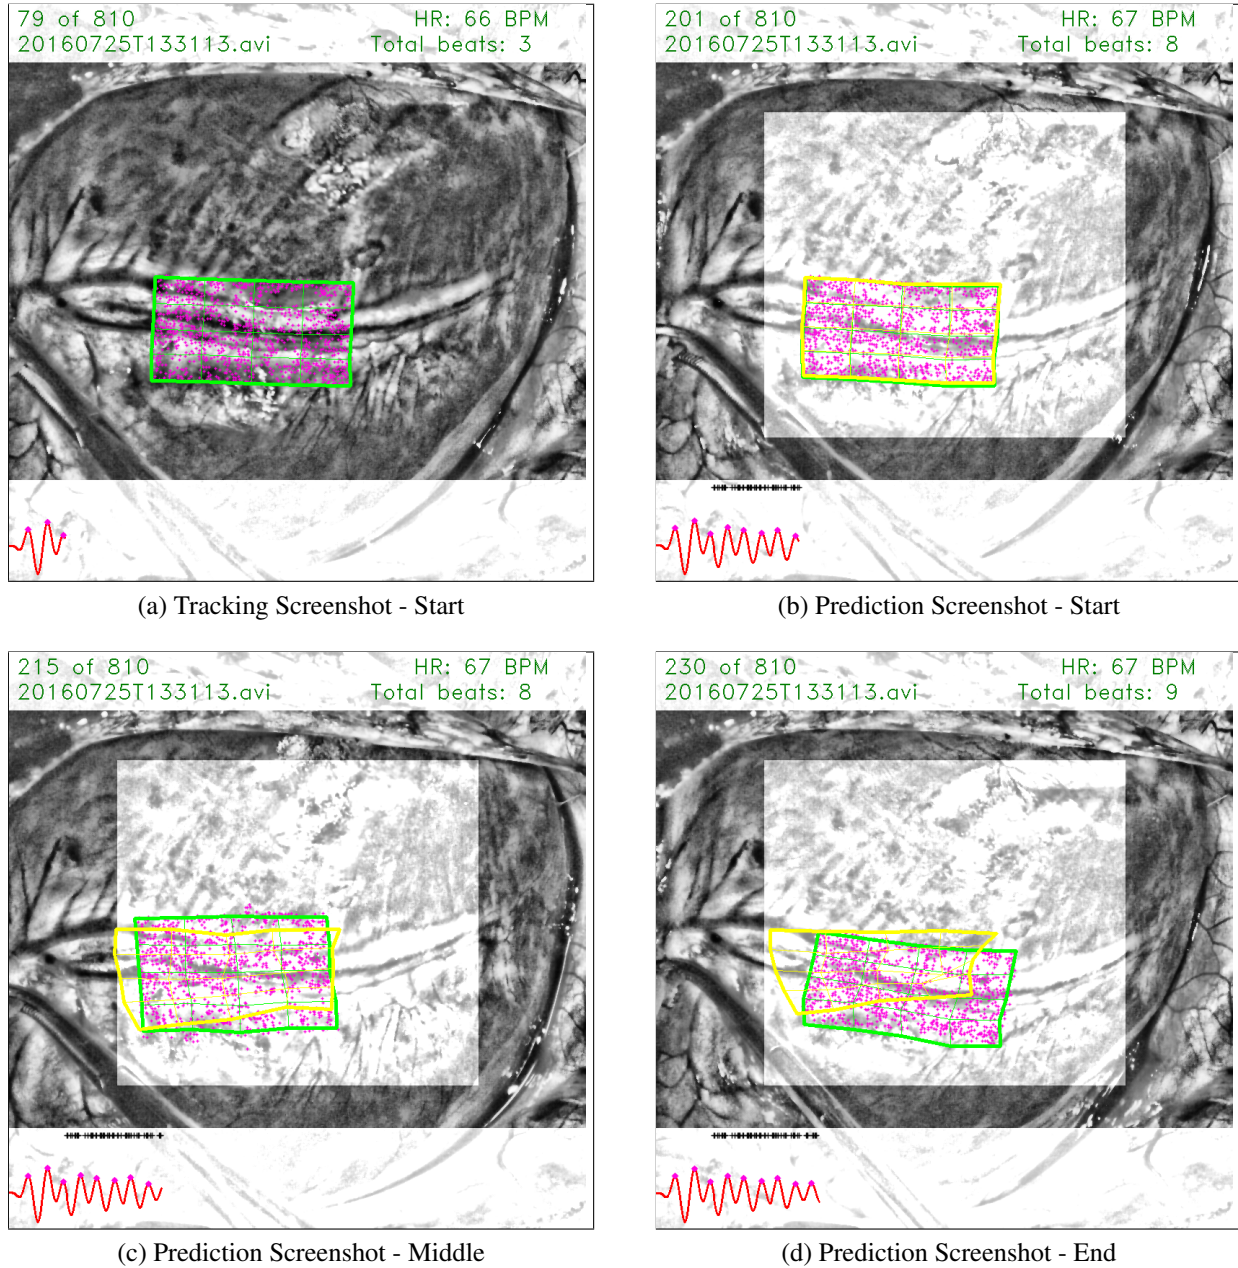


Figure 4.13: Representative Screenshots (20160725T133113.avi).

Figure 4.13 shows that the prediction begins close to the tracked ROI and continues to deviate as the prediction continues. In the middle of the prediction, the predicted grid has moved in a

similar direction as the ground truth and shows a moderate amount of deformation. By the last prediction frame, the predicted grid has a large error and deformation.

Table 4.5: Tracking Performance Summary (20160725T133113.avi).

| Parameter | Total | % of Total |
|------------------|-------|------------|
| Total Frames | 810 | - |
| Tracked Frames | 731 | 90.2 % |
| Diastole Peaks | 33 | 4.1 % |
| Tracking Resets | 6 | 0.8 % |
| Flag Raised | 297 | 40.6 % |
| Predicted Frames | 30 | 4.0 % |




Figure 4.14: Analysis Timeline (20160725T133113.avi).

Figure 4.14 shows that there are a large number of frames where the optical flow has not passed both of the tracking constraint criteria (flag raised) but only 6 frames where the tracking was forced to reset.

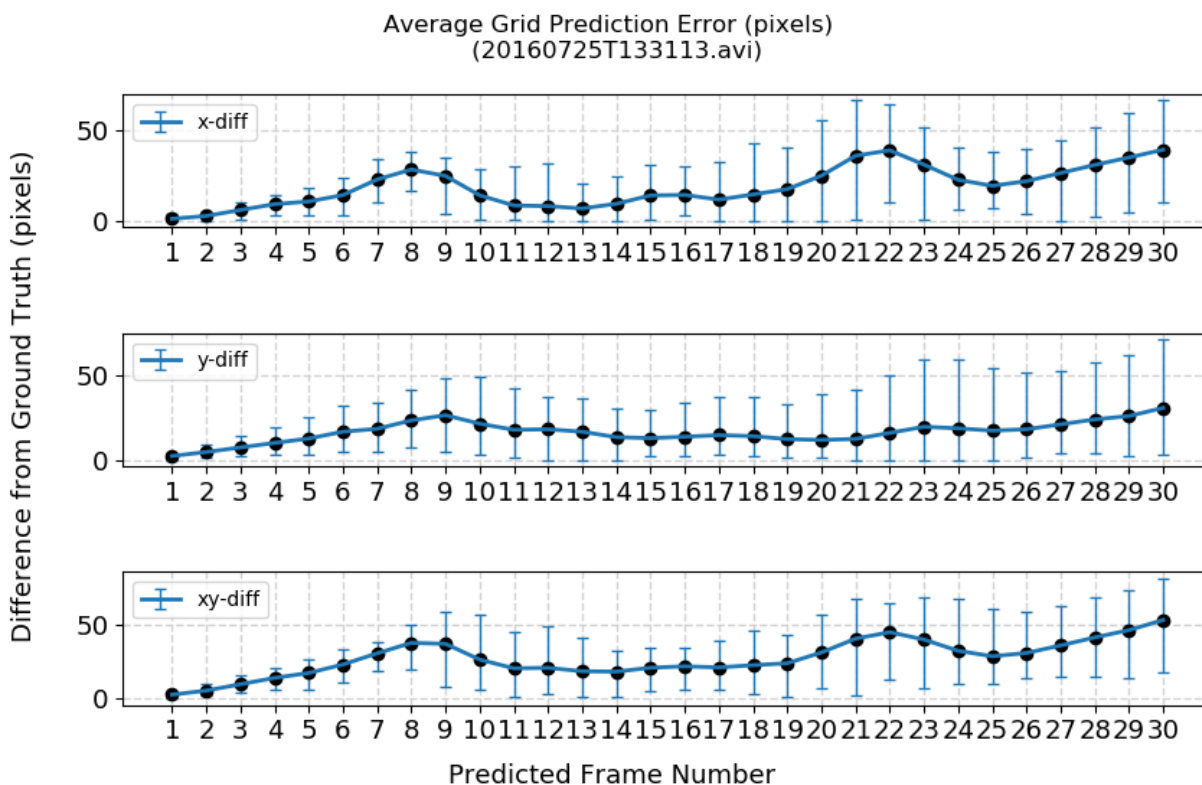


Figure 4.15: Average Grid Prediction Error (x, y, and x-y [absolute]). Error bars show max and min error values. Image is 800 x 800 pixels. Initial ROI is ~160 x 280 pixels. (20160725T133113.avi).

Figure 4.15 shows that the average difference, as well as the spread of the differences, generally increases as the prediction time increases. However, the graph shows that there are periods during which the errors decrease as the cyclic motion changes direction. This is more evident in the video output files. The average error remains below 50 pixels throughout the prediction. The errors in the x- and y-direction are of similar magnitude, but the x-direction shows more cyclical motion.

4.6 Neural Network Training

The neural network was trained for 5 epochs with a batch size of 512. After 5 epochs, the losses on the training and testing data sets were approximately equal as shown in Figure 4.16. The

accuracy of both data sets begins to plateau after 4 epochs as shown in Figure 4.17. Training time was a total of 14 minutes. See Section 3.5.2 for a discussion of the loss and accuracy calculation.




Figure 4.16: Neural Network Training Loss.




Figure 4.17: Neural Network Training Accuracy.

4.7 Speed of Execution

Table 4.6 provides a summary of the run time for each file. The file read/write time was subtracted from the total time and frame rate calculation to give a better estimate of the performance that might be expected from a live camera with better frame buffering (See Section 5.4.4).

Table 4.6: Execution Speed Summary.

| Data Set | Execution Time (s) | Approximate Frame Rate (FPS) |
|---------------------|--------------------|------------------------------|
| 20160725T124323.avi | 55.1 | 13.1 |
| 20160725T124609.avi | 47.0 | 15.5 |
| 20160725T124909.avi | 51.4 | 13.9 |
| 20160725T132840.avi | 49.1 | 14.4 |
| 20160725T133113.avi | 47.2 | 15.5 |
| Average | 50.0 | 14.5 |

CHAPTER 5

DISCUSSION

This chapter discusses the results presented in Chapter 4.

5.1 Tracking Performance

While a direct comparison to the original research is not made (due to the difference in the selected ROI and change in constraint failure limit), the results from Chapter 4 demonstrate that the tracking failure rate is improved in the new implementation. Table 5.1 below gives a summary and comparison of the failure rates between the new and original algorithm.

Table 5.1: Tracking Failure Summary and Comparison.

| Data Set | Tracking Resets | | |
|---------------------|----------------------|----------|-----------------|
| | Original Failures | Diastole | New Failures |
| 20160725T124323.avi | 22.2% | 4.0% | 0.6% |
| 20160725T124609.avi | 16.8% | 4.2% | 0.1% |
| 20160725T124909.avi | 17.2% | 4.3% | 0.8% |
| 20160725T132840.avi | 30.8% | 5.0% | 0.6% |
| 20160725T133113.avi | 11.7% | 4.1% | 0.8% |
| Average | 19.7% | 4.3% | 0.6% |

Table 5.1 demonstrates that the changes to re-initializing tracking at every diastole peak greatly reduces the overall reset rate from almost 20% to 5%. If tracking can be maintained through

multiple diastole peaks without resetting, the reset rate may be further reduced. The equalization, denoising, and failure limit changes drive the resets due to failure close to zero.

5.2 Prediction Performance

The results shown in Chapter 4 demonstrate that the motion of cardiac images contained in the data set was learned by the neural network model. Table 5.2 below gives a summary and comparison of the x-y error across the range of predicted frames.

Table 5.2: Grid Prediction Error Summary.

| Data Set | x-y Errors (pixels) | | | | | | | | |
|---------------------|---------------------|-----|-----|----------|------|------|----------|------|------|
| | Frame 1 | | | Frame 15 | | | Frame 30 | | |
| | Min | Avg | Max | Min | Avg | Max | Min | Avg | Max |
| 20160725T124323.avi | 1.0 | 3.2 | 5.1 | 1.0 | 19.4 | 22.5 | 5.1 | 30.6 | 48.0 |
| 20160725T124609.avi | 2.2 | 3.8 | 5.7 | 3.2 | 26.7 | 61.6 | 4.2 | 26.7 | 59.0 |
| 20160725T124909.avi | 0.0 | 1.7 | 3.6 | 1.0 | 22.0 | 35.4 | 3.0 | 21.5 | 45.1 |
| 20160725T132840.avi | 2.2 | 4.7 | 7.1 | 9.9 | 17.3 | 24.8 | 2.8 | 26.5 | 50.8 |
| 20160725T133113.avi | 0.0 | 2.8 | 4.5 | 5.4 | 20.9 | 34.0 | 18.0 | 53.2 | 81.5 |
| Average | 1.1 | 3.2 | 5.2 | 4.1 | 21.3 | 35.7 | 6.6 | 31.7 | 56.9 |

Table 5.2 shows that the neural network model was able to predict the grid point locations for the overall data set with an average maximum distance error of 56.9 pixels after a 30 frame prediction period. The mean average error for all of the data sets gradually increased throughout the prediction period from 3.2 pixels to 31.7 pixels. The overall maximum error was below 81.5 pixels across the prediction period for all data sets.

The errors quickly accumulate and become large across the prediction period which may limit the practical usefulness of the prediction depending on the application. An average 50 pixel offset could mean the grid point is off by ~15-30% of the grid dimensions. In a precise tumor tracking

application this offset may not be acceptable, but in longer term monitoring of tissue strain, for example, it may be acceptable. However, further training with better data or different network configurations may reduce the errors.

5.3 Speed of Execution

One of the original goals of this research was to improve the implementation of the baseline algorithm so that it could run in real-time on a live camera feed. The original research did not include any information on the algorithm execution time, but, in informal testing, the frame rate was around 5 fps. The results shown in Section 4.7 demonstrate that the new implementation improves the frame rate to around 14.5 fps.

By examining the output of the code profiling tool included with Python, other improvements can be identified. Table 5.3 shows a summary of the most time consuming functions in the algorithm grouped by type for one video file. Results were typical across the data set.

Table 5.3: Profiler Summary for 20160725T124323.avi.

| Function Name | Number of Calls | Cumulative Time (s) | Total Time (s) |
|-------------------------------|---------------------|---------------------|----------------|
| TensorFlow Functions: | | | |
| TFE_Py_execute | 792 | 15.076 | 15.076 |
| TFE_Py_FastPathExecute | 2560 | 1.116 | 1.116 |
| NumPy Functions: | | | |
| implement_array_function | 7085963/ 2741625 | 18.85 | 5.416 |
| append | 1932354 | 10.892 | 1.73 |
| ravel | 2082699 | 3.821 | 1.155 |
| array | 4824238 | 2.408 | 2.404 |
| reduce | 465898 | 1.014 | 1.011 |
| OpenCV Functions: | | | |
| calcOpticalFlowPyrLK | 1592 | 4.561 | 4.561 |
| Write cv2.VideoCapture object | 809 | 4.534 | 4.534 |
| Bilateral Filter | 810 | 2.72 | 2.72 |
| norm | 6472 | 1.787 | 1.787 |
| detectAndCompute | 35 | 1.69 | 1.69 |
| Read cv2.VideoCapture object | 811 | 1.654 | 1.654 |
| Other: | | | |
| update_patch_flow | 783 | 25.359 | 2.428 |

The second column shows the total number of times the function is called. The third column shows the total time for that function along with all of the functions called within that function. The last column shows the total time used within that function only.

Overall, the NumPy functions use the most amount of runtime with millions of calls to functions that manipulate the feature point lists. The "update_patch_flow" function shown on the last row is where many of these operations occur.

The TensorFlow inference operations also take a significant amount of the total run time as inference is performed one point at a time.

The last category of time-consuming functions relate to OpenCV. Two of these functions are used to read from the input video file and write to a new file and may be unnecessary with a live camera feed. The optical flow tracker, feature detection, denoising filter, and norm operation (used for calculating the pulse tracker optical flow) are the most intensive image processing operations.

5.4 Limitations / Future Improvements

Several improvements could be added in future research to improve the tracking algorithm, neural network prediction accuracy, and execution speed.

5.4.1 Tracking Algorithm

Further work could be done to develop additional constraints for the tracking algorithm. Additional criteria could be added to help distinguish non-rigid motion from the motion of occluding objects to trigger a switch from tracking to prediction or detect corner cases where the tracking algorithm does not fail but has an undesired shape (grid lines overlapping, etc.).

The smoothness constraint threshold could be improved by addition of an extreme event detection algorithm using Extreme Value Theory (EVT) which could continuously monitor and update the threshold range dynamically as tracking progresses [59, 60]. The current method assumes that the motion during the second diastole is representative of the rest of the video and may require periodic threshold resets as motion changes over longer periods.

A tracking re-initialization that occurs within the 10-step sequence history just prior to a prediction could cause an abrupt change within the sequence that may introduce additional errors. An adjustment to smooth large re-initialization changes may reduce errors and improve the predictions.

The tracked feature points can sometimes exhibit unstable jumps to other image areas when details are suddenly lost (e.g. glare) or when objects enter the ROI. A more robust method of handling outlier motion could help to screen out these points and improve the tracking performance.

The switch to triggering re-initialization during cardiac diastole peaks required some knowledge of the period and type of tissue motion. Future implementations for other organs may require changes specific to that tissue motion.

5.4.2 Data Collection

As can be seen in Figure 3.7, the extracted data point features tended to cluster over the length of the tracking sequence. Performing a periodic re-initialization of the features to reduce this clustering could result in better training data. However, because these sequences were sliced into relatively small sequences, any long term clustering effects would only minimally affect the final sequences. There could be a bias of more sequences based on the clustered areas, but this might not necessarily be detrimental since the clusters are likely to form in more well-defined areas of the image.

A large number of points (10,000 maximum) were used in the feature detection step in order to obtain more points along the coronary artery. This area did not show many detected features without a high maximum point threshold setting. Further changes to the feature detection routine to improve detection along the artery, divide the initial detection into zones with different detection settings, or eliminate excess feature points in certain regions could reduce the size of the data set. The total number of point sequences obtained from the combined data set after slicing the 30 second time series numbered approximately 28 million sequences. Out of memory errors were encountered during training with this data set so only 40% of the collected sequences (after shuffling) were used.

Valid predictions would be limited to the central area of the image due to the circular feature extraction area chosen. The feature extraction method should be tailored to each specific tissue and data set as necessary.

Finally, all videos in this data set had a similar field of view and similar, regular heart rates. An increase in the variation of angles, heart rates, and beat patterns would allow better prediction in other scenarios. On another data set with different heart rates, prediction accuracy might be reduced due to the difference in distance moved (tissue speed) during each frame. More data could be collected at several heart rates or, assuming the cyclic motion is similar across all heart rates, a scaling algorithm could be developed to scale the predicted output based on the detected heart rate. Collection of data with no foreign objects in the frame would also improve the quality of the data and require less post-processing.

5.4.3 Neural Network

Adopting a different neural network structure may yield improved results by better capturing the spatial relationships between grid points. As described in Section 2.2.3, a CNN network can learn these spatial relationships, but it typically performs operations on entire images or fixed grids unlike the moving, user-selected ROI grids created in this research. However, a CNN-LSTM combination trained on the entire set of image frames could predict regional optical flows and adjust grid point locations [61, 62]. Other recent research into using Convolutional LSTMs or LSTMs with attention mechanisms for analysis of non-uniform sensor data may provide a framework for using the variable grids generated in this research with a CNN [63, 64]. Further exploration of the addition of dropout layers may yield better, more generalized prediction results.

In this research, a simple 80% (training) / 20% (testing) data split was performed after shuffling the entire data set. Thus, the network model accuracy calculated during training is limited to this single data split configuration. A more robust approach to calculating accuracy is to use k-fold cross-validation in which several data splits (folds) are generated and used to train the model

multiple times. This helps to ensure that no single portion of data is over-emphasized in the results. The average accuracy of multiple training runs on a single model is also better suited for use in comparisons to different network models with different hyper-parameters [20]. In this instance, 20% of the data is 1/5 of the data set, hence, a 5-fold cross-validation would train the model 5 times with different 20% test data split.

5.4.4 Speed of Execution

Achieving real-time performance for use during a surgery is an eventual goal for this type of tissue tracking system. While real-time performance was not achieved in this research due to several factors (See Section 5.3), several improvements can be added that would significantly increase performance. Simply adding additional computing power (faster CPUs, GPUs) may provide improved performance (and additional hardware cost), but optimization of the programming routines could provide a similar or greater performance boost while allowing the algorithm to run on less powerful hardware.

Various programming improvements could increase the speed of calculations. Loops should be re-examined and optimized to take advantage of NumPy vector operations when possible as vectorization can speed up array operations by orders of magnitude [65]. Additional performance may be possible on large data sets by using CyPy, a CUDA accelerated version of NumPy [66, 67]. Conversion of double-precision values to single-precision may improve performance if the reduction in precision can be tolerated [31]. Image processing functions should generally perform better on GPU hardware, but care must be taken to balance the performance decrease due to the transfer of data to and from the GPU with the performance improvement of a specific routine [31]. However, various OpenCV functions may or may not necessarily always benefit from being executed on a GPU [68]. Parallel threads to perform the grid point prediction may greatly improve the speed of inference and would be necessary for a real-time grid prediction.

One bottleneck identified in Section 5.3 is in the file reading and writing routines. This is consistent with the experience of other software developers and is a known issue when reading image frames from video files [69]. Similar issues may be encountered when quickly reading frames from a video camera [70]. Frame buffering and multi-threaded operations are recommended to improve the input/output performance in both cases [69, 70].

A final way to dramatically increase speed is to simply reduce the frame rate to a lower value. For example, a frame rate reduction of 30 fps to 20 fps should reduce the overall amount of computation by about $1/3$. The frame rate would need to be high enough that each motion step does not increase to a level that causes tracking to fail and the neural network would need to be retrained on data adjusted to that frame rate.

CHAPTER 6

CONCLUSIONS

A method was successfully developed to predict a non-rigid tracking grid location from a neural network model that has learned similar types of motion. This method was applied to a cardiac video data set and integrated with an existing non-rigid tissue tracking algorithm.

Several improvements were made to increase the performance of the baseline algorithm. A complete rewrite from Matlab into Python code was finished to simplify the code and increase readability. This rewrite also allowed the use of industry-standard, open-source software libraries.

A pulse tracking routine based on optical flow was added to aid in feature detection and re-initialization of tracking. This gives the additional benefit of adding an estimate of heart rate to the algorithm. If developed further, this optical flow measurement could be used as an additional diagnostic parameter during surgery to monitor overall tissue strain as described in Chapter 1.

The final speed of execution was about one-half of that required for a real-time system. The performance bottlenecks identified in Section 5.4.4 have great potential for optimization suggesting that real-time performance is achievable on the current generation of high-end desktop/laptop computers.

Other future improvements identified in Section 5.4 include developing additional tracking constraints, improving occlusion detection, improving the quality of the data collection, and experimenting with more neural network configurations.

Finally, this technique need not be limited to cardiac images. It could be applied to other types of non-rigid motion tracking outside of biological tissues such as ocean surface tracking for weather and energy research or fabric motion tracking for augmented reality applications.

REFERENCES

- [1] Intuitive Surgical Inc. Intuitive surgical, inc. - annual report 2019, 2019. <https://investor.intuitivesurgical.com/static-files/31b5c428-1d95-4c01-9c85-a7293bac5e05>, Accessed 05/21/2020.
- [2] Koichiro Yamakado. Image-guided ablation of adrenal lesions. *Seminars in Interventional Radiology*, 31(2):149–156, Jun 2014.
- [3] Nebojša Mujović, Milan Marinković, Radoslaw Lenarczyk, Roland Tilz, and Tatjana Potpara. Catheter ablation of atrial fibrillation: An overview for clinicians. *Advances in Therapy*, 34(8):1897–1917, Aug 2017.
- [4] Marlène C. H. Hekman, Mark Rijpkema, Johan F. Langenhuijsen, Otto C. Boerman, Egbert Oosterwijk, and Peter F. A. Mulders. Intraoperative imaging techniques to support complete tumor resection in partial nephrectomy. *European Urology Focus*, 4(6):960–968, 2017.
- [5] Michele Solis. New frontiers in robotic surgery: The latest high-tech surgical tools allow for superhuman sensing and more. *IEEE Pulse*, 7(6):51–55, Nov 2016.
- [6] O. Lorton, P. C. Guillemain, N. Mori, L. A. Crowe, S. Boudabbous, S. Terraz, C. D. Becker, P. Cattin, R. Salomir, and L. Gui. Self-scanned hifu ablation of moving tissue using real-time hybrid us-mr imaging. *IEEE transactions on bio-medical engineering*, 66(8):2182–2191, August 01 2019.
- [7] T. Bader, A. Wiedemann, K. Roberts, and U. D. Hanebeck. Model-based motion estimation of elastic surfaces for minimally invasive cardiac surgery. pages 2261–2266, 2007.
- [8] A. Soltani, J. Lahti, K. Jarvela, S. Curtze, J. Laurikka, M. Hokka, and V. T. Kuokkala. An optical method for the in-vivo characterization of the biomechanical response of the right ventricle. *Scientific reports*, 8(1):6831–z, May 01 2018.
- [9] Archie Hughes-Hallett, Erik K. Mayer, Hani J. Marcus, Thomas P. Cundy, Philip J. Pratt, Ara W. Darzi, and Justin A. Vale. Augmented reality partial nephrectomy: Examining the current status and future perspectives. *Urology*, 83(2):266–273, 2014.

- [10] T. R. McCarty and T. Rustagi. New indications for endoscopic radiofrequency ablation. *Clinical gastroenterology and hepatology : the official clinical practice journal of the American Gastroenterological Association*, 16(7):1007–1017, July 01 2018.
- [11] R. Richa, A. P. Bo, and P. Poignet. Towards robust 3d visual tracking for motion compensation in beating heart surgery. *Medical image analysis*, 15(3):302–315, June 01 2011.
- [12] G. R. Hale, F. Pesapane, S. Xu, I. Bakhutashvili, N. Glossop, B. Turkbey, P. A. Pinto, and B. J. Wood. Tracked foley catheter for motion compensation during fusion image-guided prostate procedures: a phantom study. *European radiology experimental*, 4(1):24–4, April 16 2020.
- [13] A. Soltani, J. Lahti, K. Jarvela, J. Laurikka, V. T. Kuokkala, and M. Hokka. Characterization of the anisotropic deformation of the right ventricle during open heart surgery. *Computer methods in biomechanics and biomedical engineering*, 23(3):103–113, February 01 2020.
- [14] Bryant Tucker. Development of a heart motion tracking system using non-invasive imaging data. Master’s thesis, 2017. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works.; Last updated - 2017-11-07.
- [15] Jianbo Shi and Tomasi. Good features to track. In *1994 Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 593–600, 1994.
- [16] Basic concepts of the homography explained with code. https://docs.opencv.org/master/d9/dab/tutorial_homography.html, Accessed 05/21/2020.
- [17] A. Sharaff and M. Choudhary. Comparative analysis of various stock prediction techniques. In *2018 2nd International Conference on Trends in Electronics and Informatics (ICOEI)*, pages 735–738, 2018.
- [18] N. Muralidhar, S. Muthiah, K. Nakayama, R. Sharma, and N. Ramakrishnan. Multivariate long-term state forecasting in cyber-physical systems: A sequence to sequence approach. In *2019 IEEE International Conference on Big Data (Big Data)*, pages 543–552, 2019.
- [19] Peter J. Brockwell and Richard A. Davis. *Introduction to time series and forecasting*. Springer, New York, 2002.
- [20] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [21] J. Brownlee. *Deep Learning With Python: Develop Deep Learning Models on Theano and TensorFlow Using Keras*. Machine Learning Mastery, 2017. v1.8.

- [22] Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *CoRR*, abs/1406.1078, 2014.
- [23] R. Yamashita, M. Nishio, R. K. G. Do, and K. Togashi. Convolutional neural networks: an overview and application in radiology. *Insights into imaging*, Jun 22, 2018.
- [24] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1, NIPS'12*, page 1097–1105, Red Hook, NY, USA, 2012. Curran Associates Inc.
- [25] M. Amaresh and S. Chitrakala. Video captioning using deep learning: An overview of methods, datasets and metrics. In *2019 International Conference on Communication and Signal Processing (ICCSPP)*, pages 0656–0661, 2019.
- [26] L. Zhang and P. N. Suganthan. Visual tracking with convolutional neural network. In *2015 IEEE International Conference on Systems, Man, and Cybernetics*, pages 2072–2077, 2015.
- [27] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis, 2018.
- [28] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [29] Choosing mex applications - matlab. https://www.mathworks.com/help/matlab/matlab_external/choosing-mex-applications.html, Accessed 05/21/2020.
- [30] Bill Chou. Implement deep learning applications for nvidia gpus with gpu coder. <https://www.mathworks.com/videos/implement-deep-learning-applications-for-nvidia-gpus-with-gpu-coder-1512748950189.html>, Accessed 05/21/2020.
- [31] Mathworks. Measure and improve gpu performance. <https://www.mathworks.com/help/parallel-computing/measure-and-improve-gpu-performance.html>, Accessed 05/21/2020.
- [32] Python. <https://www.python.org/>.
- [33] Tom Radcliffe. Python vs. c/c++ in embedded systems, 08/29/16 08/29/16. <https://opensource.com/life/16/8/python-vs-cc-embedded-systems>, Accessed 05/21/2020.
- [34] Cython. <https://www.cython.org>.

- [35] Numpy. <https://numpy.org/>.
- [36] Scipy. <https://www.scipy.org/>.
- [37] Scikit-learn. <https://scikit-learn.org/>.
- [38] Matplotlib. <https://matplotlib.org/>.
- [39] Opencv: About. <https://opencv.org/about/>, Accessed 05/21/2020.
- [40] Cuda zone | nvidia developer. <https://developer.nvidia.com/cuda-zone>, Accessed 05/21/2020.
- [41] Opencv: Opencv. <https://opencv.org/opencv/>, Accessed 05/21/2020.
- [42] Why tensorflow. <https://www.tensorflow.org/about>, Accessed 05/21/2020.
- [43] Mathworks. Enhance contrast using histogram equalization - matlab histeq. <https://www.mathworks.com/help/images/ref/histeq.html>, Accessed 05/21/2020.
- [44] OpenCV. Histograms - 2: Histogram equalization. https://docs.opencv.org/master/d5/daf/tutorial_py_histogram_equalization.html, Accessed 05/21/2020.
- [45] Opencv: Smoothing images. https://docs.opencv.org/master/dc/dd3/tutorial_gaussian_median_blur_bilateral_filter.html, Accessed 05/21/2020.
- [46] Mathworks. Detect corners using minimum eigenvalue algorithm and return cornerpoints object - matlab detectmineigenfeatures. <https://www.mathworks.com/help/vision/ref/detectmineigenfeatures.html>, Accessed 05/21/2020.
- [47] OpenCV. Shi tomasi corner detector and good features to track. https://docs.opencv.org/master/d4/d8c/tutorial_py_shi_tomasi.html, Accessed 05/21/2020.
- [48] Mathworks. Track points in video using kanade-lucas-tomasi (klt) algorithm - matlab. <https://www.mathworks.com/help/vision/ref/vision.pointtracker-system-object.html>, Accessed 05/21/2020.
- [49] Bruce D. Lucas and Takeo Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence - Volume 2, IJCAI'81*, page 674–679, San Francisco, CA, USA, 1981. Morgan Kaufmann Publishers Inc.
- [50] OpenCV. Opencv: Object tracking. https://docs.opencv.org/master/dc/d6b/group__video__track.html#ga473e4b886d0bcc6b65831eb88ed93323, Accessed 05/21/2020.

- [51] OpenCV. Opencv: Optical flow. https://docs.opencv.org/master/d4/dee/tutorial_optical_flow.html, Accessed 05/21/2020.
- [52] Jason Brownlee. Time series forecasting with the long short-term memory network in python, April 7, 2017. <https://machinelearningmastery.com/time-series-forecasting-long-short-term-memory-network-python/>, Updated August 5, 2019, Accessed 05/21/2020.
- [53] Stacey Ronaghan. Deep learning: Which loss and activation functions should i use?, July 26, 2018. <https://towardsdatascience.com/deep-learning-which-loss-and-activation-functions-should-i-use-ac02f1c56aa8>, Accessed 05/21/2020.
- [54] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1):1929–1958, January 2014.
- [55] Jason Brownlee. How to diagnose overfitting and underfitting of lstm models, September 1, 2017. <https://machinelearningmastery.com/diagnose-overfitting-underfitting-lstm-models/>, Updated on January 8, 2020, Accessed 05/21/2020.
- [56] Jason Brownlee. Time series forecasting performance measures with python, February 1, 2017. <https://machinelearningmastery.com/time-series-forecasting-performance-measures-with-python/>, Updated on August 21, 2019, Accessed 05/21/2020.
- [57] John Pomerat, Aviv Segev, and Rituparna Datta. On neural network activation functions and optimizers in relation to polynomial regression. pages 6183–6185, 12 2019.
- [58] Rochak Agrawal. Optimization algorithms for deep learning, July 23, 2019. <https://medium.com/analytics-vidhya/optimization-algorithms-for-deep-learning-1f1a2bd4c46b>, Accessed 05/21/2020.
- [59] Daizong Ding, Mi Zhang, Xudong Pan, Min Yang, and Xiangnan He. Modeling extreme events in time series prediction. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, page 1114–1122, 2019.
- [60] Alban Siffer, Pierre-Alain Fouque, Alexandre Termier, and Christine Largouet. Anomaly detection in streams with extreme value theory. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, page 1067–1075, 2017.
- [61] Pengpeng Liu, Michael Lyu, Irwin King, and Jia Xu. Selfflow: Self-supervised learning of optical flow, 2019.

- [62] A. M. Fiorito, A. Østvik, E. Smistad, S. Leclerc, O. Bernard, and L. Lovstakken. Detection of cardiac events in echocardiography using 3d convolutional recurrent neural networks. In *2018 IEEE International Ultrasonics Symposium (IUS)*, pages 1–4, 2018.
- [63] Chaoyun Zhang, Marco Fiore, Iain Murray, and Paul Patras. Cloudlstm: A recurrent neural model for spatiotemporal point-cloud stream forecasting, 2019.
- [64] Yuxuan Liang, Songyu Ke, Junbo Zhang, Xiuwen Yi, and Yu Zheng. Geoman: Multi-level attention networks for geo-sensory time series prediction. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 3428–3434. International Joint Conferences on Artificial Intelligence Organization, 7 2018.
- [65] George Seif. One simple trick for speeding up your python code with numpy, Jun 5, 2019. <https://towardsdatascience.com/one-simple-trick-for-speeding-up-your-python-code-with-numpy-1afc846db418>, Accessed 05/21/2020.
- [66] Ryosuke Okuta, Yuya Unno, Daisuke Nishino, Shohei Hido, and Crissman Loomis. Cupy: A numpy-compatible library for nvidia gpu calculations. In *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*, 2017.
- [67] George Seif. Here’s how to use cupy to make numpy over 10x faster, August 22, 2019. <https://towardsdatascience.com/heres-how-to-use-cupy-to-make-numpy-700x-faster-4b920dda1f56>, Accessed 05/21/2020.
- [68] James Bowley. Opencv 3.4 gpu cuda performance comparison (nvidia vs intel), February 28, 2018. <https://jamesbowley.co.uk/opencv-3-4-gpu-cuda-performance-comparison-nvidia-vs-intel/>, Accessed 05/21/2020.
- [69] Adrian Rosebrock. Faster video file fps with cv2.videocapture and opencv, February 6, 2017. <https://www.pyimagesearch.com/2017/02/06/faster-video-file-fps-with-cv2-videocapture-and-opencv/>, Accessed 05/21/2020.
- [70] Najam R Syed. Multithreading with opencv-python to improve video processing performance, July 5, 2018. <https://nrsyed.com/2018/07/05/multithreading-with-opencv-python-to-improve-video-processing-performance>, Accessed 05/21/2020.

APPENDICES

APPENDIX A
SOURCE CODE

A.1 Main Algorithm

The following code is a re-factored version of the original MATLAB algorithm [14] into Python with the various improvements described in Chapter 3.

A.1.1 Overview

Algorithm 1: Main Tracking.

Result: Track a region-of-interest through an image sequence

```
begin
  pre-load neural network model and scaler
  open data set input video
  while new frame available do
    pre-process frame (equalize histogram, de-noise)
    if first frame then
      begin pulse tracker
    else if first diastole detected then
      prompt for ROI selection determine number of patches
    else if second diastole detected then
      set neighbor flow difference thresholds
    else if normal tracking is started (third diastole detected) then
      update optical flow
      TODO add criteria to identify occlusion and switch to prediction
      if neural net prediction frames active then
        run neural net grid point prediction
      else
        check constraints
        re-initialize tracking on each diastole or constraint failure
      end
    record grid point history sequences
    write and display current frame with overlays
  end
end
  close data set input video
  close tracking output video
end
```

A.1.2 Code

```
1 # =====
2 #
3 # file:    main.py
4 # version: v1.0.0
5 # author:  Billy Hempstead
6 #
7 # summary: The purpose of this program is to loop through an entire data set
8 #          to track the boundary edges of a selected ROI, fix failed estimate,
9 #          and quantify current track estimate.
10 #
11 #         Piecewise tracking algorithm loop for data set with ML tracking
12 #         improvements
13 #
14 #         Based on the code from the thesis work:
15 #         "Development of a Heart Motion Tracking System using Non-invasive
16 #         Imaging Data" by Bryent Tucker, July, 2017
17 #
18 # =====
19
20 # Initialization =====
21 import cv2
22 import tensorflow as tf
23 import numpy as np
24 import matplotlib.pyplot as plt
25 import os, time, math
26 import functions as f
27 import classes as c
28 import cProfile, io, pstats
29 from scipy import stats as spstats
30 from sklearn.externals import joblib
31 from tkinter import filedialog
32 from tkinter import *
33
34 # Global Variables =====
35
36 # algorithm selection
37 # the 'Canny' and 'approximate' methods are not supported on a GPU.
38 edge_method = "Canny"
39 edge_threshold = (0.025, 0.3)
40
```

```

41 # plot options
42 edges_on = True
43 patches_on = True
44 points_on = False
45 ROI_on = True
46
47 # tracking reset parameters
48 tracking_enabled = True
49 n_patches_max = 36
50 n_patch_rows = 3
51 n_patches_found = False
52 ML_enabled = True
53 pulse_detection_enabled = False
54 min_corners_in_patch = 8
55 sufficient_corners = False
56 last_peak = 0
57
58 # pulse tracking parameters
59 pulse_max_corners = 10
60 pulse_quality_level = 0.5
61 pulse_min_distance = 25
62 pulse_block_size = 3
63 pulse_use_harris = 0
64 pulse_k = 0.04
65
66 # optical flow settings
67 OF_win_size = (15, 15)
68 OF_max_level = 2
69
70 # colors
71 colors = c.Colors()
72
73 # drawing options
74 show_contours = False
75
76 # Main Loop =====
77
78 # start profiler
79 pr = cProfile.Profile()
80 pr.enable()
81
82 # get filename

```



```

83 root = Tk()
84 root.fullpath = filedialog.askopenfilename(initialdir="./data",
85                                             filetypes=(("AVI File", "*.avi"),("All Files", "*.*")),
86                                             title = "Choose a file.")
87 fullpath = root.fullpath
88 root.destroy()
89
90 pathname = os.path.dirname(fullpath)
91 input_file = os.path.basename(fullpath)
92 basename = os.path.splitext(input_file)[0]
93 output_path = "./output/"
94 output_file = output_path + basename + "_output.avi"
95 text_output_file = output_path + basename + "_summary.txt"
96 ML_model_path = "./models/sequence_model.h5"
97 scaler_model_path = "./models/scaler_model.pkl"
98
99 # initialize local variables
100 corners = []
101 corners_pulse = []
102 peaks = []
103 flow_avg = []
104 contours = []
105 hierarchy = []
106 diastole_frames = []
107 new_peak = False
108 main_tracking_start = False
109 main_tracking_start_frame = -1
110 peak_count = 0
111 screenshot_count = 0
112 t_read = 0
113
114 # filtering parameters
115 lowcut = 30/60 # Hz
116 highcut = 70/60 # Hz
117
118 # initialize constraints object
119 constraints = c.Constraints()
120
121 # set length of sequence for prediction
122 sequence_history_length = 10
123
124 # pre-load neural network

```

```

125 if ML_enabled == True:
126     model = tf.keras.models.load_model(ML_model_path)
127     scaler = joblib.load(scaler_model_path)
128
129 # load input / initialize output video
130 cap, output_vid, frame_width, frame_height, n_last_frame, FPS = \
131     f.import_video(fullpath, output_file)
132
133 # create region of interest rectangle and mask
134 ROI = c.Rect(100, 100, 200, 200) # (x, y, w, h)
135
136 # define area of occlusion and frame range
137 occluded_area = c.Rect(150, 150, 500, 450)
138 occluded_frames = (201, 230)
139
140 # define frame range for analysis
141 frame_range = ( 0, n_last_frame )
142 n_frames = frame_range[1] - frame_range[0]
143 cap.set(cv2.CAP_PROP_POS_FRAMES, frame_range[0])
144
145 # loop through frames
146 while(cap.isOpened()):
147
148     # get current frame
149     frame_count = int(cap.get(cv2.CAP_PROP_POS_FRAMES))
150
151     # update progress
152     # os.system('clear') # slows down execution a lot
153     print('Processing frame: ', int(frame_count), ' of ', int(n_last_frame))
154
155     # exit loop early if desired
156     if frame_count > frame_range[1]:
157         break
158
159     # get a new frame from video
160     t1_read = time.time()
161     ret, frame = cap.read()
162     t2_read = time.time()
163     t_read = t_read+(t2_read-t1_read)
164
165     if ret == True:
166

```

```

167     # equalize and de-noise frame
168     frame_gray = f.preprocess_frame(frame, 'adaptive', 'bilateral')
169     frame_gray_original = np.copy(frame_gray)
170
171     if main_tracking_start == False:
172         if frame_count == frame_range[0]:
173
174             # initialize corners for pulse tracking on first frame
175             corners_pulse = cv2.goodFeaturesToTrack(frame_gray,
176                                                     maxCorners = pulse_max_corners,
177                                                     qualityLevel = pulse_quality_level,
178                                                     minDistance = pulse_min_distance, mask = None,
179                                                     blockSize = pulse_block_size,
180                                                     useHarrisDetector = pulse_use_harris, k = pulse_k )
181
182         else:
183
184             # update pulse detection tracker
185             hr, flow_avg, flow_avg_scaled, peaks, new_peak, corners_pulse =
186                 f.detect_pulse(frame_gray, frame_gray_old, corners_pulse,
187                               peaks, flow_avg, lowcut, highcut, FPS,
188                               frame_count, n_frames)
189
190             # determine number of patches based on first diastole
191             if len(peaks) == 1:
192
193                 # save reference frame and SURF features for reinitialization
194                 if 'frame_ref' not in locals():
195                     frame_ref = np.copy(frame_gray)
196                     detector = cv2.xfeatures2d.SURF_create( hessianThreshold =
197                                                             1000, nOctaves = 4,
198                                                             nOctaveLayers = 3,
199                                                             extended = True,
200                                                             upright = True)
201                     pts_ref, dsc_ref = detector.detectAndCompute(frame_ref, None)
202
203                 t1_roi = time.time()
204                 # select tracking ROI
205                 ROI = f.ROI_selection(frame_gray, ROI)
206                 t2_roi = time.time()
207                 t_roi = t2_roi - t1_roi

```

```

205         if n_patches_found == False:
206             if 'frame_diastole_1' not in locals():
207                 # detect corner points for use in tracker on first peak
208                 frame_diastole_1 = frame_gray_old
209                 diastole_frames.append(frame_count)
210
211             patches, n_patch_rows, n_patches, n_patches_found =
                f.determine_n_patches(frame_diastole_1, frame_ref, ROI,
                n_patch_rows, n_patches_max )
212
213             if n_patches_found == True:
214                 last_peak = len(peaks)
215                 x_grid_ref, y_grid_ref = f.create_grid(patches)
216                 x_grid_ML = np.zeros_like(x_grid_ref)
217                 y_grid_ML = np.zeros_like(y_grid_ref)
218                 x_grid_hist = [ [] for _ in range(x_grid_ref.size)]
219                 y_grid_hist = [ [] for _ in range(y_grid_ref.size)]
220                 x_grid_ML_hist = [ [] for _ in range(x_grid_ref.size)]
221                 y_grid_ML_hist = [ [] for _ in range(y_grid_ref.size)]
222
223             # set neighbor flow threshold based on cycle after next diastole
224             if len(peaks) > last_peak and n_patches_found == True:
225
226                 if 'frame_diastole_2' not in locals():
227                     # reinitialize on peak
228                     diastole_frames.append(frame_count)
229                     frame_diastole_2 = frame_gray_old
230                     patches, corners, _, _, _ =
                        f.reinitialize_tracking(frame_ref, frame_gray,
                        frame_gray_old, patches, x_grid_ref, y_grid_ref,
                        pts_ref, dsc_ref)
231
232                 # set neighbor flow thresholds
233                 patches = f.set_flow_thresholds(frame_gray, frame_gray_old,
                corners, patches)
234
235             # start main tracker on next diastole
236             if len(peaks) > last_peak + 1:
237
238                 main_tracking_start = True
239                 main_tracking_start_frame = frame_count+1
240                 diastole_frames.append(frame_count)

```

```

241
242     # initialize timers
243     t0_sec = time.time()
244     t0 = time.ctime()
245     print(t0)
246     t_write = 0
247     t_read = 0
248     frame_start_tracking = frame_count
249
250     patches, corners, _, _, _ = f.reinitialize_tracking(frame_ref,
        frame_gray, frame_gray_old, patches, x_grid_ref, y_grid_ref,
        pts_ref, dsc_ref)
251
252     # initialize grid sequences
253     grid_sequences = [ [] for _ in range(x_grid_ref.size)]
254
255     frame_final_RGB = f.write_frame(frame_gray_original, ROI, [], [],
256                                     [], [], output_vid, [], [],
257                                     [], [], frame_count, frame_range,
258                                     occluded_frames, n_frames,
259                                     n_last_frame,
260                                     peaks, flow_avg_scaled, hr,
261                                     constraints, basename)
262
263     elif main_tracking_start == True:
264
265         last_peak = len(peaks)
266
267         # draw occluded rectangle on prediction frames
268         if occluded_frames[0] <= frame_count <= occluded_frames[1] and
269             "occluded_frames" in locals():
270             frame_gray_original = f.rectangle_occlusion(frame_gray_original,
271                                                         occluded_area, alpha = 0.5 )
272
273         # analyze frame
274         patches, corners, grid_sequences, constraints, x_grid, y_grid,
275         x_grid_ML, y_grid_ML = \
276             f.analyze_frame(frame_gray, frame_gray_old, frame_ref, ROI,
277                             patches,
278                             corners, grid_sequences, constraints,
279                             frame_count, cap,

```

```

273             ML_enabled, model, scaler,
                sequence_history_length, occluded_frames,
274             x_grid_ref, y_grid_ref, x_grid_ML, y_grid_ML,
                pts_ref, dsc_ref)
275
276     # detect edges
277     if show_contours == True:
278         contours, hierarchy = f.find_contours(frame_gray, x_grid, y_grid)
279
280     # update pulse detection tracker
281     hr, flow_avg, flow_avg_scaled, peaks, new_peak, corners_pulse =
        f.detect_pulse(frame_gray, frame_gray_old, corners_pulse,
282                     peaks, flow_avg, lowcut, highcut, FPS, frame_count,
                        n_frames)
283
284     # reset corners on diastole
285     if len(peaks) > last_peak:
286         patches, corners, corners_new, x_grid, y_grid =
            f.reinitialize_tracking(frame_ref,
287                                 frame_gray, frame_gray_old, patches, x_grid_ref,
                                    y_grid_ref, pts_ref, dsc_ref)
288         diastole_frames.append(frame_count)
289
290     # save corner sequences
291     if occluded_frames[0] <= frame_count <= occluded_frames[1]:
292         grid_sequences = f.append_grid_sequences(grid_sequences, x_grid_ML,
                                                    y_grid_ML)
293         x_grid_hist, y_grid_hist, x_grid_hist_ML, y_grid_hist_ML =
            f.append_grid_prediction_history(x_grid, y_grid, x_grid_ML,
294                                           y_grid_ML, x_grid_hist, y_grid_hist,
                                                    x_grid_ML_hist, y_grid_ML_hist)
295     else:
296         grid_sequences = f.append_grid_sequences(grid_sequences, x_grid,
                                                    y_grid)
297
298     # write frame to file
299     t1_write = time.time()
300
301     frame_final_RGB = f.write_frame(frame_gray_original, ROI, patches,
        corners,
302                                     contours, hierarchy, output_vid,
303                                     x_grid, y_grid, x_grid_ML, y_grid_ML,

```

```

304         frame_count, frame_range, occluded_frames,
           n_frames,
305         n_last_frame, peaks, flow_avg_scaled,
306         hr, constraints, basename)
307
308         t2_write = time.time()
309         t_write = t_write + (t2_write - t1_write)
310
311         # show current frame
312         if frame_count == frame_range[0]:
313             cv2.imshow('Current Frame', frame_gray)
314         else:
315             cv2.imshow('Current Frame', frame_final_RGB)
316             if cv2.waitKey(1) & 0xFF == ord('q'):
317                 break
318
319         # save frame at desired frames
320         if frame_count in [main_tracking_start_frame, occluded_frames[0],
           occluded_frames[1], int((occluded_frames[1]+occluded_frames[0])/2)]:
321             screenshot_count +=1
322             screenshot_filename = output_path + 'plots/' + 'screenshot_' +
           str(screenshot_count) + '_' + basename + '.png'
323             cv2.imwrite(screenshot_filename, frame_final_RGB)
324
325         # save current frame for next cycle
326         frame_gray_old = frame_gray.copy()
327
328         # break the loop
329         else:
330             break
331
332         # Post-processing =====
333
334         # close video files
335         cap.release()
336         cv2.destroyAllWindows()
337         output_vid.release()
338
339         # update timer
340         t1_sec = time.time()
341         t1 = time.ctime()

```

```

342 total_secs = t1_sec - t0_sec - t_roi - t_read - t_write # subtract out frame read
      / write time
343 actual_FPS = (n_frames-frame_start_tracking)/total_secs
344
345 # print info
346 os.system('clear')
347
348 n_tracked_frames = int(n_frames-main_tracking_start_frame+1)
349 n_diastole_frames = len(diastole_frames)
350 n_failed_frames = len(constraints.failed_frames)
351 n_reinit_frames = len(constraints.reinitialized_frames)
352 n_predicted_frames = int(occluded_frames[1]-occluded_frames[0]+1)
353
354 fout = open(text_output_file,'w')
355 print("Filename:", input_file, file=fout)
356 print("\n", file=fout)
357 print("ROI: Upper Left:", ROI.start_pt, "Lower Right:", ROI.end_pt)
358 print("End Time:", t1, file=fout)
359 print("Total Tracking Time:", round(total_secs,1), " secs", file=fout)
360 print("Tracking Framerate:", round(actual_FPS, 1), " FPS", file=fout)
361 print("Percent of Full Speed:", round(actual_FPS/FPS*100, 1), "%", file=fout)
362 print("\n", file=fout)
363 print("Total Frames:", n_frames, file=fout)
364 print("Tracked Frames:", n_tracked_frames, '(',
      round(n_tracked_frames/n_frames*100, 1), '% )', file=fout)
365 print("Re-Initialized Frames (Diastole):", n_diastole_frames, '(',
      round(n_diastole_frames/n_frames*100, 1), '% )', file=fout)
366 print("Re-Initialized Frames (Failures):", n_reinit_frames, '(',
      round(n_reinit_frames/n_tracked_frames*100, 1), '% )', file=fout)
367 print("Failed Frames (Diastole):", n_failed_frames, '(',
      round(n_failed_frames/n_tracked_frames*100, 1), '% )', file=fout)
368 print("Predicted Frames:", n_predicted_frames, '(',
      round(n_predicted_frames/n_tracked_frames*100, 1), '% )', file=fout)
369 print("ROI: Upper Left:", ROI.start_pt, "Lower Right:", ROI.end_pt)
370 fout.close()
371
372 print("Filename:", input_file)
373 print("\n")
374 print("ROI: Upper Left:", ROI.start_pt, "Lower Right:", ROI.end_pt)
375 print("End Time:", t1)
376 print("Total Tracking Time:", round(total_secs,1), " secs")
377 print("Tracking Framerate:", round(actual_FPS, 1), " FPS")

```



```

378 print("Percent of Full Speed:", round(actual_FPS/FPS*100, 1), "%")
379 print("\n")
380 print("Total Frames:", n_frames)
381 print("Tracked Frames:", n_tracked_frames, '(' ,
        round(n_tracked_frames/n_frames*100, 1), '% )')
382 print("Re-Initialized Frames (Diastole):", n_diastole_frames, '(' ,
        round(n_diastole_frames/n_frames*100, 1), '% )')
383 print("Re-Initialized Frames (Failures):", n_reinit_frames, '(' ,
        round(n_reinit_frames/n_tracked_frames*100, 1), '% )')
384 print("Failed Frames (Diastole):", n_failed_frames, '(' ,
        round(n_failed_frames/n_tracked_frames*100, 1), '% )')
385 print("Predicted Frames:", n_predicted_frames, '(' ,
        round(n_predicted_frames/n_tracked_frames*100, 1), '% )')
386
387 # create plots
388 f.create_grid_plots(x_grid_hist, y_grid_hist, x_grid_ML_hist, y_grid_ML_hist,
        basename, output_path, text_output_file)
389 f.create_timeline_plots(constraints, n_last_frame, diastole_frames,
        main_tracking_start_frame, occluded_frames, basename, output_path)
390
391 # end profiler
392 pr.disable()
393 s = io.StringIO()
394 ps = pstats.Stats(pr, stream=s).sort_stats('tottime')
395 ps.print_stats()
396
397 with open(text_output_file, 'a') as fout:
398     fout.write(s.getvalue())

```

A.2 Feature Point Sequence Extraction

A.2.1 Overview

Algorithm 2: Feature Point Sequence Extraction.

Result: Obtain a sequence of feature points for each motion cycle in a video and save to file

```
begin
|
| open data set input video
| while new frame available do
| | pre-process frame (equalize histogram, denoise)
| | if first frame then
| | | begin pulse tracker
| | | if first diastole detected then
| | | | start feature tracker
| | | update optical flow for pulse and feature tracker
| | | append new feature locations to sequence array
| | | write and display current frame with overlays
| | end
| | save sequence array to file
| | close data set input video
| | close feature extraction output video
| end
end
```

A.2.2 Code

```
1 # =====
2 #
3 # file:    feature_extraction.py
4 # version: v1.0.0
5 # author:  Billy Hempstead
6 #
7 # summary: Opens an .avi file, extracts the feature points, and saves to .csv
8 #
9 # =====
10
11 # Imports =====
12
13 import cv2
14 import numpy as np
15 import matplotlib.pyplot as plt
16 import time, os, math
17 import functions as f
18 import classes as c
19 from sklearn import preprocessing
20 from scipy.signal import find_peaks
21 from tkinter import filedialog
22 from tkinter import *
23
24 # Global Variables =====
25
26 # ROI size
27 ROI_radius = 250
28
29 # optical flow settings
30 OF_win_size = (15, 15)
31 OF_max_level = 2
32
33 # filtering parameters
34 lowcut = 30/60 # Hz
35 highcut = 70/60 # Hz
36
37 # colors
38 colors = c.Colors()
39
40 # output folder
```

```

41 output_path = './data_feature_extraction/'
42
43 # initialize heart rate parameters
44 hr = 0
45 new_peak = False
46 peaks = []
47 corners_1 = []
48 corners_new_1 = []
49 corners_new_pulse = []
50 status_1 = []
51 status_1_lost_index = []
52 err_1 = []
53 flow_avg = []
54
55 # initialize sequence lists
56 sequences = []
57
58 # tracking parameters
59 trk_max_corners = 10000
60 trk_quality_level = 0.001
61 trk_min_distance = 3
62 trk_block_size = 3
63 trk_use_harris = 0
64 trk_k = 0.004
65
66 # pulse tracking parameters
67 pulse_max_corners = 10
68 pulse_quality_level = 0.005
69 pulse_min_distance = 25
70 pulse_block_size = 3
71 pulse_use_harris = 0
72 pulse_k = 0.004
73
74 # Main Loop =====
75
76 # get filename
77 root = Tk()
78 root.fullpath = filedialog.askopenfilename(initialdir="./data/",
79                                           filetypes=(("AVI File", "*.avi"),("All Files","*.*")),
80                                           title = "Choose a file.")
81 fullpath = root.fullpath
82 root.destroy()

```

```

83
84 pathname = os.path.dirname(fullpath)
85 input_file_name = os.path.basename(fullpath)
86 basename = os.path.splitext(input_file_name)[0]
87 output_csv = pathname + '/' + basename + '.csv'
88 output_npy = output_path + basename + '.npy'
89 input_file = fullpath
90 output_file = output_path + basename + '_w_features.avi'
91
92 # load input video and initialize output video
93 cap, output_vid, frame_width, frame_height, n_frames, FPS = \
94     f.import_video(input_file, output_file)
95
96 # set sampling frequency
97 fs = FPS
98
99 # initialize timer
100 t0_sec = time.time()
101 t0 = time.ctime()
102 print(t0)
103
104 # create ROI mask
105 mask = np.zeros((frame_width, frame_height), dtype=np.uint8)
106 mask = cv2.circle(mask, center = (int(frame_width/2), int(frame_height/2)),
107     radius = ROI_radius, color = colors.white,
108     thickness = -1)
109
110 # loop through frames
111 while(cap.isOpened()):
112     # get current frame
113     frame_count = int(cap.get(cv2.CAP_PROP_POS_FRAMES))
114
115     # update progress
116     os.system('clear')
117     print('Processing frame: ', int(frame_count), ' of ', int(n_frames))
118
119     # get a new frame from video
120     ret, frame = cap.read()
121
122     if ret == True:
123

```

```

124     # adjust frame
125     frame_gray = f.preprocess_frame(frame, 'adaptive', 'bilateral')
126
127     if frame_count == 0:
128
129         # detect pulse corners
130         corners_pulse = cv2.goodFeaturesToTrack( frame_gray, pulse_max_corners,
131             pulse_quality_level, pulse_min_distance, mask, pulse_block_size,
132             pulse_use_harris, pulse_k )
133
134         # convert first frame to RGB
135         frame_RGB = cv2.cvtColor(frame_gray, cv2.COLOR_GRAY2RGB)
136
137         # save current frame for pulse calculation
138         frame_gray_old = frame_gray.copy()
139
140     elif len(peaks) == 1 and new_peak == True:
141
142         # detect feature corners on first diastole
143         corners_1 = cv2.goodFeaturesToTrack( frame_gray, trk_max_corners,
144             trk_quality_level, trk_min_distance, mask, trk_block_size,
145             trk_use_harris, trk_k )
146
147         # convert frame to RGB
148         frame_RGB = cv2.cvtColor(frame_gray, cv2.COLOR_GRAY2RGB)
149
150         # initialize corner sequences
151         corners_1_seq = [ [] for _ in range(corners_1.shape[0])]
152
153         for i, corner in enumerate(corners_1):
154             corners_1_seq[i] = np.append(corners_1_seq[i], corner[0])
155
156     elif frame_count > 0:
157
158         # update pulse tracker
159         corners_new_pulse, _, _ = cv2.calcOpticalFlowPyrLK(frame_gray_old,
160             frame_gray, corners_pulse, OF_win_size, OF_max_level)
161
162         # update feature tracker
163         if len(corners_1) > 0:
164             corners_new_1, status_1, err_1 = cv2.calcOpticalFlowPyrLK(

```

```

165         frame_gray_old, frame_gray, corners_1, OF_win_size,
           OF_max_level)
166
167     # collect lost feature indices
168     if status_1.min() == 0:
169         status_1_lost_index = np.append(status_1_lost_index,
           np.where(status_1 == 0)[0])
170         status_1_lost_index = np.unique(status_1_lost_index)
171
172     # append new corners to sequence
173     for i, corner in enumerate(corners_new_1):
174         corners_1_seq[i] = np.append(corners_1_seq[i], corner)
175
176     # save new corners for next cycle
177     corners_1 = corners_new_1
178     corners_pulse = corners_new_pulse
179
180     # calculate flow, heart rate, phase
181     hr, flow_avg, flow_avg_scaled, peaks, new_peak, corners = f.detect_pulse(
182         frame_gray, frame_gray_old, corners_pulse, peaks, flow_avg,
183         lowcut, highcut, fs, frame_count, n_frames )
184
185     # draw overlays
186     frame_RGB = f.draw_overlays_extraction(frame_gray, corners_new_1,
187         status_1, frame_width, frame_height, frame_count,
188         n_frames, peaks, new_peak, flow_avg_scaled, basename, hr,
           ROI_radius)
189
190     # save current frame for next cycle
191     frame_gray_old = frame_gray.copy()
192
193     # write frame to file
194     output_vid.write(frame_RGB)
195
196     # show converted frame
197     cv2.imshow('Current Frame',frame_RGB)
198     if cv2.waitKey(1) & 0xFF == ord('q'):
199         break
200
201     # break the loop
202 else:
203     break

```

```

204
205 # Post-processing =====
206
207 # close video files
208 cap.release()
209 cv2.destroyAllWindows()
210 output_vid.release()
211
212 # stop timer
213 t1_sec = time.time()
214 t1 = time.ctime()
215 total_secs = t1_sec-t0_sec
216
217 # reshape sequences
218 sequence_length = int(len(corners_1_seq[0])/2)
219 for i, sequence in enumerate(corners_1_seq):
220     corners_1_seq[i] = sequence.reshape((sequence_length, 2))
221
222 # remove lost sequences
223 for ele in sorted(status_1_lost_index, reverse = True):
224     del corners_1_seq[int(ele)]
225
226 # save data
227 np.save(output_npy, corners_1_seq)
228
229 # print info
230 os.system('clear')
231 print("Number of Point Sequences Extracted:", len(corners_1_seq))
232 print("End Time:", t1)
233 print("Total Seconds:", round(total_secs,3))
234 print('\n')

```

A.3 Data Set Creation

A.3.1 Overview

Algorithm 3: Data Set Creation.

Result: Consolidate and process extracted data into training and testing data sets

begin

- open each point sequence file
- append to data set array
- shuffle data set
- create array of shorter sub-sequences
- shuffle data set
- split data set into training and test data sets
- fit a Standard scaler to training data set only
- scale training and test data set values
- separate training and test data sets into features and labels
- save training and test data sets and scaler model

end

A.3.2 Code

```
1 # =====
2 #
3 # file:    create_dataset.py
4 # version: v1.0.0
5 # author:  Billy Hempstead
6 #
7 # summary: Compiles extracted data sequences into a single dataset.
8 #
9 # =====
10
11 # Initialization =====
12
13 from __future__ import absolute_import, division, print_function, unicode_literals
14
15 import numpy as np
16
17 from sklearn.model_selection import train_test_split
18 from sklearn.externals import joblib
19 from sklearn import preprocessing
20 from sklearn.utils import shuffle
21
22 import os, argparse, platform, math, time, io, glob
23
24 # Functions =====
25
26 def load_dataset(data_file_dir):
27
28     files = sorted(glob.glob(data_file_dir + '/*.npy'))
29     arrays = []
30     data_lengths = []
31
32     for f in files:
33         temp_data = np.load(f, allow_pickle=True)
34         data_lengths.append(temp_data.shape[1])
35
36     min_data_length = min(data_lengths)
37
38     for f in files:
39         temp_data = np.load(f, allow_pickle=True)
40         temp_data = temp_data[:, :min_data_length]
```

```

41     arrays.append(temp_data)
42
43     data = np.concatenate(arrays)
44
45     return data
46
47 # -----
48
49 def get_scaler(data, scaler_type):
50
51     os.system('clear')
52     print("Fitting Scaler")
53
54     step_size = data[0].shape[0]
55
56     if scaler_type == "minmax":
57         scaler = preprocessing.MinMaxScaler()
58     elif scaler_type == "standard":
59         scaler = preprocessing.StandardScaler()
60     elif scaler_type == "robust":
61         scaler = preprocessing.RobustScaler()
62
63     data_flat = np.concatenate(data)
64     data_flat = np.concatenate(data_flat)
65
66     scaler.fit(data_flat.reshape(-1,1))
67
68     return scaler
69
70 # -----
71
72 def scale_dataset(data, scaler, transform_type):
73
74     os.system('clear')
75     print("Scaling Data")
76
77     data_flat = np.concatenate(data)
78     data_flat = np.concatenate(data_flat)
79
80     if transform_type == 'transform':
81         data_flat = scaler.transform(data_flat.reshape(-1,1))
82

```

```

83     elif transform_type == 'inverse':
84         data_flat = scaler.inverse_transform(data_flat.reshape(-1,1))
85
86     os.system('clear')
87
88     return data_flat
89
90 # -----
91
92 def shift_dataset(data, step_size):
93
94     # create shifted versions of each sequence
95     n_seq = data.shape[0]
96     data_shifted = []
97     sequence_size = int(step_size*2)
98
99     for i, sequence in enumerate(data):
100
101         os.system('clear')
102         print("Creating shifted copies - Processing sequence number", i+1, 'of',
103             n_seq)
104
105         for j in range(len(sequence)):
106             shifted_seq = sequence[ j:j+step_size ]
107             if shifted_seq.size == sequence_size:
108                 data_shifted.append(shifted_seq)
109
110         os.system('clear')
111
112     return data_shifted
113 # -----
114
115 def split_data(data, validation_split):
116
117     # should be done after shuffling
118
119     data_length = len(data)
120     split_point = int(data_length * (1-validation_split))
121
122     train_data = data[:split_point]
123     test_data = data[split_point:]

```

```

124
125     return train_data, test_data
126
127 # -----
128
129 def separate_dataset(data, step_size):
130
131     n_seq = int(len(data)/step_size/2)
132     data = data.reshape((n_seq, step_size, 2))
133
134     features = data[:, :-1]
135     labels = data[:, -1]
136
137     return features, labels
138
139 # Main =====
140
141 # display version info
142 print('Python Version: ', platform.python_version())
143 print('Numpy Version: ', np.__version__)
144
145 step_size = 10
146 data_reduction_ratio = 0.4
147
148 data_file_dir = './data_feature_extraction'
149
150 # load and compile datafiles
151 data = load_dataset(data_file_dir)
152
153 # create shifted versions of sequences
154 data_split = int(data_reduction_ratio*len(data))
155 data = shuffle(data[:data_split])
156 data = shift_dataset(data, step_size = step_size)
157
158 # shuffle all of the sequences
159 data = shuffle(data, random_state = 0 )
160
161 # separate shuffled sequences into training and testing data
162 train_data, test_data = split_data(data, validation_split = 0.2)
163
164 # scale data based on training data only

```

```
165 # train_split = int(train_split_ratio*len(train_data)) # reduce amount of data to
    scaler to avoid out of memory
166 scaler = get_scaler(train_data, 'standard')
167 train_data = scale_dataset(train_data, scaler, 'transform')
168 test_data = scale_dataset(test_data, scaler, 'transform')
169
170 # separate features and labels
171 train_data_x, train_data_y = separate_dataset(train_data, step_size)
172 test_data_x, test_data_y = separate_dataset(test_data, step_size)
173
174 # save files
175 joblib.dump(scaler, './models/scaler_model.pkl')
176 joblib.dump([train_data_x, train_data_y, test_data_x, test_data_y],
    './models/dataset.pkl')
177
178 # print sample
179 sample_number = np.random.randint(0, train_data_x.shape[0])
180 print('Features: ', '\n', train_data_x.shape)
181 print(train_data_x[sample_number])
182 print('Labels: ', '\n', train_data_y.shape)
183 print(train_data_y[sample_number])
184
185 print('Features: ', '\n', train_data_x.shape)
186 # print(scale_dataset(train_data_x[sample_number], scaler, 'inverse'))
187 print(scaler.inverse_transform(train_data_x[sample_number]))
188
189 print('Labels: ', '\n', train_data_y.shape)
190 # print(scale_dataset(train_data_y[sample_number], scaler, 'inverse'))
191 print(scaler.inverse_transform(train_data_y[sample_number].reshape(1,-1)))
```

A.4 Neural Network Training

A.4.1 Overview

Algorithm 4: Neural Network Training.

Result: Train the neural network on the point sequence data set

```
begin
  open training and test data sets and scaler model
  create NN model
  for number of epochs do
    train model with training data set
    check performance on test data set
  end
  run a sample prediction to check output
  save trained model
end
```

A.4.2 Code

```
1 # =====
2 #
3 # file:    nn_training.py
4 # version: v1.0.0
5 # author:  Billy Hempstead
6 #
7 # summary: Neural Network training code for extracted heart feature point
8 #          sequences.
9 #
10 # =====
11
12 # Initialization =====
13
14 from __future__ import absolute_import, division, print_function, unicode_literals
15
16 import tensorflow as tf
17 import numpy as np
18 import matplotlib.pyplot as plt
19
20 from sklearn.model_selection import train_test_split
21 from sklearn.externals import joblib
22 from sklearn import preprocessing
23 from sklearn.utils import shuffle
24
25 import os, argparse, platform, math, time, io, glob
26
27 # Main =====
28
29 # display version info
30 print('Python Version: ', platform.python_version())
31 print('Tensorflow Version:', tf.__version__)
32 print('Numpy Version:   ', np.__version__)
33
34 earllystop_callback = tf.keras.callbacks.EarlyStopping(monitor='val_loss',
35               patience=3)
36
37 # import dataset and scaler
38 train_data_x, train_data_y, test_data_x, test_data_y =
39     joblib.load('./models/dataset.pkl')
40 scaler = joblib.load('./models/scaler_model.pkl')
```



```

39
40 # start timer
41 t0_sec = time.time()
42
43 # create model
44 model = tf.keras.models.Sequential()
45 model.add(tf.keras.layers.GRU(64, input_shape=(None, 2), return_sequences = True))
46 # model.add(tf.keras.layers.Dropout(0.2))
47 model.add(tf.keras.layers.GRU(32))
48 # model.add(tf.keras.layers.Dropout(0.2))
49 model.add(tf.keras.layers.Dense(2, activation = 'linear'))
50
51 # train model
52 # opt = tf.keras.optimizers.SGD(lr=0.01, momentum=0.9, decay=0.01)
53 opt = tf.keras.optimizers.Adam()
54 model.compile(loss='mean_squared_error', optimizer=opt, metrics=['accuracy'])
55 model.summary()
56 history = model.fit(train_data_x, train_data_y, validation_data=(test_data_x,
    test_data_y), batch_size=512, epochs=5)#, callbacks=[earlystop_callback])
57
58 # update timer
59 t1_sec = time.time()
60 total_min = (t1_sec - t0_sec)/60
61
62 # save model
63 model.save('./models/sequence_model.h5')
64
65 # plot loss during training
66 plt.title('Loss / Mean Squared Error')
67 plt.plot(history.history['loss'], label='train')
68 plt.plot(history.history['val_loss'], label='test')
69 plt.legend()
70 plt.xlabel('Epochs')
71 plt.ylabel('Loss')
72 plt.savefig('./models/loss.png')
73 plt.show()
74
75 # plot accuracy during training
76 plt.title('Accuracy')
77 plt.plot(history.history['accuracy'], label='train')
78 plt.plot(history.history['val_accuracy'], label='test')
79 plt.legend()

```

```
80 plt.xlabel('Epochs')
81 plt.ylabel('Accuracy')
82 plt.savefig('./models/accuracy.png')
83 plt.show()
84
85 # check output
86 sample_number = np.random.randint(0, test_data_x.shape[0])
87 sample = test_data_x[sample_number].reshape((1,
      test_data_x[sample_number].shape[0], test_data_x[sample_number].shape[1]))
88 pred = model.predict(sample)
89 sample = scaler.inverse_transform(sample[-1])
90 pred = scaler.inverse_transform(pred)
91
92 print(sample)
93 print(pred)
94
95 # print info
96 print()
97 print("Total Training Time:", round(total_min,1), "min")
```

A.5 Supporting Classes / Functions

The following code provides classes and supporting functions that are called in the source files in previous sections.

A.5.1 Classes

```
1 # =====
2 #
3 # file:    classes.py
4 # version: v1.0.0
5 # author:  Billy Hempstead
6 #
7 # summary: Contains classes for use with the main piecewise tracking algorithm
8 #
9 # =====
10
11 import numpy as np
12 import math
13
14 class Patch:
15
16     def __init__(self):
17
18         self.index = 1           # overall patch number
19
20         # corners points of patch (UL, UR, LR, LL)
21         self.corners = np.empty([4,2], dtype='int')
22
23         # index of points inside patch
24         self.point_idx = np.empty([0], dtype='int')
25
26         self.u_avg = 0           # average horizontal flow
27         self.v_avg = 0           # average vertical flow
28         self.u_avg_neighbors = 0
29         self.v_avg_neighbors = 0
30         self.u_diff = 0
31         self.v_diff = 0
32         self.u_diff_min = 0
33         self.u_diff_max = 0
34         self.v_diff_min = 0
35         self.v_diff_max = 0
```

```

36     self.normal_corner_dist = False      # do points have normal dist.?
37     self.continuous_neighbor_flow = False # is neighbor flow continuous?
38     self.u_std_dev = 0
39     self.v_std_dev = 0
40     self.u_outliers = 0
41     self.v_outliers = 0
42
43     # method to move the corners of each patch by the average of the optical flow
44     def update_corners(self):
45
46         if not math.isnan(self.u_avg):
47             for i in range(4):
48                 self.corners[i][0] = round(float(self.corners[i][0]) + self.u_avg)
49
50         if not math.isnan(self.v_avg):
51             for i in range(4):
52                 self.corners[i][1] = round(float(self.corners[i][1]) + self.v_avg)
53
54         return self.corners
55
56     class Rect:
57
58         def __init__(self, x, y, width, height):
59             self.x = x
60             self.y = y
61             self.width = width
62             self.height = height
63             self.start_pt = (self.x, self.y)
64             self.end_pt = (self.x + self.width, self.y + self.height)
65
66     class Constraints:
67
68         def __init__(self):
69             self.all_normal = False
70             self.neighbor_flow_continuous = False
71             self.passed = False
72             self.failed_frames = []
73             self.reinitialized_frames = []
74             self.fails = 0
75             self.fail_limit = 5
76
77     class Colors:

```

```
78
79     def __init__(self):
80         self.blue      = (255, 0, 0)
81         self.red       = (0, 0, 255)
82         self.green     = (0, 255, 0)
83         self.white     = (255, 255, 255)
84         self.black     = (0, 0, 0)
85         self.dark_green = (0, 125, 0)
86         self.magenta   = (225, 0, 255)
87         self.cyan      = (255, 255, 0)
88         self.yellow    = (0, 255, 255)
```

A.5.2 Functions

```
1 # =====
2 #
3 # file:    functions.py
4 # version: v1.0.0
5 # author:  Billy Hempstead
6 #
7 # summary: Contains functions for use with the main piecewise tracking algorithm
8 #           and feature extraction
9 #
10 # =====
11
12 # Initialization =====
13
14 import cv2
15 import tensorflow as tf
16 import numpy as np
17 import matplotlib.pyplot as plt
18 from matplotlib.ticker import FormatStrFormatter
19 from matplotlib.ticker import PercentFormatter
20 import time
21 import os
22 import math
23 import classes as c
24 from scipy import stats as spstats
25 from scipy import ndimage
26 from scipy.signal import butter, lfilter
27 from sklearn import preprocessing
28 from scipy.signal import find_peaks
29 from scipy.spatial import distance
30
31 # Global Variables =====
32
33 # colors
34 colors = c.Colors()
35
36 # SURF detector and matcher
37 detector = cv2.xfeatures2d.SURF_create(hessianThreshold = 1000, nOctaves = 4,
38                                       nOctaveLayers = 3, extended = True,
39                                       upright = True)
40
```

```

41 matcherBF = cv2.BFMatcher(cv2.NORM_L2, crossCheck=False)
42
43 # detection parameters
44 det_max_corners = 1000
45 det_quality_level = 0.001
46 det_min_distance = 3
47 det_block_size = 3
48 det_use_harris = 0
49 det_k = 0.04
50
51 # convolution kernel for neighbor flow average
52 kernel = np.ones((3, 3))
53 kernel[1, 1] = 0
54
55 # equalization / blur settings
56 CLAHE_clip_limit = 30.0
57 CLAHE_tile_size = (20, 20)
58 median_blur_level = 5
59 BLF_dia = 5
60 BLF_sigma_color = 60
61 BLF_sigma_space = 60
62
63 # tracker settings
64 OF_win_size = (15, 15)
65 OF_max_level = 2
66
67 # constraint settings
68 flow_threshold_margin = 0.1
69 min_corners_in_patch = 8
70
71 # contours settings
72 contours_max_level = 2
73 contours_lower_th = 0
74 contours_upper_th = 255
75
76 # grid settings
77 grid_outline_thickness = 4
78 grid_color_normal = colors.green
79 grid_color_ML = colors.yellow
80
81 # drawing options
82 draw_ROI = False

```

```

83 draw_grid = True
84 draw_grid_ML = True
85 draw_corners = True
86 draw_contours = False
87 draw_patch_corners = False
88 draw_overlays = True
89 draw_failed_frames = True
90
91 # Functions =====
92
93 # === Input/Output =====
94
95 def import_video(input_file, output_file):
96
97     # open input file
98     cap = cv2.VideoCapture(input_file)
99
100    # check if file opened successfully
101    if (cap.isOpened() == False):
102        print("Error opening video stream or file")
103
104    # get video properties
105    frame_width = int(cap.get(cv2.CAP_PROP_FRAME_WIDTH))
106    frame_height = int(cap.get(cv2.CAP_PROP_FRAME_HEIGHT))
107    n_last_frame = int(cap.get(cv2.CAP_PROP_FRAME_COUNT))
108    FPS = cap.get(cv2.CAP_PROP_FPS)
109
110    # create writer object
111    fourcc = cv2.VideoWriter_fourcc(*'XVID')
112    out = cv2.VideoWriter(output_file,fourcc, FPS, (frame_height, frame_width) )
113
114    return cap, out, frame_width, frame_height, n_last_frame, FPS
115
116 # === Image Processing =====
117
118 def preprocess_frame(frame, equalization_type, blur_type):
119
120    # operations performed on each frame before tracking
121
122    # convert frame to grayscale
123    if len(frame.shape) == 3:
124        frame = cv2.cvtColor(frame, cv2.COLOR_RGB2GRAY)

```



```

125
126 # equalize frame
127 if equalization_type == 'normal':
128     frame = cv2.equalizeHist(frame)
129
130 elif equalization_type == "adaptive":
131     # apply Contrast Limited Adaptive Histogram Equalization (CLAHE)
132     clahe = cv2.createCLAHE(clipLimit = CLAHE_clip_limit, tileGridSize =
133         CLAHE_tile_size)
134     frame = clahe.apply(frame)
135
136 else:
137     frame = frame
138
139 # blur to reduce noise to focus on actual features
140 if blur_type == 'median':
141     frame = cv2.medianBlur(frame, median_blur_level)
142
143 elif blur_type == 'bilateral':
144     # better retains edges
145     frame = cv2.bilateralFilter(frame, BLF_dia, BLF_sigma_color,
146         BLF_sigma_space)
147
148 else:
149     frame = frame
150
151 return frame
152
153 # === Analysis =====
154 def analyze_frame(frame_curr, frame_prev, frame_ref, ROI, patches, corners,
155     grid_sequences, constraints, frame_count, cap, ML_enabled, model,
156     scaler, sequence_history_length, occluded_frames,
157     x_grid_ref, y_grid_ref, x_grid_ML, y_grid_ML, pts_ref, dsc_ref):
158 # tracking/flow update and patch analysis for current frame
159
160 # update tracker
161 corners_new, _, _ = cv2.calcOpticalFlowPyrLK(frame_prev, frame_curr, corners,
162     OF_win_size, OF_max_level)
163
164 # calculate flow of features between frames

```

```

164 patches = update_patch_flow(patches, corners, corners_new)
165
166 # check constraints
167 constraints.all_normal, constraints.neighbor_flow_continuous =
    check_constraints(patches)
168
169 if constraints.all_normal or constraints.neighbor_flow_continuous == False:
170     constraints.passed = False
171     constraints.fails += 1
172     constraints.failed_frames.append(frame_count)
173
174 else:
175     constraints.passed = True
176     constraints.fails = 0
177
178 # trigger reinitialization
179 if constraints.fails == constraints.fail_limit:
180     constraints.reinitialized_frames.append(frame_count)
181     patches, corners, corners_new, _, _ = reinitialize_tracking(frame_ref,
182                                                                frame_curr, frame_prev, patches,
                                                                x_grid_ref, y_grid_ref, pts_ref,
                                                                dsc_ref)
183
184 # move patch boundaries by flow
185 for patch in patches:
186     patch.update_corners()
187
188 # save corners for next cycle
189 corners = corners_new
190
191 # get updated grid locations
192 x_grid, y_grid = create_grid(patches)
193
194 # ML prediction loop
195 if occluded_frames[0] <= frame_count <= occluded_frames[1]+1:
196
197     x_grid_ML = np.zeros(len(grid_sequences))
198     y_grid_ML = np.zeros_like(x_grid_ML)
199     n_rows = int(np.sqrt(len(x_grid_ML)))
200
201     for i, sequence in enumerate(grid_sequences):
202         print('Predicting point ', i, ' of ', len(grid_sequences) )

```

```

203         output = run_NN(sequence[-(int(sequence_history_length*2)):], model,
                scaler)
204         x_grid_ML[i] = output[0][0]
205         y_grid_ML[i] = output[0][1]
206
207         x_grid_ML = x_grid_ML.reshape(n_rows, n_rows).astype(int)
208         y_grid_ML = y_grid_ML.reshape(n_rows, n_rows).astype(int)
209
210     return patches, corners, grid_sequences, constraints, x_grid, y_grid,
        x_grid_ML, y_grid_ML
211
212     # -----
213
214     def ROI_selection(frame, ROI):
215
216         # prompts user for selection of ROI
217
218         cv2.destroyAllWindows()
219         ROI_sel = cv2.selectROI("Select Region-of-Interest", frame)
220         cv2.waitKey(0)
221         cv2.destroyAllWindows()
222
223         ROI.x = ROI_sel[0]
224         ROI.y = ROI_sel[1]
225         ROI.width = ROI_sel[2]
226         ROI.height = ROI_sel[3]
227         ROI.start_pt = (ROI_sel[0], ROI_sel[1])
228         ROI.end_pt = (ROI_sel[0]+ROI_sel[2], ROI_sel[1]+ROI_sel[3])
229
230         return ROI
231
232     # -----
233
234     def determine_n_patches(frame_diastole, frame_next, ROI, n_patch_rows,
        n_patches_max ):
235
236         # gradually increases the number of patches until all patches meet constraint 1
237
238         print('Determining Number of Patches within ROI')
239
240         n_patch_rows +=1
241         patches, corners, n_patches = create_patches(frame_next, ROI, n_patch_rows)

```

```

242
243 # calculate optical flow between diastole frame and prev frame
244 corners_new, _, _ = cv2.calcOpticalFlowPyrLK(frame_diastole, frame_next,
        corners, OF_win_size, OF_max_level)
245
246 # update patch flow
247 patches = update_patch_flow(patches, corners, corners_new)
248
249 constraint_1, _ = check_constraints(patches)
250
251 for patch in patches:
252     if len(patch.point_idx) >= min_corners_in_patch:
253         sufficient_corners = True
254     else:
255         sufficient_corners = False
256         break
257
258 # constraint 1 and minimum number of points per patch must be met
259 # continue on if constraints met or number of patches exceeds max
260 # TODO determine how to handle constraints never met
261 if constraint_1 == True and sufficient_corners == True or n_patches ==
        n_patches_max:
262     n_patches_found = True
263 else:
264     n_patches_found = False
265
266 return patches, n_patch_rows, n_patches, n_patches_found
267
268 # -----
269
270 def set_flow_thresholds(frame_curr, frame_prev, corners, patches):
271
272 # adjusts the neighbor difference thresholds based on latest tracking update
273
274     print('Setting Flow Thresholds')
275
276     # calculate optical flow
277     corners_new, _, _ = cv2.calcOpticalFlowPyrLK(frame_curr, frame_prev, corners,
        OF_win_size, OF_max_level)
278
279     # update patch flow
280     patches = update_patch_flow(patches, corners, corners_new)

```

```

281
282     for i, patch in enumerate(patches):
283         patch.u_diff = patch.u_avg - patch.u_avg_neighbors
284         patch.v_diff = patch.v_avg - patch.v_avg_neighbors
285
286         if patch.u_diff < patch.u_diff_min:
287             patch.u_diff_min = patch.u_diff * (1 + flow_threshold_margin)
288             #TODO should subtract margin in next version
289         elif patch.u_diff > patch.u_diff_max:
290             patch.u_diff_max = patch.u_diff * (1 + flow_threshold_margin)
291
292         if patch.v_diff < patch.v_diff_min:
293             patch.v_diff_min = patch.v_diff * (1 + flow_threshold_margin)
294             #TODO should subtract margin in next version
295         elif patch.v_diff > patch.v_diff_max:
296             patch.v_diff_max = patch.v_diff * (1 + flow_threshold_margin)
297
298     return patches
299
300 # -----
301
302 def detect_pulse( frame_curr, frame_prev, corners, peaks, flow_avg, lowcut,
303                 highcut,
304                 fs, frame_count, n_frames ):
305     # updates pulse detection history
306
307     frame_height = frame_curr.shape[0]
308     frame_width = frame_curr.shape[1]
309
310     flow = 0
311     hr = 0
312
313     # calculate optical flow
314     corners_new, _, _ = cv2.calcOpticalFlowPyrLK(frame_prev, frame_curr, corners,
315         OF_win_size, OF_max_level)
316
317     # calculate average point flow
318     for i in range(len(corners_new)):
319         p1 = corners_new[i]
320         p2 = corners[i]
321         dist = cv2.norm(p1,p2)

```

```

321     flow += dist
322
323     flow_avg = np.append(flow_avg, flow/len(corners_new))
324
325     # filter/scale average flow
326     flow_avg_filtered = butter_bandpass_filter(flow_avg, lowcut, highcut, fs,
327                                               order=5)
328
329     # plot graph
330     # normalization to fit to canvas dimension
331     flow_avg_scaled = np.asarray(flow_avg_scaled, dtype = np.float32)
332
333     # detect peaks
334     num_peaks = len(peaks)
335     peaks, _ = find_peaks(-flow_avg_scaled, height = 0.1)
336
337     if len(peaks) > num_peaks:
338         new_peak = True
339     else:
340         new_peak = False
341
342     # calculate heart rate
343     if len(peaks) > 1:
344         average_frames_per_cycle = sum(np.diff(peaks))/(len(peaks)-1)
345         time_per_frame = 1/fs
346         time_per_cycle = average_frames_per_cycle * time_per_frame
347         hr = int(60 / time_per_cycle) # bpm
348
349     return hr, flow_avg, flow_avg_scaled, peaks, new_peak, corners_new
350
351 # -----
352
353 def find_contours(frame, x_grid, y_grid):
354
355     # finds the contours (edges) within the ROI
356
357     pts = get_grid_outline(x_grid, y_grid)
358     pts = np.asarray( pts, dtype = np.int32).reshape((-1,1,2))
359
360     mask = np.zeros((frame.shape[0], frame.shape[1], 1), np.uint8)
361     mask = cv2.polylines(mask, [pts], True, colors.white, thickness = 2)

```

```

362     mask = cv2.fillPoly(mask, [pts], [255])
363
364     frame = cv2.bitwise_and(frame, frame, mask = mask)
365
366     # # apply blur
367     # frame = cv2.medianBlur(frame, median_blur_level)
368
369     # apply thresholding
370     retval, frame = cv2.threshold(frame, contours_lower_th, contours_upper_th,
        cv2.THRESH_OTSU)
371
372     # detect contours
373     contours, hierarchy = cv2.findContours(frame, cv2.RETR_CCOMP,
        cv2.CHAIN_APPROX_SIMPLE)
374
375     return contours, hierarchy
376
377 # === Filtering =====
378
379 def butter_bandpass(lowcut, highcut, fs, order=5):
380     nyq = 0.5 * fs
381     low = lowcut / nyq
382     high = highcut / nyq
383     b, a = butter(order, [low, high], btype='band')
384     return b, a
385
386
387 def butter_bandpass_filter(data, lowcut, highcut, fs, order=5):
388     b, a = butter_bandpass(lowcut, highcut, fs, order=order)
389     y = lfilter(b, a, data)
390     return y
391
392 # === Neural Network =====
393
394 def run_NN(sequence, model, scaler):
395
396     # run a new point prediction
397
398     # scale sequence
399     x_input = np.asarray(sequence)
400     x_input = scaler.transform(x_input.reshape(-1,1))
401     x_input = x_input.reshape(1, int(len(sequence)/2), 2)

```

```

402
403     # run prediction
404     yhat = model.predict(tf.cast(x_input,tf.float32), verbose=0)
405
406     # inverse scale prediction
407     yhat = scaler.inverse_transform(yhat)
408
409     return yhat
410
411 # === Re-initialization =====
412
413 def reinitialize_tracking(frame_ref, frame_curr, frame_prev, patches,
414                          x_grid_ref, y_grid_ref, pts_ref, dsc_ref):
415
416     # performs a homography operation to regain tracking
417
418     # detect features
419     pts_prev, dsc_prev = detector.detectAndCompute(frame_prev, None)
420
421     # match descriptors
422     matches = matcherBF.match(dsc_prev, dsc_ref)
423     matches_d_sort = sorted(matches, key = lambda x:x.distance)
424
425     # extract the matched keypoints
426     pts_prev_match = np.float32([pts_prev[m.queryIdx].pt for m in
427                                 matches_d_sort]).reshape(-1,1,2)
427     pts_ref_match = np.float32([pts_ref[m.trainIdx].pt for m in
428                                 matches_d_sort]).reshape(-1,1,2)
429
430     # find homography matrix using RANSAC
431     H, _ = cv2.findHomography(pts_ref_match, pts_prev_match, \
432                               method = cv2.RANSAC, ransacReprojThreshold = 100.0)
433
434     # transform grid points with homography matrix
435     x = x_grid_ref.flatten()
436     y = y_grid_ref.flatten()
437     grid = np.array([list(zip(x,y))], dtype=np.float32)
438     grid_warp = cv2.perspectiveTransform(grid, H)
439
440     # unravel warped combined grid back to x and y grid arrays
441     w = x_grid_ref.shape[0]
442     h = x_grid_ref.shape[1]

```



```

442     x_grid_new = grid_warp[0][:,0]
443     y_grid_new = grid_warp[0][:,1]
444     x_grid_new = np.reshape(x_grid_new.astype(int), (w,h))
445     y_grid_new = np.reshape(y_grid_new.astype(int), (w,h))
446
447     # mask to patch border
448     pts = get_grid_outline(x_grid_new, y_grid_new)
449     pts = np.asarray( pts, dtype = np.int32).reshape((-1,1,2))
450
451     mask = np.zeros((frame_prev.shape[0], frame_prev.shape[1], 1), np.uint8)
452     mask = cv2.polylines(mask, [pts], True, colors.white, thickness = 2 )
453     mask = cv2.fillPoly(mask, [pts], [255])
454
455     corners = cv2.goodFeaturesToTrack(frame_prev, maxCorners=det_max_corners,
456                                     qualityLevel=det_quality_level,
457                                     minDistance=det_min_distance,
458                                     blockSize = det_block_size,
459                                     useHarrisDetector = det_use_harris, mask=mask,
460                                     k = det_k)
461
462     corners_new, _, _ = cv2.calcOpticalFlowPyrLK(frame_prev, frame_curr, corners,
463                                                  OF_win_size, OF_max_level)
464
465     # reassign corner points after adjusting patch boundaries
466     patches = move_patch_corners(patches, x_grid_new, y_grid_new)
467     patches = assign_patch_points(patches, corners)
468     patches = update_patch_flow(patches, corners, corners_new)
469
470     return patches, corners, corners_new, x_grid_new, y_grid_new
471
472 # === Patch-Related =====
473 def assign_corners_to_patches(corners, ROI, n_patches):
474
475     # calculates initial boundary points of patches based on region of interest
476     # and number of patches
477
478     patches = []
479
480     # calculate the width and height of each patch
481     patch_width = math.floor(ROI.width / np.sqrt(n_patches))
482     patch_height = math.floor(ROI.height / np.sqrt(n_patches))

```

```

483
484     # assume the number of rows and columns are equal
485     n_cols = np.sqrt(n_patches)
486     n_rows = n_cols
487     row_ct = 0
488     col_ct = 0
489
490     # assign point locations
491     for i in range(n_patches):
492         patches.append( c.Patch() )
493         patches[i].index = i + 1
494         patches[i].corners[0][0] = ROI.x + patch_width * col_ct
495         patches[i].corners[0][1] = ROI.y + patch_height * row_ct
496         patches[i].corners[1][0] = patches[i].corners[0][0] + patch_width
497         patches[i].corners[1][1] = patches[i].corners[0][1]
498         patches[i].corners[2][0] = patches[i].corners[1][0]
499         patches[i].corners[2][1] = patches[i].corners[1][1] + patch_height
500         patches[i].corners[3][0] = patches[i].corners[0][0]
501         patches[i].corners[3][1] = patches[i].corners[0][1] + patch_height
502
503         if col_ct == n_cols-1:
504             col_ct = 0
505             row_ct += 1
506         else:
507             col_ct += 1
508
509     return patches
510
511 # -----
512
513 def assign_patch_points(patches, corners):
514
515     # checks which corners are within the patches
516
517     for idx, point in enumerate(corners):
518         for patch in patches:
519
520             if idx == 0:
521                 patch.point_idx = np.empty([0], dtype='int')
522
523                 point_in_patch = cv2.pointPolygonTest(patch.corners, tuple(point[0]),
524                                                         False)

```

```

525
526     # 1 means inside, 0 means on edge, -1 means outside
527     if point_in_patch == 1 or 0:
528         patch.point_idx = np.append(patch.point_idx, idx)
529
530     return patches
531
532 # -----
533
534 def update_patch_flow(patches, corners, corners_new):
535
536     # updates the optical flow based on the new tracked corner values and checks
537     # constraints
538
539     # initialize flow vectors
540     u_point_all = []
541     v_point_all = []
542
543     # generate list of horizontal and vertical optical flow values for each corner
544     for i, point in enumerate(corners):
545         u_point_all = np.append(u_point_all, corners_new[i][0][0] -
546                                corners[i][0][0])
547         v_point_all = np.append(v_point_all, corners_new[i][0][1] -
548                                corners[i][0][1])
549
550     # calculate parameters for each patch
551     # only update flow if there are sufficient number of samples (10)
552     for patch in patches:
553         # get points within current patch
554         if patch.point_idx.size > 0:
555             u_p_value = 0
556             v_p_value = 0
557
558             u_point_patch = u_point_all[patch.point_idx]
559             v_point_patch = v_point_all[patch.point_idx]
560
561             # calculate patch flow after excluding outliers
562             z_u = np.abs(spstats.zscore(u_point_patch))
563             z_v = np.abs(spstats.zscore(v_point_patch))
564

```

```

565     u_outlier_idx = np.where(z_u < 2)[0]
566     v_outlier_idx = np.where(z_v < 2)[0]
567
568     u_point_patch = u_point_patch[u_outlier_idx]
569     v_point_patch = v_point_patch[v_outlier_idx]
570
571     if len(u_point_patch) > 0:
572         patch.u_avg = np.mean(u_point_patch)
573     else:
574         patch.u_avg = 0
575     if len(v_point_patch) > 0:
576         patch.v_avg = np.mean(v_point_patch)
577     else:
578         patch.v_avg = 0
579
580     # calculate stats if there are enough points for comparison
581     # test for normal distribution
582
583     if u_point_patch.size > 8: # skewtest minimum
584         _, u_p_value = spstats.skewtest(u_point_patch)
585
586     if v_point_patch.size > 8: # skewtest minimum
587         _, v_p_value = spstats.skewtest(v_point_patch)
588
589     alpha = 0.005
590
591     if u_p_value >= alpha and v_p_value >= alpha:
592         patch.normal_corner_dist = True
593     else:
594         patch.normal_corner_dist = False
595
596     # set the patch flow to zero if fewer than specified points in patch
597     else:
598         patch.u_avg = 0
599         patch.v_avg = 0
600         patch.u_std_dev = 0
601         patch.v_std_dev = 0
602         patch.normal_corner_dist = True
603
604     # set patches with zero flow to average neighbor flow after all patch flows
605     # have been calculated
606

```

```

607     # create matrices for u and v avg
608     x_dim = int(np.sqrt(len(patches)))
609     u_avg_all = np.zeros(len(patches))
610     v_avg_all = np.zeros(len(patches))
611
612     for i, patch in enumerate(patches):
613         u_avg_all[i] = patch.u_avg
614         v_avg_all[i] = patch.v_avg
615
616     u_avg_all = np.reshape(u_avg_all, (x_dim, x_dim))
617     v_avg_all = np.reshape(v_avg_all, (x_dim, x_dim))
618
619     # find average of neighbors by convolution
620     u_avg_neighbors = ndimage.generic_filter(u_avg_all, np.nanmean,
        footprint=kernel, mode='constant', cval=np.NaN)
621     v_avg_neighbors = ndimage.generic_filter(v_avg_all, np.nanmean,
        footprint=kernel, mode='constant', cval=np.NaN)
622
623     # flatten for iteration
624     u_avg_neighbors = np.ndarray.flatten(u_avg_neighbors)
625     v_avg_neighbors = np.ndarray.flatten(v_avg_neighbors)
626
627     # set zero flows to neighbor average
628     for i, patch in enumerate(patches):
629         patch.u_avg_neighbors = u_avg_neighbors[i]
630         patch.v_avg_neighbors = v_avg_neighbors[i]
631
632         if patch.u_avg == 0 and patch.v_avg == 0:
633             patch.u_avg = patch.u_avg_neighbors
634             patch.v_avg = patch.u_avg_neighbors
635
636     # check for continuous flow between neighbors
637     for patch in patches:
638         patch.u_diff = patch.u_avg - patch.u_avg_neighbors
639         patch.v_diff = patch.v_avg - patch.v_avg_neighbors
640
641         if patch.u_diff_min <= patch.u_diff <= patch.u_diff_max and \
642             patch.v_diff_min <= patch.v_diff <= patch.v_diff_max:
643
644             patch.continuous_neighbor_flow = True
645         else:
646             patch.continuous_neighbor_flow = False

```

```

647
648     return patches
649
650 # -----
651
652 def create_patches(frame, ROI, n_patch_rows):
653
654 # creates a patch object with corners
655
656     patches = []
657
658     n_patches = int(n_patch_rows**2)
659     # rows, cols are always equal; done for convenience in this project but
660     # could be changed to allow non-square grid; would require more changes to the
        functions
661
662     # detect initial corners
663     mask = np.zeros(frame.shape, dtype=np.uint8)
664     mask = cv2.rectangle(mask, ROI.start_pt, ROI.end_pt, (255,255,255), -1)
665
666     corners = cv2.goodFeaturesToTrack(frame, maxCorners = det_max_corners,
667                                     qualityLevel = det_quality_level,
668                                     minDistance = det_min_distance,
669                                     blockSize = det_block_size,
670                                     useHarrisDetector = det_use_harris,
671                                     mask = mask,
672                                     k = det_k)
673
674     # run function to determine initial patch corner points
675     patches = assign_corners_to_patches(corners, ROI, n_patches)
676
677     # assign points to patches
678     patches = assign_patch_points(patches, corners)
679
680     return patches, corners, n_patches
681
682 # -----
683
684 def move_patch_corners(patches, x_grid, y_grid):
685
686 # updates patch boundaries to match grid
687

```

```

688     n_rows = int(np.sqrt(len(patches)))
689     n_cols = n_rows
690
691     row = 0
692     col = 0
693
694     for patch in patches:
695         patch.corners[0] = [ x_grid[row, col ], y_grid[row, col] ]
696         patch.corners[1] = [ x_grid[row, col+1], y_grid[row, col+1] ]
697         patch.corners[2] = [ x_grid[row+1, col+1], y_grid[row+1, col+1] ]
698         patch.corners[3] = [ x_grid[row+1, col ], y_grid[row+1, col] ]
699
700         if patch.index % n_rows == 0:
701             row +=1
702             col = 0
703         else:
704             col += 1
705
706     return patches
707
708 # -----
709
710 def check_constraints(patches):
711
712 # checks that the conditions for piecewise flow are met
713
714     # initialize constraints
715     constraint_1 = True
716     constraint_2 = True
717
718     # check constraints
719     for patch in patches:
720
721         # 1. all corners within each patch must be normally distributed
722         if patch.normal_corner_dist == False:
723             constraint_1 = False
724             break
725
726         # 2. the average flow measure between neighboring patches is continuous
727         if patch.continuous_neighbor_flow == False:
728             constraint_2 = False
729             break

```

```

730
731     return constraint_1, constraint_2
732
733 # === Grid-Related =====
734
735 def append_grid_sequences(grid_sequences, x_grid, y_grid):
736
737 # adds the current grid points to the sequence history array
738
739     x_grid = np.ndarray.flatten(x_grid)
740     y_grid = np.ndarray.flatten(y_grid)
741
742     for i, sequence in enumerate(grid_sequences):
743         sequence.append(x_grid[i])
744         sequence.append(y_grid[i])
745
746     return grid_sequences
747
748 # -----
749
750 def get_grid_outline( x_grid, y_grid ):
751
752 # returns the outline shape (points) of the grid
753
754     n_rows = x_grid.shape[0]
755     n_cols = x_grid.shape[1]
756
757     pts = []
758
759     # top row
760     for col in range(n_cols):
761         pts.append([x_grid[0, col], y_grid[0, col]])
762     # right col
763     for row in range(n_rows):
764         pts.append([x_grid[row, n_cols-1], y_grid[row, n_cols-1]])
765     # bottom row
766     for col in reversed(range(n_cols)):
767         pts.append([x_grid[n_rows-1, col], y_grid[n_rows-1, col]])
768     # left col
769     for row in reversed(range(n_rows)):
770         pts.append([x_grid[row, 0], y_grid[row, 0]])
771

```



```

772     return pts
773
774 # -----
775
776 def append_grid_prediction_history(x_grid, y_grid, x_grid_ML, y_grid_ML,
777                                 x_grid_hist, y_grid_hist,
778                                 x_grid_ML_hist, y_grid_ML_hist, ):
779
780 # creates history arrays for the ground truth grid and predicted grid
781
782     rows = x_grid.shape[0]
783     cols = x_grid.shape[1]
784
785     x_grid = x_grid.flatten()
786     y_grid = y_grid.flatten()
787     x_grid_ML = x_grid_ML.flatten()
788     y_grid_ML = y_grid_ML.flatten()
789
790     for i in range(x_grid.size):
791         x_grid_hist[i] = np.append(x_grid_hist[i], x_grid[i])
792         y_grid_hist[i] = np.append(y_grid_hist[i], y_grid[i])
793         x_grid_ML_hist[i] = np.append(x_grid_ML_hist[i], x_grid_ML[i])
794         y_grid_ML_hist[i] = np.append(y_grid_ML_hist[i], y_grid_ML[i])
795
796     return x_grid_hist, y_grid_hist, x_grid_ML_hist, y_grid_ML_hist
797
798 # -----
799
800 def create_grid(patch):
801
802 # creates a grid array from the patch boundaries
803
804     # determine size of grid - 1 more dimension than number of patches
805     n_patches = len(patch)
806     n_rows_patch = int(np.sqrt(n_patches))
807     n_cols_patch = n_rows_patch
808     n_rows_grid = n_rows_patch + 1
809     n_cols_grid = n_rows_grid
810     x_grid = np.zeros((n_rows_grid, n_cols_grid), dtype = 'int')
811     y_grid = np.zeros((n_rows_grid, n_cols_grid), dtype = 'int')
812
813     # start counting rows/cols of patch

```

```

814     patch_row = 1
815     patch_col = 1
816
817     # assign grid points based on corners in each patch
818     for i, patch in enumerate(patch_list):
819
820         patch_num = i + 1
821
822         # grid point indices start at 0
823         grid_row = patch_row - 1
824         grid_col = patch_col - 1
825
826         # assign top left point
827         x_grid[grid_row, grid_col] = patch.corners[0][0]
828         y_grid[grid_row, grid_col] = patch.corners[0][1]
829
830         # assign bottom left point on last row
831         if patch_row == n_rows_patch:
832             x_grid[grid_row+1, grid_col] = patch.corners[3][0]
833             y_grid[grid_row+1, grid_col] = patch.corners[3][1]
834
835         # assign top right point on last column
836         if patch_col == n_cols_patch:
837             x_grid[grid_row, grid_col+1] = patch.corners[1][0]
838             y_grid[grid_row, grid_col+1] = patch.corners[1][1]
839
840         # assign bottom left point at end
841         if patch_num == n_patches:
842             x_grid[grid_row+1, grid_col+1] = patch.corners[2][0]
843             y_grid[grid_row+1, grid_col+1] = patch.corners[2][1]
844             break
845
846         # increment counters (patchNum indexed to 0)
847         if not patch_num == n_patches:
848             if patch_num % n_cols_patch == 0:
849                 patch_col = 1
850                 patch_row += 1
851             else:
852                 patch_col += 1
853
854     return x_grid, y_grid
855

```

```

856 # -----
857
858 def update_grid_flow(patches, x_grid, y_grid):
859
860 # moves grid points based on patch movements
861
862     n_patches = len(patches)
863     n_cols = int(np.sqrt(n_patches))
864     n_rows = n_cols
865
866     # "current" patch will be lower left relative to the grid
867     curr_patch = 0
868
869     # move patch corners by average patch U and V flow
870     for row in range(n_rows):
871         for col in range(n_cols):
872
873             neighbors = []
874             u_vals = 0
875             v_vals = 0
876
877             # find neighboring patches (LR, LL, UR, UL)
878
879             # top row left
880             if row == 0 and col == 0:
881                 neighbors.append(curr_patch)
882
883             # top row center
884             elif row == 0 and 0 < col < n_cols:
885                 neighbors.append(curr_patch)
886                 neighbors.append(curr_patch-1)
887
888             # top row right
889             elif row == 0 and col == n_cols:
890                 neighbors.append(curr_patch-1)
891
892             # interior points
893             elif 0 < row < n_rows and 0 < col < n_cols:
894                 neighbors.append(curr_patch)
895                 neighbors.append(curr_patch-1)
896                 neighbors.append(curr_patch-1-n_cols)
897                 neighbors.append(curr_patch-n_cols)

```

```

898
899     # bottom row left
900     elif row == n_rows and col == 0:
901         neighbors.append(curr_patch-n_cols)
902
903     # bottom row center
904     elif row == n_rows and 0 < col < n_cols:
905         neighbors.append(curr_patch-n_cols)
906         neighbors.append(curr_patch-n_cols)
907
908     # bottom row right
909     elif row == n_rows and col == n_cols:
910         neighbors.append(curr_patch-n_cols-1)
911
912     # left column center
913     elif col == 0 and 0 < row < n_rows:
914         neighbors.append(curr_patch)
915         neighbors.append(curr_patch-n_rows)
916
917     # left column right
918     elif col == n_cols and 0 < row < n_rows:
919         neighbors.append(curr_patch-1)
920         neighbors.append(curr_patch-1-n_rows)
921
922     # get average flow
923     for n in neighbors:
924         u_vals += patches[n].u_avg
925         v_vals += patches[n].v_avg
926
927     x_grid[row, col] += u_vals/len(neighbors)
928     y_grid[row, col] += v_vals/len(neighbors)
929
930     curr_patch += 1
931
932     return x_grid, y_grid
933
934 # === Drawing =====
935
936 def rectangle_occlusion(frame, ROI, alpha):
937
938 # adds a rectangle overlay to a frame
939

```

```

940     # convert frame to RGB
941     if len(frame.shape) == 2:
942         frame = cv2.cvtColor(frame,cv2.COLOR_GRAY2RGB)
943
944     overlay = np.zeros(frame.shape, dtype=np.uint8)
945     overlay = cv2.rectangle(overlay, ROI.start_pt, ROI.end_pt, color =
        colors.white, thickness = -1)
946
947     frame = cv2.addWeighted(overlay, alpha, frame, 1, 0, frame)
948
949     return frame
950
951 # -----
952
953 def draw_overlays_main(frame, frame_count, frame_range, n_frames, n_last_frame,
        peaks,
954                     flow_avg_scaled, hr, constraints, basename):
955
956 # draws top and bottom information overlays for main analysis
957
958     frame_height = frame.shape[0]
959     frame_width = frame.shape[1]
960     font_scale = 1
961
962     # apply the transparent header/footer shaded overlay
963     overlay = np.zeros(frame.shape, dtype=np.uint8)
964     overlay = cv2.rectangle(overlay, (0, 660), (frame_width, frame_height), color
        = colors.white, thickness = -1)
965     overlay = cv2.rectangle(overlay, (0, 0), (frame_width, 80), color =
        colors.white, thickness = -1)
966     frame = cv2.addWeighted(overlay, 0.7, frame, 1, 0, frame)
967
968     # scale pulse values and draw to frame
969     t = []
970     flow_avg_scaled_max = max(flow_avg_scaled)
971
972     # had some issues with divide by zero here hence the conditional statements
973     # scale to 5% of frame height then offset to bottom
974     if abs(flow_avg_scaled_max) > 0:
975         flow_avg_scaled = frame_height*0.05 * flow_avg_scaled /
            flow_avg_scaled_max + (frame_height - 50)
976     else:

```

```

977     flow_avg_scaled = frame_height*0.05 * flow_avg_scaled / 0.1 +
          (frame_height - 50)
978
979     for i, point in enumerate(flow_avg_scaled):
980         t_increment = (frame_width * i / n_frames)
981         t = np.append(t, t_increment)
982
983     flow_avg_scaled = np.vstack((t,flow_avg_scaled)).T.astype(np.int)
984
985     frame = cv2.polylines(frame, [flow_avg_scaled], color = colors.red, thickness
          = 2, isClosed = False)
986
987     # plot cycle peaks and sequence numbers
988     if len(peaks) > 0:
989         for i, peak in enumerate(peaks):
990             x = flow_avg_scaled[peak][0]
991             y = flow_avg_scaled[peak][1]
992             frame = cv2.drawMarker(frame, tuple(flow_avg_scaled[peak]),
          colors.magenta, cv2.MARKER_CROSS, 5, 4, cv2.LINE_AA)
993
994     # draw failed frame markers
995     if draw_failed_frames == True:
996         if len(constraints.failed_frames) > 0:
997             for frame_number in constraints.failed_frames:
998                 frame_number_scaled = round((frame_number -
          frame_range[0])/n_frames * frame.shape[0])
999                 cv2.drawMarker(frame, (frame_number_scaled, 670 ), colors.black,
          cv2.MARKER_CROSS, 8, 1, cv2.LINE_AA)
1000
1001         if len(constraints.reinitialized_frames) > 0:
1002             for frame_number in constraints.reinitialized_frames:
1003                 frame_number_scaled = round((frame_number -
          frame_range[0])/n_frames * frame.shape[0])
1004                 cv2.drawMarker(frame, (frame_number_scaled, 680 ), colors.red,
          cv2.MARKER_TRIANGLE_UP, 8, 1, cv2.LINE_AA)
1005
1006     # print header text
1007     # heart rate
1008     text_hr = "HR: " + str(hr) + ' BPM'
1009     frame = cv2.putText(frame, text_hr, (560,30), fontFace =
          cv2.FONT_HERSHEY_DUPLEX,
1010                         fontScale = font_scale, color = colors.dark_green)

```

```

1011
1012 # frame number
1013 text_fn = str(frame_count) + " of " + str(n_last_frame)
1014 frame = cv2.putText(frame, text_fn, (10,30), fontFace =
      cv2.FONT_HERSHEY_DUPLEX,
1015                       fontStyle = font_scale, color = colors.dark_green)
1016
1017 text_tb = "Total beats: " + str(len(peaks))
1018 frame = cv2.putText(frame, text_tb, (500,65), fontFace =
      cv2.FONT_HERSHEY_DUPLEX,
1019                       fontStyle = font_scale, color = colors.dark_green)
1020
1021 text_fn = basename + ".avi"
1022 frame = cv2.putText(frame, text_fn, (10,65), fontFace =
      cv2.FONT_HERSHEY_DUPLEX,
1023                       fontStyle = font_scale, color = colors.dark_green)
1024
1025 return frame
1026
1027 # -----
1028
1029 def draw_overlays_extraction_old(frame, corners_1, corners_2, status_1, status_2,
      frame_width, frame_height,
1030     frame_count, n_frames, peaks, new_peak, flow_avg_scaled,
      basename, hr, ROI_radius):
1031
1032 # draws top and bottom information overlays for feature extraction
1033
1034     frame_height = frame.shape[0]
1035     frame_width = frame.shape[1]
1036     font_scale = 1
1037
1038     # separate points into lost and found
1039     if new_peak == False and len(corners_1) > 0 and len(corners_2) > 0:
1040         status_found_idx_1 = np.where(status_1 == 1)[0]
1041         status_lost_idx_1 = np.where(status_1 == 0)[0]
1042
1043         corners_found_1 = corners_1[status_found_idx_1]
1044         corners_lost_1 = corners_1[status_lost_idx_1]
1045
1046         if len(peaks) > 1:
1047             status_found_idx_2 = np.where(status_2 == 1)[0]

```

```

1048         status_lost_idx_2 = np.where(status_2 == 0)[0]
1049
1050         corners_found_2 = corners_2[status_found_idx_2]
1051         corners_lost_2 = corners_2[status_lost_idx_2]
1052
1053     else:
1054         if len(peaks > 0):
1055             corners_found_1 = corners_1
1056             corners_lost_1 = []
1057
1058             if len(peaks) > 1:
1059                 corners_found_2 = corners_2
1060                 corners_lost_2 = []
1061
1062     # draw corners within ROI
1063     frame = cv2.cvtColor(frame, cv2.COLOR_GRAY2RGB)
1064
1065     if len(peaks) > 0:
1066         if (len(corners_found_1) > 0):
1067             for point in corners_found_1:
1068                 cv2.drawMarker(frame, tuple(point[0]), colors.magenta,
1069                               cv2.MARKER_CROSS, 5, 1, cv2.LINE_AA)
1070
1071         if (len(corners_lost_1) > 0):
1072             for point in corners_lost_1:
1073                 cv2.drawMarker(frame, tuple(point[0]), colors.green,
1074                               cv2.MARKER_CROSS, 20, 2, cv2.LINE_AA)
1075
1076     if len(peaks) > 1:
1077         if (len(corners_found_2) > 0):
1078             for point in corners_found_2:
1079                 cv2.drawMarker(frame, tuple(point[0]), colors.cyan,
1080                               cv2.MARKER_CROSS, 5, 1, cv2.LINE_AA)
1081
1082         if (len(corners_lost_2) > 0):
1083             for point in corners_lost_2:
1084                 cv2.drawMarker(frame, tuple(point[0]), colors.green,
1085                               cv2.MARKER_CROSS, 20, 2, cv2.LINE_AA)
1086
1087     # draw ROI
1088     frame = cv2.circle(frame,
1089                       center = (int(frame_height/2), int(frame_width/2)),

```



```

1086         radius = ROI_radius,
1087         color = colors.green,
1088         thickness = 2)
1089
1090     # apply the header/footer shaded overlay
1091     overlay = np.zeros(frame.shape, dtype=np.uint8)
1092     overlay = cv2.rectangle(overlay, (0, 660), (frame_width, frame_height), color
        = colors.white, thickness = -1)
1093     overlay = cv2.rectangle(overlay, (0, 0), (frame_width, 80), color =
        colors.white, thickness = -1)
1094
1095     alpha = 0.7
1096     frame = cv2.addWeighted(overlay, alpha, frame, 1, 0, frame)
1097
1098
1099     # scale pulse values to frame
1100     t = []
1101     flow_avg_scaled_max = max(flow_avg_scaled)
1102
1103     # had some issues with divide by zero here hence the conditional statements
1104     # scale to 5% of frame height then offset to bottom
1105     if abs(flow_avg_scaled_max) > 0:
1106         flow_avg_scaled = frame_height*0.05 * flow_avg_scaled /
            flow_avg_scaled_max + (frame_height - 50)
1107     else:
1108         flow_avg_scaled = frame_height*0.05 * flow_avg_scaled / 0.1 +
            (frame_height - 50)
1109
1110     for i, point in enumerate(flow_avg_scaled):
1111         t_increment = (frame_width * i / n_frames)
1112         t = np.append(t, t_increment)
1113
1114     flow_avg_scaled = np.vstack((t,flow_avg_scaled)).T.astype(np.int)
1115     frame = cv2.polylines(frame, [flow_avg_scaled], color = colors.red, thickness
        = 2, isClosed = False)
1116
1117     # plot cycle peaks and sequence numbers
1118     if len(peaks) > 0:
1119         for i, peak in enumerate(peaks):
1120             x = flow_avg_scaled[peak][0]
1121             y = flow_avg_scaled[peak][1]
1122

```

```

1123     if i % 2 == 0:
1124         peak_color = colors.magenta
1125         frame = cv2.line(frame, (x, y-10), (x, 670), colors.black, 2)
1126         if i != 0:
1127             frame = cv2.putText(frame, str(int(i-1)), (x-5,795), fontFace =
1128                 cv2.FONT_HERSHEY_DUPLEX,
1129                 fontScale = .5, color = colors.black)
1130             frame = cv2.drawMarker(frame, tuple(flow_avg_scaled[peak]),
1131                 peak_color, cv2.MARKER_CROSS, 5, 4, cv2.LINE_AA)
1132
1133     else:
1134         peak_color = colors.cyan
1135         frame = cv2.line(frame, (x, y+10), (x, 800), colors.black, 2)
1136         frame = cv2.putText(frame, str(int(i-1)), (x-5,680), fontFace =
1137             cv2.FONT_HERSHEY_DUPLEX,
1138             fontScale = .5, color = colors.black)
1139         frame = cv2.drawMarker(frame, tuple(flow_avg_scaled[peak]),
1140             peak_color, cv2.MARKER_CROSS, 5, 4, cv2.LINE_AA)
1141
1142     # print header text
1143     text_fn = str(frame_count) + " of " + str(n_frames)
1144     frame = cv2.putText(frame, text_fn, (10,30), fontFace =
1145         cv2.FONT_HERSHEY_DUPLEX,
1146         fontScale = font_scale, color = colors.dark_green)
1147
1148     text_hr = "HR: " + str(hr) + ' BPM'
1149     frame = cv2.putText(frame, text_hr, (560,30), fontFace =
1150         cv2.FONT_HERSHEY_DUPLEX,
1151         fontScale = font_scale, color = colors.dark_green)
1152
1153     text_tb = "Total beats: " + str(len(peaks))
1154     frame = cv2.putText(frame, text_tb, (500,65), fontFace =
1155         cv2.FONT_HERSHEY_DUPLEX,
1156         fontScale = font_scale, color = colors.dark_green)
1157
1158     text_fn = basename + ".mat"
1159     frame = cv2.putText(frame, text_fn, (10,65), fontFace =
1160         cv2.FONT_HERSHEY_DUPLEX,
1161         fontScale = font_scale, color = colors.dark_green)
1162
1163     return frame

```

```

1157
1158 # -----
1159
1160 def draw_overlays_extraction(frame, corners_1, status_1, frame_width, frame_height,
1161                             frame_count, n_frames, peaks, new_peak, flow_avg_scaled,
1162                             basename, hr, ROI_radius):
1163 # draws top and bottom information overlays for feature extraction
1164
1165     frame_height = frame.shape[0]
1166     frame_width = frame.shape[1]
1167     font_scale = 1
1168
1169     # separate points into lost and found
1170     if len(corners_1) > 0 and new_peak == False:
1171         status_found_idx_1 = np.where(status_1 == 1)[0]
1172         status_lost_idx_1 = np.where(status_1 == 0)[0]
1173
1174         corners_found_1 = corners_1[status_found_idx_1]
1175         corners_lost_1 = corners_1[status_lost_idx_1]
1176
1177     else:
1178         if len(peaks > 0):
1179             corners_found_1 = corners_1
1180             corners_lost_1 = []
1181
1182     # draw corners within ROI
1183     frame = cv2.cvtColor(frame, cv2.COLOR_GRAY2RGB)
1184
1185     if len(peaks) > 0:
1186         if (len(corners_found_1) > 0):
1187             for point in corners_found_1:
1188                 cv2.drawMarker(frame, tuple(point[0]), colors.magenta,
1189                               cv2.MARKER_CROSS, 5, 1, cv2.LINE_AA)
1189
1190         if (len(corners_lost_1) > 0):
1191             for point in corners_lost_1:
1192                 cv2.drawMarker(frame, tuple(point[0]), colors.green,
1193                               cv2.MARKER_CROSS, 20, 2, cv2.LINE_AA)
1193
1194     # draw ROI
1195     frame = cv2.circle(frame,

```

```

1196         center = (int(frame_height/2), int(frame_width/2)),
1197         radius = ROI_radius,
1198         color = colors.green,
1199         thickness = 2)
1200
1201     # apply the header/footer shaded overlay
1202     overlay = np.zeros(frame.shape, dtype=np.uint8)
1203     overlay = cv2.rectangle(overlay, (0, 660), (frame_width, frame_height), color
        = colors.white, thickness = -1)
1204     overlay = cv2.rectangle(overlay, (0, 0), (frame_width, 80), color =
        colors.white, thickness = -1)
1205
1206     alpha = 0.7
1207     frame = cv2.addWeighted(overlay, alpha, frame, 1, 0, frame)
1208
1209     # scale pulse values to frame
1210     t = []
1211     flow_avg_scaled_max = max(flow_avg_scaled)
1212
1213     # had some issues with divide by zero here hence the conditional statements
1214     # scale to 5% of frame height then offset to bottom
1215     if abs(flow_avg_scaled_max) > 0:
1216         flow_avg_scaled = frame_height*0.05 * flow_avg_scaled /
            flow_avg_scaled_max + (frame_height - 50)
1217     else:
1218         flow_avg_scaled = frame_height*0.05 * flow_avg_scaled / 0.1 +
            (frame_height - 50)
1219
1220     for i, point in enumerate(flow_avg_scaled):
1221         t_increment = (frame_width * i / n_frames)
1222         t = np.append(t, t_increment)
1223
1224     flow_avg_scaled = np.vstack((t, flow_avg_scaled)).T.astype(np.int)
1225     frame = cv2.polylines(frame, [flow_avg_scaled], color = colors.red, thickness
        = 2, isClosed = False)
1226
1227     # plot cycle peaks and sequence numbers
1228     if len(peaks) > 0:
1229         for i, peak in enumerate(peaks):
1230             x = flow_avg_scaled[peak][0]
1231             y = flow_avg_scaled[peak][1]
1232

```

```

1233     peak_color = colors.magenta
1234     frame = cv2.drawMarker(frame, tuple(flow_avg_scaled[peak]), peak_color,
                             cv2.MARKER_CROSS, 5, 4, cv2.LINE_AA)
1235
1236     # print header text
1237     text_fn = str(frame_count) + " of " + str(n_frames)
1238     frame = cv2.putText(frame, text_fn, (10,30), fontFace =
                           cv2.FONT_HERSHEY_DUPLEX,
1239                           fontScale = font_scale, color = colors.dark_green)
1240
1241     text_hr = "HR: " + str(hr) + ' BPM'
1242     frame = cv2.putText(frame, text_hr, (560,30), fontFace =
                           cv2.FONT_HERSHEY_DUPLEX,
1243                           fontScale = font_scale, color = colors.dark_green)
1244
1245     text_tb = "Total beats: " + str(len(peaks))
1246     frame = cv2.putText(frame, text_tb, (500,65), fontFace =
                           cv2.FONT_HERSHEY_DUPLEX,
1247                           fontScale = font_scale, color = colors.dark_green)
1248
1249     text_fn = basename + ".avi"
1250     frame = cv2.putText(frame, text_fn, (10,65), fontFace =
                           cv2.FONT_HERSHEY_DUPLEX,
1251                           fontScale = font_scale, color = colors.dark_green)
1252
1253     return frame
1254
1255     # -----
1256
1257     def draw_grid_lines(frame, x_grid, y_grid, color):
1258
1259         # draw gridlines on frame
1260
1261         # plot grid lines over an image
1262         n_cols = len(x_grid)
1263         n_rows = n_cols
1264
1265         if n_cols > 1:
1266
1267             for row in range(n_rows):
1268                 for col in range(n_cols):
1269

```

```

1270         # draw vertical lines
1271     if row < n_rows - 1:
1272         cv2.line(frame,
1273                 tuple( (x_grid[row, col], y_grid[row, col]) ),
1274                 tuple( (x_grid[row+1, col], y_grid[row+1, col]) ),
1275                 color)
1276
1277         # draw left/right outline
1278     if col == 0 or col == n_cols - 1:
1279         cv2.line(frame,
1280                 tuple( (x_grid[row, col], y_grid[row, col]) ),
1281                 tuple( (x_grid[row+1, col], y_grid[row+1, col]) ),
1282                 color, grid_outline_thickness)
1283
1284
1285         # draw horizontal lines
1286     if col < n_cols - 1:
1287         cv2.line(frame,
1288                 tuple( (x_grid[row, col], y_grid[row, col]) ),
1289                 tuple( (x_grid[row, col+1], y_grid[row, col+1]) ),
1290                 color)
1291
1292         # draw top/bottom outline
1293     if row == 0 or row == n_rows - 1:
1294         cv2.line(frame,
1295                 tuple( (x_grid[row, col], y_grid[row, col]) ),
1296                 tuple( (x_grid[row, col+1], y_grid[row, col+1]) ),
1297                 color, grid_outline_thickness)
1298
1299     return frame
1300
1301 # -----
1302
1303 def create_grid_plots(x_grid_hist, y_grid_hist, x_grid_ML_hist, y_grid_ML_hist,
1304                     basename, output_path, text_output_file):
1305     # calculations
1306
1307     # absolute distances
1308     x_grid_hist = np.asarray(x_grid_hist)
1309     y_grid_hist = np.asarray(y_grid_hist)
1310     x_grid_ML_hist = np.asarray(x_grid_ML_hist)

```

```

1311 y_grid_ML_hist = np.asarray(y_grid_ML_hist)
1312
1313 x_grid_diff_abs = np.abs(x_grid_hist - x_grid_ML_hist)
1314 y_grid_diff_abs = np.abs(y_grid_hist - y_grid_ML_hist)
1315 grid_diff_abs = np.zeros_like(x_grid_diff_abs)
1316
1317 for i in range(len(x_grid_hist)):
1318     for j in range(len(x_grid_hist[0])):
1319         point1 = (x_grid_hist[i][j], y_grid_hist[i][j])
1320         point2 = (x_grid_ML_hist[i][j], y_grid_ML_hist[i][j])
1321         grid_diff_abs[i][j] = np.abs(distance.euclidean(point1, point2))
1322
1323 x_max_idx = np.argmax(x_grid_hist, axis=0)
1324
1325 x_grid_avg_diff_abs = np.mean(x_grid_diff_abs, axis=0)
1326 x_grid_max_diff_abs = np.max(x_grid_diff_abs, axis=0)
1327 x_grid_min_diff_abs = np.min(x_grid_diff_abs, axis=0)
1328
1329 y_grid_avg_diff_abs = np.mean(y_grid_diff_abs, axis=0)
1330 y_grid_max_diff_abs = np.max(y_grid_diff_abs, axis=0)
1331 y_grid_min_diff_abs = np.min(y_grid_diff_abs, axis=0)
1332
1333 grid_avg_diff_abs = np.mean(grid_diff_abs, axis=0)
1334 grid_max_diff_abs = np.max(grid_diff_abs, axis=0)
1335 grid_min_diff_abs = np.min(grid_diff_abs, axis=0)
1336
1337 f_nos = np.arange(1, len(x_grid_avg_diff_abs)+1)
1338
1339 yerr_x_abs = [np.abs(x_grid_min_diff_abs-x_grid_avg_diff_abs),
1340              np.abs(x_grid_max_diff_abs-x_grid_avg_diff_abs)]
1341 yerr_y_abs = [np.abs(y_grid_min_diff_abs-y_grid_avg_diff_abs),
1342              np.abs(y_grid_max_diff_abs-y_grid_avg_diff_abs)]
1343 yerr_abs = [np.abs(grid_min_diff_abs-grid_avg_diff_abs),
1344            np.abs(grid_max_diff_abs-grid_avg_diff_abs)]
1345
1346 # plot grid differences - average
1347 fig, axs = plt.subplots(3, figsize=(8.5,5.5))
1348 title_string = 'Average Grid Prediction Error (pixels) \n (' + basename +
1349               '.avi)'
1350 fig.suptitle(title_string)
1351 axs[0].errorbar(f_nos, x_grid_avg_diff_abs, yerr = yerr_x_abs, label =
1352               'x-diff', linewidth=2, elinewidth = 1, capsize = 3)

```

```

1348     axs[0].plot(f_nos, x_grid_avg_diff_abs, 'ok')
1349     axs[0].set_xticks(f_nos, minor=False)
1350     axs[0].tick_params(labelsize=13)
1351     axs[0].legend(loc = 2, ncol = 1)
1352     axs[0].grid(color = 'lightgray', linestyle = '--', linewidth = 1)
1353
1354     axs[1].errorbar(f_nos, y_grid_avg_diff_abs, yerr = yerr_y_abs, label =
        'y-diff', linewidth=2, elinewidth = 1, capsize = 3)
1355     axs[1].plot(f_nos, y_grid_avg_diff_abs, 'ok')
1356     axs[1].set_xticks(f_nos, minor=False)
1357     axs[1].tick_params(labelsize=13)
1358     axs[1].legend(loc = 2, ncol = 1)
1359     axs[1].grid(color = 'lightgray', linestyle = '--', linewidth = 1)
1360
1361     axs[2].errorbar(f_nos, grid_avg_diff_abs, yerr = yerr_abs, label = 'xy-diff',
        linewidth=2, elinewidth = 1, capsize = 3)
1362     axs[2].plot(f_nos, grid_avg_diff_abs, 'ok')
1363     axs[2].set_xticks(f_nos, minor=False)
1364     axs[2].tick_params(labelsize=13)
1365     axs[2].legend(loc = 2, ncol = 1)
1366     axs[2].grid(color = 'lightgray', linestyle = '--', linewidth = 1)
1367
1368     fig.text(0.5, 0.01, 'Predicted Frame Number', ha='center', fontsize=13)
1369     fig.text(0.01, 0.5, 'Difference from Ground Truth (pixels)', va='center',
        rotation='vertical', fontsize=13)
1370
1371     fig.subplots_adjust(hspace=0.8, left=0.1, right=0.99)
1372
1373     plt.savefig(output_path + 'plots/' + 'grid_diff_avg_' + basename + '.png')
1374     plt.show()
1375
1376     # plot grid differences - histogram - first frame
1377     n_bins = 10
1378     plt.hist(x_grid_diff_abs[:,0], bins='auto', facecolor='blue', alpha=0.5,
        edgecolor='black')
1379     title_string = 'First Frame Differences (pixels) \n (' + basename + '.avi)'
1380     plt.title(title_string)
1381     plt.xlabel('Difference from Ground Truth (pixels)')
1382     plt.ylabel('Number of Values')
1383     plt.savefig(output_path + 'plots/' + 'grid_diff_hist_patches_' + basename +
        '.png')
1384     plt.show()

```



```

1385
1386 # plot grid differences - boxplot - patches
1387 patch_nos = np.arange(x_grid_diff_abs.shape[0])+1
1388 plt.boxplot(np.transpose(x_grid_diff_abs), meanline = True)
1389 title_string = 'Per Patch Grid Prediction Error (pixels) \n (' + basename +
    '.avi)'
1390 plt.title(title_string)
1391 plt.grid(color = 'lightgray', linestyle = '--', linewidth = 1)
1392 plt.xlabel('Patch Number')
1393 plt.ylabel('Difference from Ground Truth (pixels)')
1394 plt.xticks(patch_nos,patch_nos-1)
1395 plt.savefig(output_path + 'plots/' + 'grid_diff_patches_' + basename + '.png')
1396 plt.show()
1397
1398 f1 = 0
1399 f3 = len(grid_avg_diff_abs)-1
1400 f2 = int(f3/2)
1401
1402 fout = open(text_output_file,'a')
1403 print("\n", file=fout)
1404 print("x-y Distances", file=fout)
1405 print("  F", f1, "|F", f2, "|F", f3, file=fout)
1406 print("Min:", round(grid_min_diff_abs[f1], 1), round(grid_min_diff_abs[f2],1),
    round(grid_min_diff_abs[f3],1), file=fout)
1407 print("Avg:", round(grid_avg_diff_abs[f1], 1), round(grid_avg_diff_abs[f2],1),
    round(grid_avg_diff_abs[f3],1), file=fout)
1408 print("Max:", round(grid_max_diff_abs[f1], 1), round(grid_max_diff_abs[f2],1),
    round(grid_max_diff_abs[f3],1), file=fout)
1409 print(" ")
1410 fout.close()
1411
1412 return
1413
1414 # -----
1415
1416 def create_timeline_plots(constraints, n_last_frame, diastole_frames,
    main_tracking_start_frame, occluded_frames, basename, output_path):
1417
1418 # calculations
1419 predicted_frames = np.arange(occluded_frames[0], occluded_frames[1])
1420 frame_numbers = np.arange(0, n_last_frame)
1421

```

```

1422     diastole_frames_y = np.empty(len(diastole_frames))
1423     diastole_frames_y.fill(2)
1424     reinitialized_frames_y = np.empty(len(constraints.reinitialized_frames))
1425     reinitialized_frames_y.fill(3)
1426     failed_frames_y = np.empty(len(constraints.failed_frames))
1427     failed_frames_y.fill(4)
1428     predicted_frames_y = np.empty(len(predicted_frames))
1429     predicted_frames_y.fill(5)
1430
1431     # plot timeline
1432     plt.figure(figsize=(8, 3))
1433     plt.plot(main_tracking_start_frame, 1, '|k', markersize = 10)
1434     plt.plot(diastole_frames, diastole_frames_y, '|r', markersize = 10)
1435     plt.plot(constraints.reinitialized_frames, reinitialized_frames_y, '|m',
1436             markersize = 10)
1437     plt.plot(constraints.failed_frames, failed_frames_y, '|b', markersize = 10)
1438     plt.plot(predicted_frames, predicted_frames_y, '|y', markersize = 10)
1439
1440     title_string = 'Analysis Timeline \n (' + basename + '.avi)'
1441     plt.title(title_string)
1442     plt.xlabel('Frame Number', fontsize = 12)
1443     ax = plt.gca()
1444     ax.xaxis.grid()
1445     plt.xticks(fontsize=12)
1446     ax.set_xlim(0, n_last_frame)
1447     ax.set_ylim(0, 6)
1448     ax.set_yticks([1, 2, 3, 4, 5])
1449     ax.set_yticklabels(['Tracking Start', 'Diastole Frames', 'Tracking Resets',
1450                       'Failed Constraints', 'Predicted Frames'], fontsize = 12)
1451     plt.tight_layout()
1452     plt.savefig(output_path + 'plots/' + 'timeline_' + basename + '.png',
1453               bbox_inches='tight')
1454     plt.show()
1455
1456     return
1457
1458 # -----
1459 def write_frame(frame, ROI, patches, corners, contours, hierarchy, vid, x_grid,
1460               y_grid, x_grid_ML, y_grid_ML, frame_count, frame_range,
1461               occluded_frames,

```

```

1459         n_frames, n_last_frame, peaks, flow_avg_scaled, hr, constraints,
           basename):
1460
1461 # draws on current frame and writes to file
1462
1463 # convert frame to RGB
1464 if len(frame.shape) == 2:
1465     frame_RGB = cv2.cvtColor(frame, cv2.COLOR_GRAY2RGB)
1466 else:
1467     frame_RGB = frame
1468
1469 # draw overlays including pulse
1470 if draw_overlays == True:
1471     frame_RGB = draw_overlays_main(frame_RGB, frame_count, frame_range,
           n_frames, n_last_frame, peaks, flow_avg_scaled, hr, constraints,
           basename)
1472
1473 # draw corner points
1474 if draw_corners == True:
1475     # draw corners within ROI
1476     if len(corners) > 0:
1477         for point in corners:
1478             cv2.drawMarker(frame_RGB, tuple(point[0]), colors.magenta,
           cv2.MARKER_CROSS, 5, 1, cv2.LINE_AA)
1479
1480 # draw contours
1481 if draw_contours == True:
1482     if len(contours) > 0:
1483         frame_RGB = cv2.drawContours(frame_RGB, contours, -1, hierarchy =
           hierarchy, \
1484                                     maxLevel = contours_max_level, color =
           colors.yellow, thickness = 1)
1485
1486 # draw patch corners
1487 if draw_patch_corners == True:
1488     for i in range(0, len(patches)):
1489         cv2.line(frame_RGB, tuple(patches[i].corners[0]),
           tuple(patches[i].corners[1]), colors.yellow, 1, 1)
1490         cv2.line(frame_RGB, tuple(patches[i].corners[1]),
           tuple(patches[i].corners[2]), colors.yellow, 1, 1)
1491         cv2.line(frame_RGB, tuple(patches[i].corners[2]),
           tuple(patches[i].corners[3]), colors.yellow, 1, 1)

```

```
1492         cv2.line(frame_RGB, tuple(patches[i].corners[3]),
1493                  tuple(patches[i].corners[0]), colors.yellow, 1, 1)
1494     # draw grids
1495     if draw_grid == True:
1496         frame_RGB = draw_grid_lines(frame_RGB, x_grid, y_grid, grid_color_normal)
1497
1498     if draw_grid_ML == True and occluded_frames[0] <= frame_count <=
1499         occluded_frames[1]:
1500         frame_RGB = draw_grid_lines(frame_RGB, x_grid_ML, y_grid_ML, grid_color_ML)
1501     # draw ROI
1502     if draw_ROI == True:
1503         cv2.rectangle(frame_RGB, ROI.start_pt, ROI.end_pt, colors.green, thickness
1504                       = 2)
1505     # write frame to file - needs RGB
1506     vid.write(frame_RGB)
1507
1508     return frame_RGB
```

A.6 MAT to AVI Conversion

```
1 % =====
2 %
3 % file:    mat2avi.m
4 % version: v1.0.0
5 % author:  Billy Hempstead
6 %
7 % summary: Converts a .mat video file to a .avi file
8 %
9 % =====
10
11 clear all
12 close all
13 clc
14
15 frameskip = 2;
16
17 % select and load data file
18 load_msg = 'Loading Datafile';
19 disp(load_msg)
20
21 [filename, pathname] = uigetfile('*.mat');
22 fullpath = strcat(pathname,filename);
23 load(fullpath)
24
25 % setup output video
26 output_file = strcat(pathname,filename(1:end-4),'.avi');
27 output_vid = VideoWriter(output_file, 'Grayscale AVI');
28 fps = IDS.CAM.FPS;
29 output_vid.FrameRate = fps/frameskip;
30 open(output_vid);
31
32 % write frames to output video
33 n_frames = length(IDS.ImgVIS);
34
35 for i = 1:frameskip:n_frames
36     clc
37     update_msg = strcat({'Processing frame '}, num2str(i), {' of '},
38         num2str(n_frames));
39     disp(update_msg)
40     frame = IDS.ImgVIS(:,:,i);
```

```
40     writeVideo(output_vid,frame);  
41 end  
42  
43 close(output_vid);
```
