AUTOMATED PREDICTION OF HEPATIC ARTERIAL STENOSIS

A THESIS IN
Computer Science

Presented to the Faculty of the University
of Missouri-Kansas City in partial fulfillment of
the requirements for the degree

MASTER OF SCIENCE

by
JUSTIN JAY LOUIS BARABOO
B.S., Truman State University, 2015

Kansas City, Missouri
2017

AUTOMATED PREDICTION OF HEPATIC ARTERIAL STENOSIS

Justin Jay Louis Baraboo, Candidate for the Master of Science Degree

University of Missouri-Kansas City, 2017

ABSTRACT

Several thousand life-saving liver transplants are performed each year. One of the causes of early transplant failure is arterial stenosis of the anastomotic junction. Early detection of transplant arterial stenosis can help prevent transplant failure and the need to re-transplant. Doppler ultrasound with manual measurements is the most common screening method, but it suffers from poor specificity when thresholded to reduce false negatives. Positive screening cases proceed to angiography, which is an invasive and expensive procedure. A more accurate test could decrease the number of normal patients who would have to undergo this invasive diagnostic procedure. Machine learning models have shown promise in determining stenosis in the carotid artery; however, they have yet to be tested on the less ideal data hepatic arteries generate. Software has been created to extract liver artery Doppler ultrasound information in an automated fashion to predict stenosis. A turnkey approach is utilized to refine the region prior to extraction. Current methods of extraction generate waveforms with an average percent error per pixel of 6.5 percent from a human drawn waveform. Single feature models and machine learning models performed similarly when predicting stenosis; however, when thresholded for high sensitivity (greater than 0.90), random forest models had the highest specificity at 1.0 sensitivity and 0.60 specificity.

APPROVAL PAGE

The faculty listed below, appointed by the Dean of the School of Computing and Engineering have examined a thesis titled "Automated Prediction of Hepatic Arterial Stenosis," presented by Justin J. Baraboo, candidate for the Master of Science degree, and certify that in their opinion it is worthy of acceptance.

Supervisory Committee

Deendayal Dinakarpandian, M.D., Ph.D., Committee Chair
Department of Computer Science

Sherwin Chan, M.D., Ph.D.
Radiologist at Children's Mercy Hospital

Appie van de Liefvoort, Ph.D.
Department of Computer Science

CONTENTS

LIST OF ILLUSTRATIONS

Figure Page

LIST OF ILLUSTRATIONS

Figure Page

CHAPTER 1

INTRODUCTION

Chronic liver disease, including cirrhosis, is one of the top 10 causes of death. It was ranked as the 8th and 10th leading cause of death in 1980 and 2014 respectively, with 24,584 deaths in 2014 [11]. Cirrhosis is the commonest indication for a liver transplant. According to the Liver Foundation [10], 6000 transplants are performed every year, with 16,000 on the waiting list. One of the most common early causes of hepatic transplant failure is stenosis of the arterial anastomosis between donor and recipient. If detected early, the stenosis can be treated, thereby preventing failure of a precious organ transplant. An illustration depicting the concept of stenosis is shown in Figure 1 [1]. The definitive diagnosis of stenosis is by conventional angiography. However, angiography is expensive, invasive (requiring puncture of the femoral artery), and the imaging contrast agent can be nephrotoxic and/or trigger immune reactions. Ultrasound imaging of blood flow based on the Doppler effect is a safe alternative for screening for stenosis. A typical Doppler sonogram is shown in Figure 2. The informative region for detecting stenosis of the Doppler ultrasound is the waveform; the blood flow through the artery (y axis) is measured as time progresses (x-axis). Three arteries, portal hepatic artery, left common hepatic artery, and right common hepatic artery, are typically screened, producing

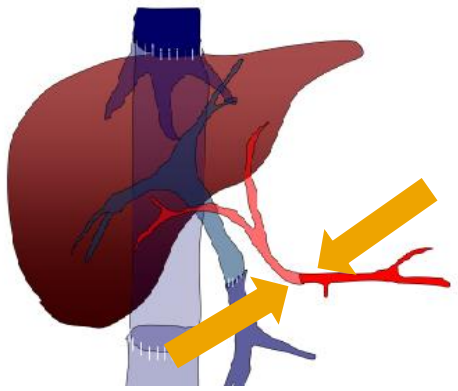

Figure 1. Stenosis is the narrowing of an artery. A stenosis between a donor and recipient artery can indicate liver transplant failure [1].

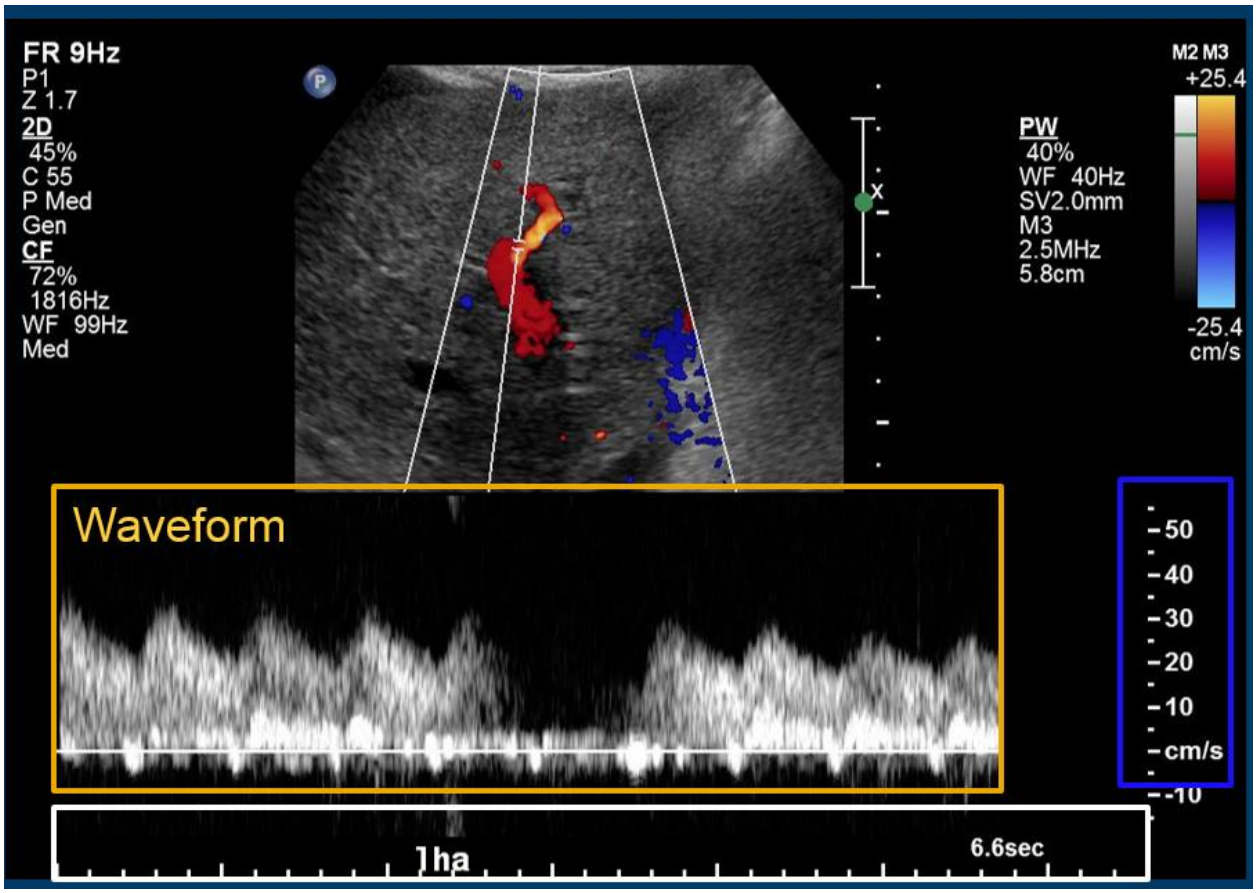

Figure 2 A typical Doppler ultrasound image. The waveform is shown bounded by the orange rectangle. The blue rectangle encloses y-axis of blood flow velocity and the white box encloses the x-axis of time.

images for each. The location of a hepatic artery can be difficult to find and patient movement can induce gaps in the Doppler ultrasound. Also, the refraction of the ultrasound can induce noise above the waveform, which can become severe in some instances. These concepts are illustrated in Figure 3. While resampling an image can be done to reduce these, generating a waveform free from any aberrations could be an intractable task for a given patient. As such, features often need to be extracted from an imperfect waveform. Features must be able to be thresholded to reduce the misclassification of stenosis for patients that have it; false negatives must be minimized, as in these cases, the patient's life and transplant are in danger. The classic features to determine stenosis only need a technologist to plot two points, the peak systolic and

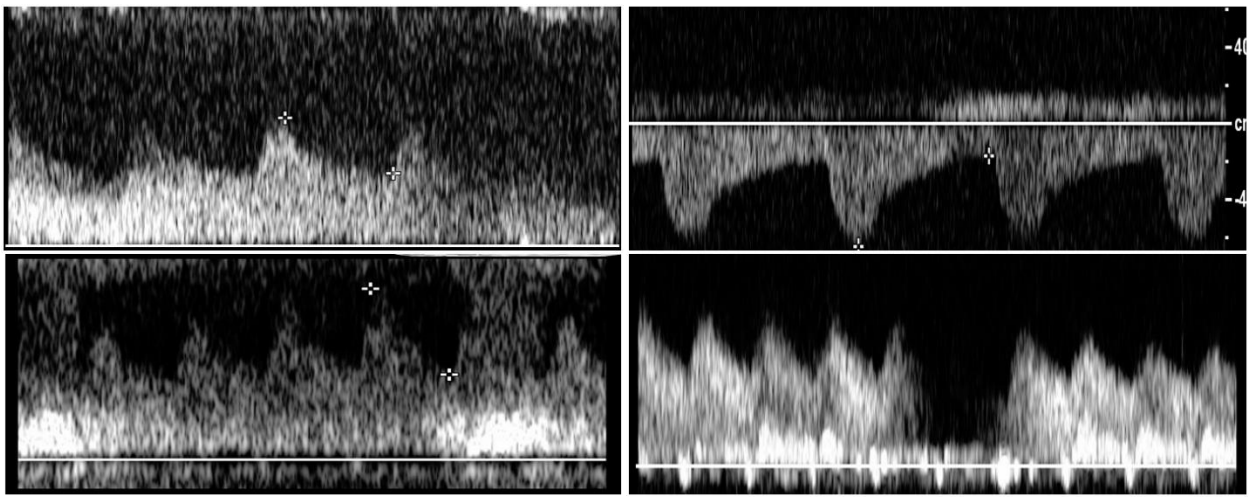

Figure 3. High noise (top left), malformations (bottom left), gaps (bottom right), and upside down waveforms (top right) impact the ability to extract a waveform easily.

diastolic velocities, on the Doppler ultrasound to compute them. Different ratios based on systolic and diastolic peak blood flow velocities, for example the resistive index (RI), have been used to predict the presence of stenosis. However, when these features minimize false negatives, the false positives garnered from a conservative classification becomes high, reducing the power of the screening process and leading to unnecessary angiography. One reason for their limited predictive accuracy is the fact that they are based on lossy representations that sample blood flow only at a few points during each cycle. It has been shown that Fourier analysis of the frequency dependent power distribution of blood flow downstream of the transplant anastomosis can be used to compute a stenosis index (SI) [3]. The SI measure takes into consideration the entire temporal profile of blood flow, which utilizes the information of the entire waveform, allowing a more information rich feature. The SI measure is currently calculated using MATLAB scripts that require intermediate decisions by a domain expert. In order to best be used in a clinical setting, this process of refinement, waveform extraction, and feature extraction needs to be automated. Machine learning models have had success in classifying stenosis in the carotid artery by using the temporal and spectral features of the waveform. This is achieved with high

sensitivity and specificity, underscoring the power of combining features together to generate a model. It is unknown how machine learning models perform for the hepatic artery for classifying stenosis though, which could be due to the imperfect data encountered with Doppler ultrasound of hepatic arteries. Thus, automation of the extraction of the waveform and features could expand further research of machine learning model building into the hepatic domain as well. So, the automated extraction of a hepatic Doppler ultrasound would ultimately allow the stenosis index to be calculated in an automated fashion, fitting better within a clinical setting, and allow machine learning models that have worked with the carotid artery which utilize much of the waveform envelope to be utilized and explored. This, in turn, would allow a better classification and prediction of stenosis, reducing the amount of unnecessary angiography while still maintaining the high sensitivity necessary to be a screening method.

CHAPTER 2

LITERATURE REVIEW

Hepatic Arterial Stenosis Detection

Hepatic Doppler ultrasound is primarily used as a screening method for detecting arterial stenosis following a liver transplant. Different indicators that a transplant is failing, such as increased liver enzymes in the blood stream, give rise to a Doppler ultrasound to be performed to see if the definitive detection method and treatment, angiography, is necessary if stenosis is present. As a screening tool, the sensitivity of classification must be high. While angiography is expensive and invasive, misclassifying someone who has stenosis as not having stenosis, a false negative, is deadly and defeats the purpose of screening. As such, the classifiers for stenosis that have emerged from Doppler imaging tend to be liberal with their false positives to ensure that mostly everyone who does have stenosis are actually classified correctly. However, the false positive rate of models can be very high, referring many patients to angiography who don't need it. Classical features tend to use only two points from the waveform to diagnosis stenosis or no stenosis from the waveform. The classical features for stenosis detection include peak systolic velocity [2], the resistive index [5], the pulsatility index [15], and the acceleration time [5]. These features are mathematically defined and illustrated in Figure 4. The peak systolic velocity of the blood flow through the artery is an indicator of stenosis as arteries that have higher stenosis will have higher and sharper peaks, thus a higher peak systolic velocity, until total occlusion where the flow stops entirely [2]. This index requires the Doppler ultrasound be to taken near the point of stenosis. Achieving this, with a threshold of diagnosing stenosis if the peak systolic velocity is greater than 2 m/s, it achieves a sensitivity of 0.73 and specificity of

Index	Equation
Peak Systolic Velocity	P
Resistive Index	$\frac{P-S}{P}$
Pulsatility Index	$\frac{P-S}{Mean}$
Acceleration Time	t

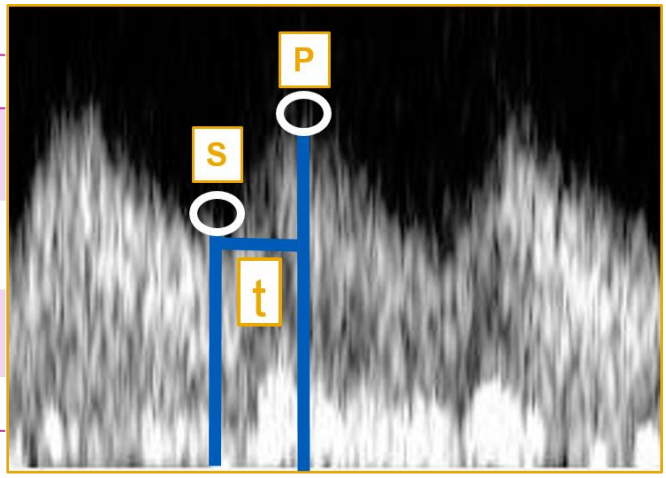

Figure 4. The classical features for arterial stenosis prediction. These features are defined for the peak systolic velocity, P, the trough systolic velocity, S, and the time it takes to reach the peak from the trough, t.

0.50-0.90 [3,6,13,20]. The resistive index ratios the difference of the peak systolic velocity and the trough systolic velocity and normalizes it by the peak systolic velocity. The resistive index has achieved around a 0.60 sensitivity and 0.78 specificity of stenosis detection [5,16,21,25]. The pulsatility index similarly finds the difference between the peak and trough systolic velocities and normalizes it by the mean of the waveform instead of the peak. The acceleration time measures the time taken from the end of a trough to the start of the peak. As a single predictor, declaring stenosis below a 1.0 second threshold, it performs with 0.73 sensitivity and 1.00 specificity [6]. As a time feature, it can be combined with the other blood flow velocity features to improve prediction [5]. These features perform acceptably when detecting mild to severe stenosis, but may fail to detect narrowing that can progress to stenosis. Ultimately, these features only use one to two points of the entire waveform to make a diagnosis, wasting information. They also highly correlate with each other as features pertaining to peak, trough, or time. This reduces the power of creating models with these features together as more independent features yield better informative models. Using more information from the waveform could improve predictive power by either using more of the waveform envelope itself

or looking at spectral features as well, which machine learning models of the carotid artery do. This information comes at a price as it would require more information to be extracted from the waveform itself, requiring more than two points to be accurately estimated from the waveform.

Currently, a technologist pinpoints two different points of the waveform in the Doppler image to generate the previously mentioned features; these pinpoints can be seen as the white cross hairs in the images in Figure 3. Human identification of healthy artery waveforms can differ in estimating both the location in the time direction (x axis) and velocity direction (y axis) due to human variation alone. This can give up to a 39 msec variation for identifying the acceleration time and a 0.08 variation in resistive index [14]. Here, 6% of the resistive index variation, 0.02 standard deviation, could be explained by different sonographers while 17% of the acceleration time variation, 16.2 msec standard deviation, could be explained by different sonographers. Other variation could be explained by variation of the waveform within itself as well as other things, such as patient waveform variation among ultrasounds. This difference can change classification of stenosis, as the threshold for the resistive index can be taken to be below 0.5 to 0.6 for screening purposes [5] while a healthy artery ranges from 0.6 to 0.7, meaning that many healthy individuals near the lower bound could be misdiagnosed due to technologist variability. Also, this could extend to arteries with stenosis being misdiagnosed through technologist variability (as only healthy patients were only used in that study). Automation could help with reducing this variability as it performs the same technique for classification each time. This idea can be seen as some images have a vague “correct peak” as the waveform could be compressing or highly distorted as shown in Figure 5. In these cases, a clear true peak isn’t readily evident in the image, so using a consistent algorithm to determine which peak to use could reduce the human induced variation.

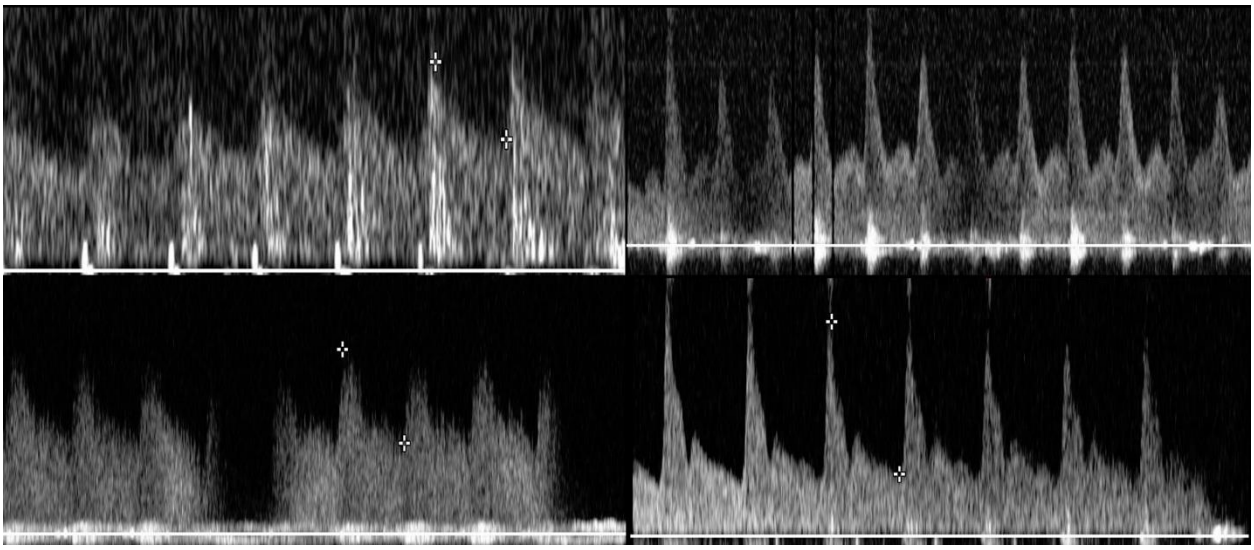

Figure 5. Peak variation with a waveform can be described by distortion (top left), natural peak variation (top right), sparsity of peaks (bottom left), and compressing peaks (bottom right). This variation can lead to misclassification or existence of a true peak, affecting classical features.

Furthermore, the use of classical features in other arteries outside livers is dubious. While renal transplant research has made uses of classical indexes with mild success [3,6,13,19], primarily the resistive index, other renal transplant studies have doubted its efficacy in both how it performs and what it assumes to measure [18,21]. It would be powerful to have a feature that can describe arterial stenosis regardless of arterial location and transplant.

Stenosis Index

The stenosis index is a measure that uses information of entire waves of the waveform to determine whether stenosis has occurred. The premise for the index is that healthy patient waveforms will be different than a patient with stenosis. Instead of looking at only two points of the wave, it breaks the waves down into their frequency components through a Fourier transform. Healthy waveforms have a greater proportion of higher frequency components compared to the fundamental frequency than patients with stenosis. By looking at the frequency components of the waveform, the ratio of the higher frequency components to the lower fundamental frequency can be used to determine the presence of stenosis. The stenosis index

employs a low pass filter to reduce noise; the fundamental frequency of a waveform, or heartbeat of the patient, falls between 40 and 120 Hz, and a low pass filter is created by only allowing frequencies between the fundamental frequency and 10 times above it. Small perturbations of points vertically due to error of extraction can add power to high frequency components, so by limiting the region considered and with a good enough extraction method, this error is eliminated.

As a predictor of stenosis, the stenosis index outperforms the resistive index. For 48 patients with stenosis and 21 patients without it, the value of 1.26 for the stenosis index performed at 0.96 sensitivity and 0.38 specificity and at 1.01 it performed at 0.60 sensitivity and 0.90 specificity. Whereas at a value of 0.50 for the resistive index, it only performed at a 0.96 sensitivity and 0.29 specificity [4]. So when thresholded to predict stenosis correctly, high sensitivity, the stenosis index and the resistive index perform similarly; however, the stenosis index has a higher specificity, meaning that less patients without stenosis are misclassified, reducing unnecessary angiography, while still keeping the predictive power of screening for stenosis.

Currently, the stenosis index is calculated within MATLAB through human guidance in the extraction of the waveform. In order to extract the waveform to calculate the stenosis index, the user must refine the region manually, help guide noise elimination, and trace the waveform. This process is tedious and can be time consuming, and, if sonographers have trouble identifying peaks and troughs consistently between themselves, calculating the same index between users can also be troublesome. As such, automation or guided automation of this process is better for a clinical setting, allowing for standardized practices to generate the same stenosis index and perform faster.

Machine Learning and Doppler Waveforms

The vast majority of literature utilizing machine learning to determine the presence of stenosis though Doppler imaging uses the carotid artery. This is due mostly to generally pristine waveforms, typically free from noise and malformations, the carotid artery produces. The carotid arteries location limits the impact of breathing, eliminating gaps. The artery is also superficial reducing the amount of tissue the ultrasound must traverse, reducing reverberation and reflection that leads to noise. Finally, the ease of finding the artery allows for easy resampling if any noise or malformations occur. The methods and algorithms provide useful insight into how to analyze pristine waveforms and can be utilized after refinement of malformations, gaps, and noisy data. There have been many machine learning techniques used to detect stenosis. These include support vector machines [7,17], k-nearest neighbors[7,8], principal component analysis[8], and neural networks[23,24]. These models often utilize the classical point based features in their feature set, such as the resistive index. Other features are often crafted as well using the waveform envelope data as well as spectral data. K-nearest neighbors and support vector machines perform around a .80 sensitivity and specificity [7]. Neural networks boast high sensitivity usually around .80 or .90 sensitivity with 1.00 specificity [23,24]. Machine learned models have also been shown to outperform single features in classification [8,17]. Again, a high sensitivity is incredibly important as Doppler imaging is used as a screening method for stenosis as false negatives are worse than false positives. While neural networks tend to perform with the highest accuracy, they lose model interpretability compared to other model frameworks. Classical features allow clinicians to relate blood flow characteristics to a model and are very interpretable, allowing the clinician to understand how the feature works

and where it could fail. Therefore, a more interpretable model is ideal, keeping as high a sensitivity and specificity as possible.

CHAPTER 3

METHODOLOGY

The goal is to classify whether stenosis exists for a given patients Doppler ultrasound. Three arteries, left hepatic artery, right hepatic artery and proper hepatic artery, are sampled for each patient. The waveforms from these images are to be extracted in an automated fashion. From this, features to predict stenosis are extracted and calculated, where they can be utilized together through a machine learning model.

The data set considered contains 77 patients, where the patients either are from a healthy control group having no stenosis, or are liver transplant patients who underwent angiography and had stenosis determined. 40 patients had no stenosis or where in the control group and 37 patients had varying degrees of mild to severe stenosis.

To get the best information from a feature, especially the peak systolic velocity and resistive index, the blood flow velocity should be ideally measured at or just beyond the arterial anastomosis of the liver transplant to detect the presence of stenosis. However, this junction is often obscured by overlying bowel gas which greatly attenuates and scatters the sound waves needed for ultrasound imaging. As an alternative, downstream blood flow within the right or left hepatic arteries is often measured. In contrast to imaging the carotid artery, which is relatively close to the surface of the body, imaging of arteries within the abdomen is more challenging. Respiration can cause the artery to move in and out of the imaging field, the orientation of flows can be both towards and away from the probe, and flow in adjacent blood vessels can be superimposed on arterial flow measurements. Automatic detection of the waveform is therefore not a trivial task. Even an expert ultrasound technician might occasionally make errors in locating the systolic peak and diastolic trough in blood velocity [14]. Features of the Doppler

sonogram that make extraction of the envelope challenging are aliasing of the waveform, ‘salt and pepper’ noise, faint or sparse signals, patchy waveforms with interior ‘holes,’ inversion of the waveform, incomplete/interrupted waveform, and wrap around start and end points as shown in Figure 6. The prevalence of these images is 45 percent good images with the rest being mildly noisy to extremely noisy with presence of gaps and malformations.

Aliasing and ‘salt and pepper’ noise make it hard to detect the junction between the crest of the wave and background. ‘Holes’ within waveforms make it challenging for an edge detection algorithm to find the correct edge. Inversion of the waveform is caused by blood flow directed away from the ultrasound probe. For a more accurate Fourier transform, a minimum of three contiguous waves is necessary. When only a subset of a sonogram is usable, the waveform detection algorithm needs to detect and extract the corresponding subsequence. Finally, in the

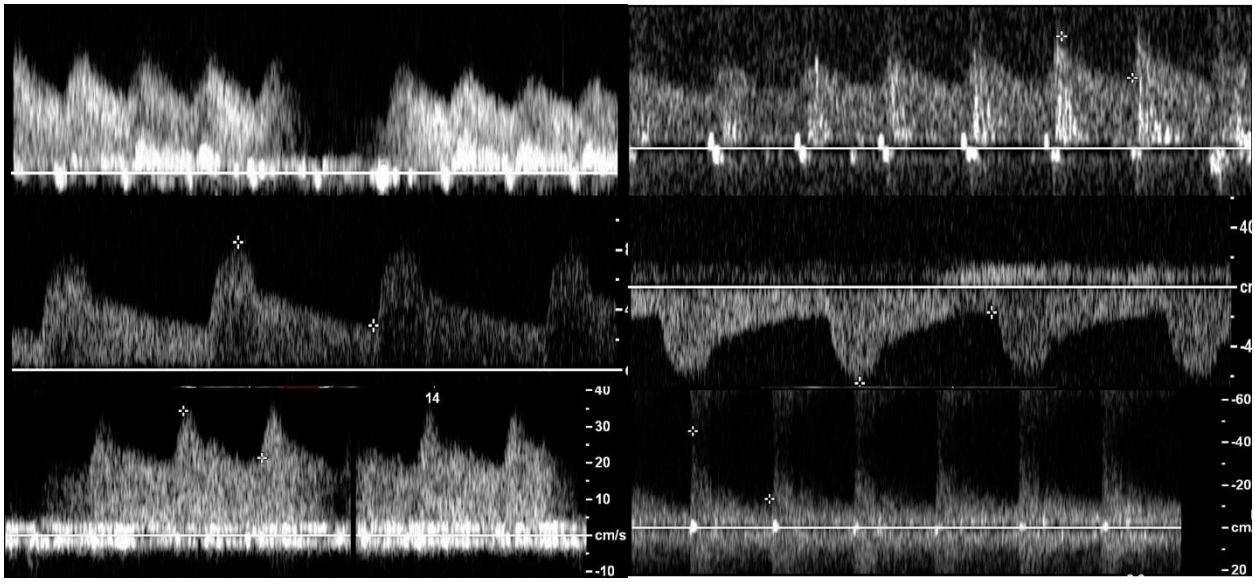

Figure 6. Features of a Doppler ultrasound that make envelope extraction difficult: Gaps and incomplete waveforms (top left), ‘salt and pepper’ noise above the waveform (top right), sparsity of the waveform (middle left), inversion of the waveform (middle right), start and end points being wrapped horizontally shown by the discontinuity in baseline (bottom left), and wrapping of the waveform vertically as shown by waveform below the baseline continuing from above (bottom right).

cases where the sonogram is useless, the detection algorithm should be able to conclude that a repeat sonogram is needed.

To achieve this, an iterative refinement, waveform extraction, and feature extraction scheme was created as shown in Figure 7. As the image quality and waveforms vary, these methods need to be able to work on pristine waveforms as well as waveforms with noise and malformations. This is achieved by iterative refinement so that the region of interest becomes better, regardless of image quality, with each method. This allows extraction to be easier and less prone to being influenced by noise, which leads to better feature extraction.

To automate the prediction of stenosis, the waveform as well as auxiliary features must be extracted from a Doppler image. The most pertinent information is the waveform, where the presence of malformations and gaps can make extracting the entire waveform undesirable; a subset of the waveform can provide wave information for determining more accurate features. Other information, such as the unit scaling for the blood flow velocity through the hepatic artery, the y-axis on the image, and the unit scaling for the time at which the blood flow is measured, the x-axis, also need to be extracted automatically. Additive methods of refinement are utilized to refine the region, with each method reducing noise, improving the waveform, or finding the best region for extraction for all image types. To extract the waveform from the image, two different methods were explored: edge detection and instantaneous blood flow velocity distribution techniques.

Doppler Image Refinement

Refining the Doppler image happens in phases: vertical refinement, horizontal refinement, and waveform improvement. While different ultrasound machines will draw the waveform at different locations near the bottom of the ultrasound image, each machine draws the

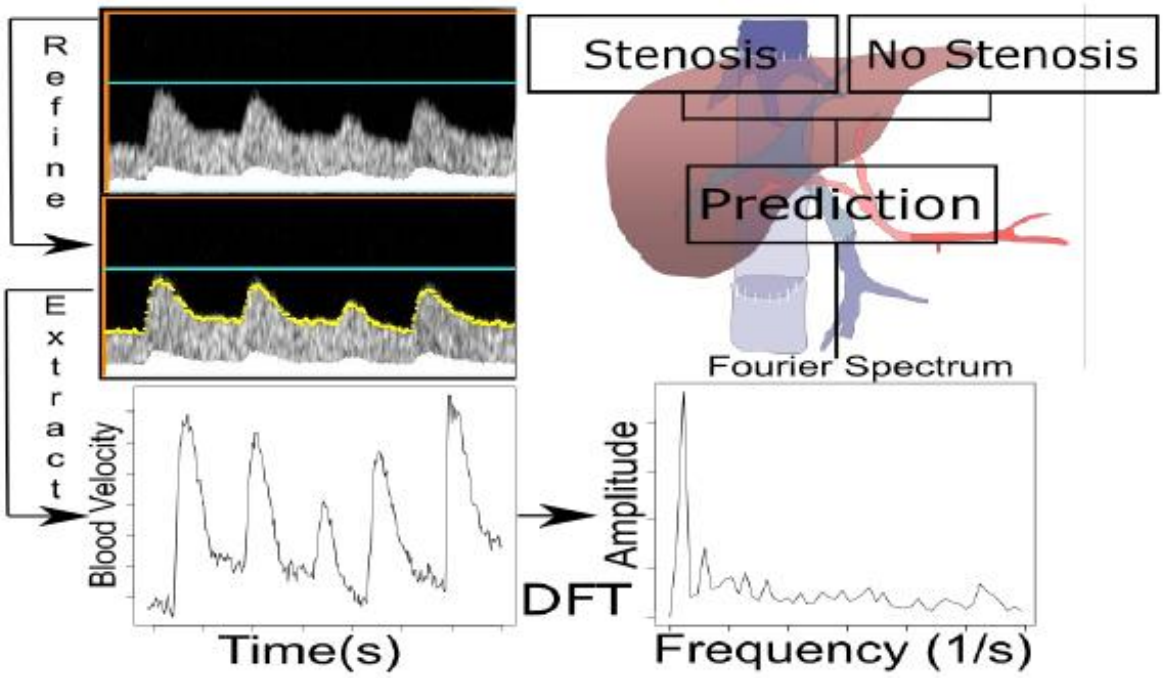

Figure 7. The workflow predicting stenosis for a patient. An image is iteratively refine, improving the region of interest through each refinement. The waveform is then extracted from the region, and then feature extraction is performed. These features can then be utilized to predict stenosis through classical or machine learned models.

waveform region in the same place each time. This location is saved as a property of the machine type and is read from a file by the software. The baseline of the image is found by the software by looking for a long line of a given color, allowing some tolerance of short gaps due to baseline wrapping. Using the machine preset area and the baseline, an outer rectangle can circumscribe the waveform region accurately. Often, the waveform can be measured below the baseline as seen in Figure 6; in these instances, the waveform region (and not the axes) is mirrored over the baseline, so the waveform region of interest will always be upright above the baseline to make extraction easier.

Depending on the blood flow velocity unit scaling and the waveform itself, a waveform may not fill up the entire region on the Doppler image, so refining vertically is employed to reduce the region of interest to just the peaks of the waveform. There can be a large vertical distance between the tops of the peaks of the waveform and the machine's preset area to draw

the waveform. Any measurement of blood flow above the peak region of the waveform corresponds to noise either due to reverberation or waveform data from below the baseline being cut off by the machine's bounding area and being drawn spilling over to the top of the area as seen in Figure 6. To reduce this, the region should be refined to just the top of the peaks. A technologist may have estimated the peaks while taking the ultrasound, which is shown with a cross hair. If available, this is utilized as human detection of peak row is generally accurate. If this is not available, the peaks are estimated by finding the row which has the best horizontal cut to restrict the area to just the peaks, and this is done by an approach on analyzing distributions. As rows of the waveform increase along the peaks, the sum of the brightness across the row decreases, as more of the trough region is taken up by dark background space. By taking a seventy-five percent difference to be the threshold, initially, for declaring a difference, the tops of the peaks of waveform can be estimated, and called the inner rectangle. If the inner rectangle row is determined to be the same as the outer rectangle row, the threshold is reduced; as in this case, the region hasn't been reduced, and no image seen has peaks that hit the exact top of the image, as the scale is adjusted to contain the waveform. The initial threshold is taken to be conservative in cutting of the tops of the peaks as it is better to overestimate the peaks than underestimate them with regard to feature extraction. An example and algorithm can be seen in Figure 8.

Horizontal refinement is also employed. Malformations and black gaps inhabit many waveforms within the data set, so by finding a subsequence of the waveform that doesn't contain

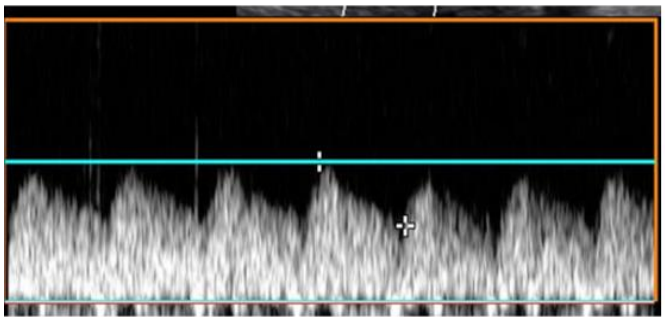

```
DeterminePeakRow(Image, OuterRectangle, Threshold)
{
   if(Image contains crosshairs)
   {
      return row of upper crosshair //human has found peak all ready
   }
   avgVel = calculate average velocity of Image bounded by
       OuterRectangle
   avgVelRow = avgVel row
   for(row = avgVelRow, row <= OuterRectangle row, row++)
   {
      if(row is OuterRectangle row)
      {
         return row //case where inner rectangle isn't found
      }
      thisRowRGB = calculate and sum Image rgb values along row
          bounded Outer Rectangle
      nextRowRGB = calculate and sum Image rgb values along row+1
          bounded Outer Recangle
      if(nextRowRGB/thisRowRGB < Threshold)
      {
         return row //case where inner rectangle is found
      }
   }
}
```

Figure 8. Vertical refinement of a Doppler ultrasound (left) and algorithm for determining the inner rectangle (right). The orange outer rectangle is the top of the preset area of the machine. The lower cyan line shows the top of the inner rectangle which is either found from a cross hair (as shown) or estimated by the marginal distribution (detailed in algorithm).

these is ideal. Also, the stenosis index requires at least three waves of the waveform to be calculated which applies a lower bound to the subsequence extracted. Many images can either have only three waves within their waveform or have only three contiguous waves without malformations or gaps separating them. As such, looking for a region of interest with three contiguous waves within the waveform, free from malformations and gaps, is utilized. This is done by an autocorrelation method finding three contingent waves within the waveform with the highest similarity between them. A copy of the waveform is shifted against itself, where when it is in phase, a minimum error is achieved and when it is out of phase, a maximum error is achieved. An example error from shifting a waveform is shown in Figure 9. The distance between subsequent minimum errors is an estimate of the global frequency of the waveform. The number of points compared is used to normalize this error, the rising error shown in Figure 9 can be explained by the gap in the image. Even though the error is normalized by the number of points, the gap within the waveform is causing error in two different locations within compared

regions, meaning it contributes more error in smaller waveforms from shifting even with point normalization. This rising nature does not affect the estimate of the global frequency though.

The global frequency is then used to form three contiguous panels on the waveform as shown in Figure 10. The similarity of these panels is computed using a modified metric instead of sum of squares errors. The panels are then shifted across the waveform, where the similarity at each shifting is also computed. Finally, the length of each panel is extended or reduced slightly based on the global frequency size and shifting and scoring is performed again normalized by the number of points considered.

The length is modified as waves within a waveform can have slightly higher or lower frequency than the global frequency, due to either error in estimating the global frequency or wave variation, so these oscillations help to find better fits for waves.

The sum of squares errors as a metric rank sparse or missing waves rank better in similarity over more pristine, near perfect waves with mild to low random , so a different metric was created. This poor region finding was due to the sum of squares errors scoring regions of darkness between waveforms highly. A metric to identify similarity between waves better was

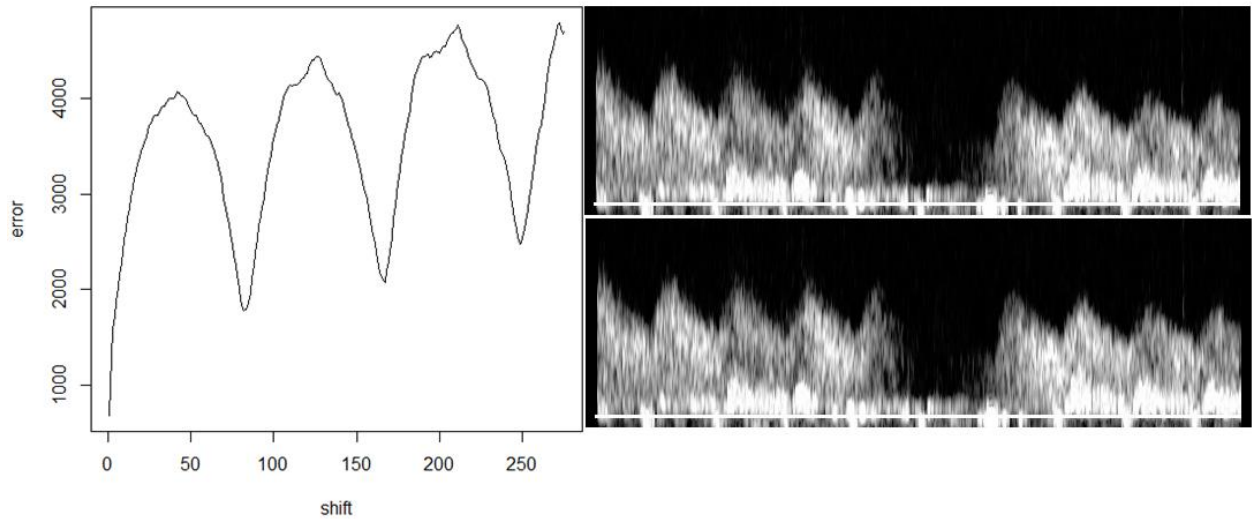

Figure 9. Error obtained from shifting the copied form against the original form.

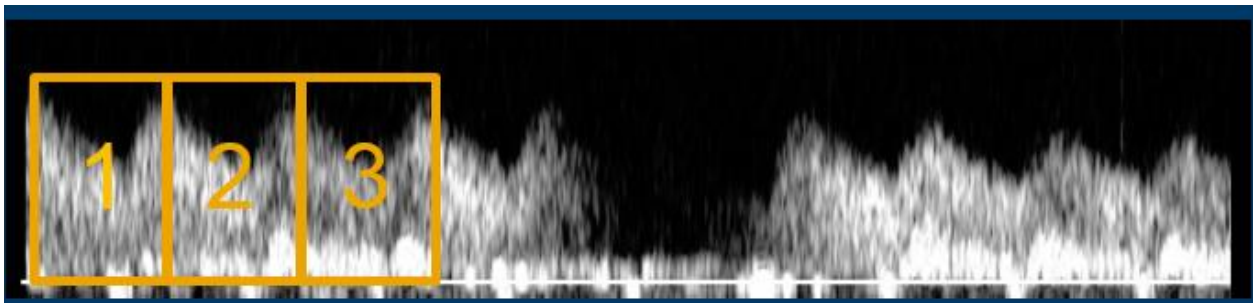

Figure 10. Panels for 3form refinement. The panels are shifted over the waveform and the similarity is calculated. The region of highest similarity is found, and the region of interest is restricted, helping to eliminate gaps and malformations.

created so that presence of similar waves among panels yielded better scores than panels with mostly background or dark regions. This idea is illustrated in Figure 11.

Pseudo counts are also added to refine the inner region of the waveform. Often, a black band is added above and below the baseline by the machine as it helps to distinguish the baseline from the waveform for the technologist. This can affect waveform characteristics, primarily the mean, so a white band a portion of each instantaneous blood flow velocity's mean is added to each column to negate this. Additionally, holes within the waveform can be mitigated by filling them in, reducing their effect on waveform extraction. These pseudo counts also help with

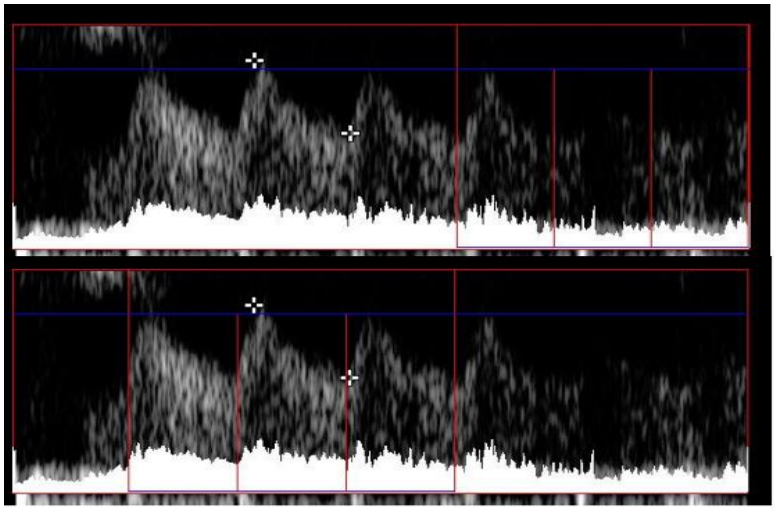

```
modifySSE(SSE error, pixel1, pixel2)
{
   ThresholdBright = lowest average rgb value that corresponds to
      waveform brightness
   ThresholdDark = highest average rgb value that corresponds to
      background darkness
   if(pixel1 avg rgb > ThresholdBright && pixel2 rgb > Threshold
      Bright)
   {
      return error //unmodified as both are waveform pixels
   }
   else if (pixel1 avg rgb < ThresholdDark && pixel2 avg rgb <
      ThresholdDark)
   {
      return error*error //both pixels are from background
   }
   else
   {
      return error*error*error //pixel mismatch
   }
}
```

Figure 11. The use of a sum of squared error as a metric rewards dark regions with no waveform in sparse signals (top left). When the SSE metric is modified to score regions of waveform similarity better than background or mismatch similarity (right), the correct region is found (bottom left).

percentile extraction. By inflating the brightness distribution close to the baseline, the percentiles are shifted further towards the true waveform, away from noise, especially in images where noise brightness is darker than the waveform.

One dimensional horizontal Gaussian blurring is also applied to the waveform, initially with a standard deviation of 3 pixels. Vertical Gaussian blurring was not utilized as it has several disadvantages. The noise above the waveform tends to be "skinny and tall," allowing horizontal blurring to reduce its brightness where vertical blurring could inflate it. Horizontal Gaussian blurring also helps to fill in any gaps across the top of the waveform while vertical blurring would exacerbate it. Also, blurring of either kind can change the waveform envelope itself, smearing the peaks, and changing the underlying frequency distribution, so the value of reducing noise and producing a truer extracted waveform must be weighed against affecting the waveform envelope itself. Asymmetric Gaussian blurring with waveform contextualization was also explored and explained later one.

An alternative refinement scheme to rid waveform regions of malformations and abundant noise was also investigated. Dubbed wave apoptosis, the global frequency found through autocorrelation methods described above is used to divide the waveform into distinct waves. The similarity of each wave is found with all other waves, forming a similarity matrix. The rows of the matrix are summed to find the total error that a wave has with all other waves. The most errored wave, with the highest summed error, and the most average wave, with the lowest number, are selected. These are used to cluster the remaining waves, so the waves with lower error, and more similar to, the average wave are clustered to it and the same for the error wave. The errored waves are removed and an average wave is generated from the remaining waves. This wave is replicated across the entire form, and this is illustrated in Figure 12.

Different frequencies between waves within a waveform can affect this method as if the peaks are slightly shifted away from each other while calculating the average wave, the average wave becomes malformed and rectangular. This outcome can also occur if the global frequency isn't correctly estimated. Peak shifting and single wave rescaling are being evaluated to fix this. Finally, this method also loses the information of the difference of waves as it replicates the same average wave. So currently, only horizontal refinement through 3form refinement is employed for reducing malformations and gaps.

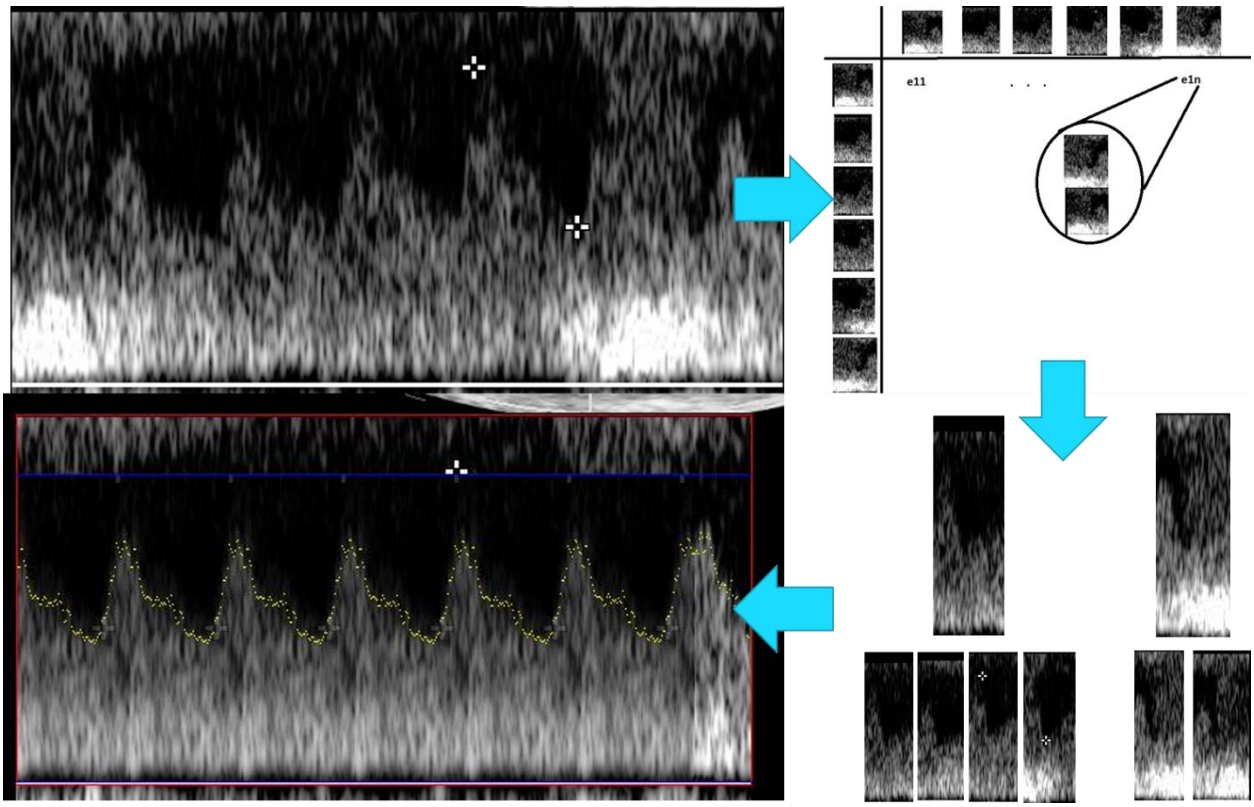

Figure 12. Illustration of wave apoptosis. The wave is broken up into waves based on the estimated frequency. The similarity between waves is then found (top right). The most similar wave to all others and the most dissimilar wave to all others are used to cluster the remaining waves (bottom right). A median wave is generated from the waves closer to the most similar wave and replicated across the waveform (bottom left).

Waveform and Feature Extraction

The waveform and features derived from it must be extracted. Edge detection and percentile based methods were explored for waveform extraction. The percentile methods were further extended to account for distance changes with percentile changes to give rise to the "Go Fish" method. Finally, object character recognition was employed to extract scaling information from the Doppler image outside the waveform to be able to rescale the values for features that aren't ratios, such as the peak systolic velocity and time features.

Edge detection methods were utilized for waveform extraction. Initially, a Prewitt filter [18] was modified to look primarily for vertical edges under the condition that the waveform is on the bottom and the black background is on the top. This gave rise to the hourglass filter as

seen in Figure 13. For a given column, the edge score of each pixel was calculated, where the point with the highest edge score was chosen. This method was further enhanced by taking into account what percentile the pixel was at within the marginal distribution of the instantaneous blood flow velocity. That is, if a pixel was high, the edge probably was due to noise, especially in the trough region, and if it was too low, it was probably an inner edge.

Percentile methods were also explored. To estimate the waveform, the column of the instantaneous blood flow velocity's distribution was examined. The 95th percentile of the column was taken as the estimate for where the waveform resides in that column. The presence of pseudo counts helped to reduce the pulling power noise had at shifting the estimation away from the waveform, especially in the trough region. This can be seen in Figure 14 as the pseudo counts pull the 95th percentile back. An example of this extraction can be seen in Figure 15.

The percentile method was further extended to account for distance changes as percentile changes. This was motivated by several aspects of 95th percentile extraction as shown in Figure 16. On a perfect waveform, this method would compress the wave as it biasedly lowers the

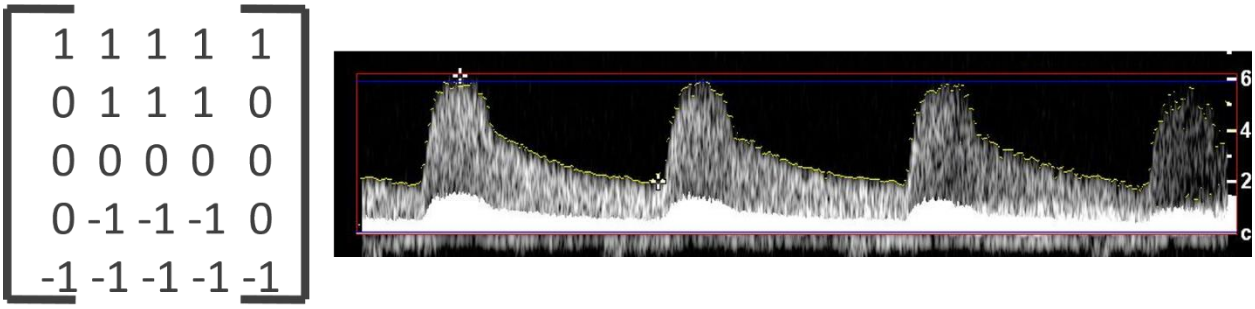

Figure 13. The hourglass metric (left) is used to generate an edge score for each pixel based on their brightness. The best edge is chosen leading to the yellow dots as extraction on the right.

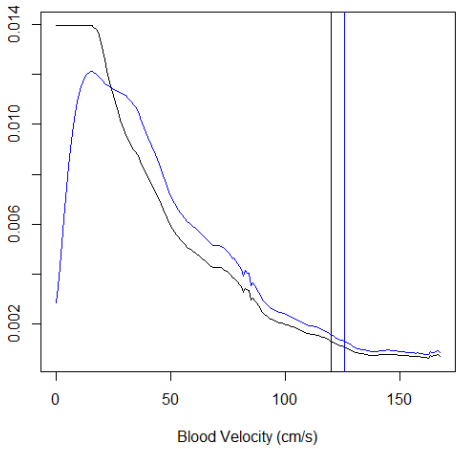

Figure 14. 95th percentile estimates for instantaneous blood flow velocity no pseudo counts (black) and with pseudo counts (blue). The presence of pseudo counts shifts the estimation closer to the waveform,, reducing the impact that noise has on extraction.

peaks of the waveform more pixels than it would the troughs of the waveform. Since this isn't a mere vertical shifting, the underlying frequency distribution can differ, affecting the stenosis index.

The percentile method idea was further enhanced to compare the pixel distances with oscillations to the percentile chosen for extraction. That is, if a change in a percentile, say from 95^{th} to 94^{th}, changed the distance greater than a threshold value, then the point has moved from a noise region closer to the waveform. If the point has moved less than that threshold value, it probably has moved further down within the waveform or further down a bright pocket of noise.

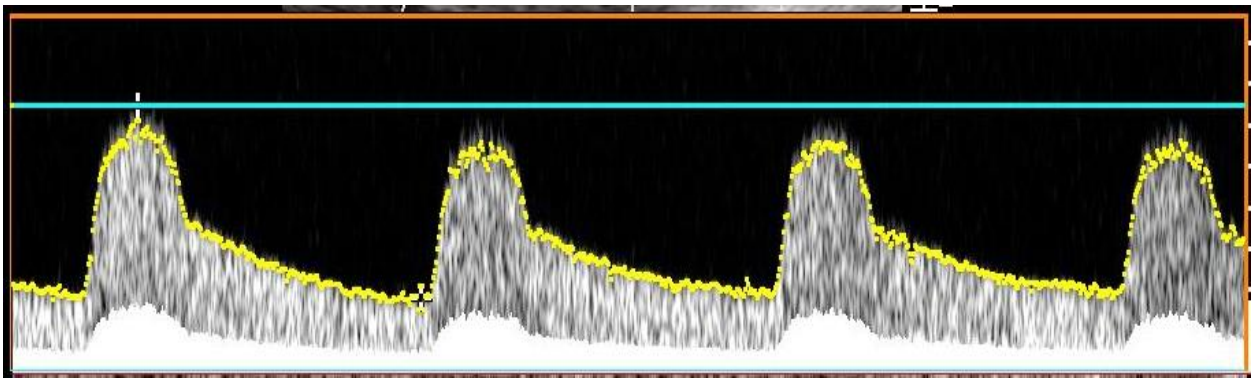

Figure 15. An example of 95^{th} percentile extraction.

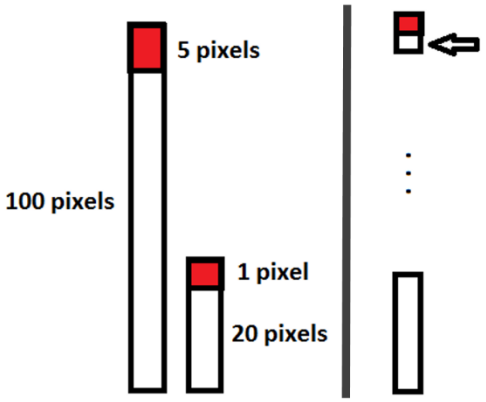

Figure 16. The 95th percentile extraction method underestimates peaks more harshly than troughs (left). The extraction technique also doesn't account for changes of distance with changes of percentile, as large drops from high percentiles to lower percentiles tend to correspond to dropping from noise to the waveform (right).

Similarly, if a point has moved from a smaller percentile, say from 80th to 81st, and it moves a distance greater than a threshold value, it is probably moving from the waveform to noise or across a hole inside the waveform. If it is less than the threshold value, it is probably moving up the waveform towards the top of the waveform and should continue. For this extraction method, a random upper and lower percentile are chosen and moved lower and upper respectively until the distance traveled is not great in the upper to lower case and is great in the lower to upper case (choosing the lower bound as the point). If the values agree, the point is extracted. Points where the extracted values are not within the threshold use neighboring points where the two estimations do agree to determine which estimation to use, choosing which estimation gives a value closer to its. This is termed "Go Fish" extraction as the extraction "fishes" for a point from above and within the waveform. The algorithm can be seen in Figure 17.

The waveform is then contextualized, that is the regions that are peaks and troughs are labeled as such. This allows re-refinement of Gaussian blurring as specific zones can be targeted for different blurring intensities. Peaks are located first by finding a region of points that contain

```
GoFish(Image, InnerRectangle)
{
   FinalWaveform = Initate a blank array of Points of size width of
      InnerRectangle
   UpperWaveform = FishFromAbove(Image, InnerRectangle, ThresholdU,
      StartingPercentileU, DeltaU)
   LowerWaveform = FishFromBelow(Image, InnerRectangle, ThresholdL,
      StartingPercentileL, DeltaL)
   //merge
   for pointU in UpperWaveform and pointL LowerWaveform
   {
      if(pointU.x == pointL.x && |pointU.y - pointL.y| < Threshold)
      {
         set the x point of FinalWaveform to (x, pointL.y)
      }
   }
   while(not all points have been set in FinalWaveform)
   {
      for(points not set in FinalWaveform)
      {
         if(neighbor point has been set)
         {
            diffU = absolute difference in UpperWaveform at current
               point and FinalWaveform at neighbor point
            diffL = absolute difference in LowerWaveform at current
               point and FinalWaveform at neighbor point
            set FinalWaveform at that point to Lower or Upper based
               on the minimum of diffL and diffU
         }
      }
   }
   return FinalWaveform
}
```

```
FishFromBelow(Image, InnerRectangle, Threshold, Starting
   Percentile, Delta)
{
   FishFromBelow = Initate a blank array of Points of size width of
      InnerRectangle
   For each column bounded by InnerRectangle in the Image
   {
      Percentile = Starting Percentile
      yBefore = find pixel location corresponding to Percentile
      Percentile = Percentile + Delta
      yAfter = find pixel location correspponding to Percentile
      while(yAfter- yBefore < Threshold)
      {
         FishFromAbove(col) = new Point(col, yAfter)
         yBefore = find pixel location corresponding to Percentile
         Percentile = Percentile + Delta
         yAfter = find pixel location correspponding to Percentile
      }
   }
   return FishFromBelow
```

```
FishFromAbove(Image, InnerRectangle, Threshold, Starting
   Percentile, Delta)
{
   FishFromAbove = Initate a blank array of Points of size width of
      InnerRectangle
   For each column bounded by InnerRectangle in the Image
   {
      Percentile = Starting Percentile
      yBefore = find pixel location corresponding to Percentile
      Percentile = Percentile - Delta
      yAfter = find pixel location correspponding to Percentile
      while(yBefore - yAfter > Threshold)
      {
         FishFromAbove(col) = new Point(col, yAfter)
         yBefore = find pixel location corresponding to Percentile
         Percentile = Percentile - Delta
         yAfter = find pixel location correspponding to Percentile
      }
   }
   return FishFromAbove
}
```

Figure 17. GoFish algorithm. "Fishing" for the correct location per column is done from above and within the waveform. The two waveforms are merged where they agree and the remaining points utilize neighboring points who have been agreed upon to decide which waveform to choose.

the highest average values without much spread between points. Troughs are then found by looking between peaks for the lowest average value without much spread between points. Finally, these zones are expanded by comparing neighbors of known peak or trough points and seeing which they should belong to. This process is shown in Figure 18.

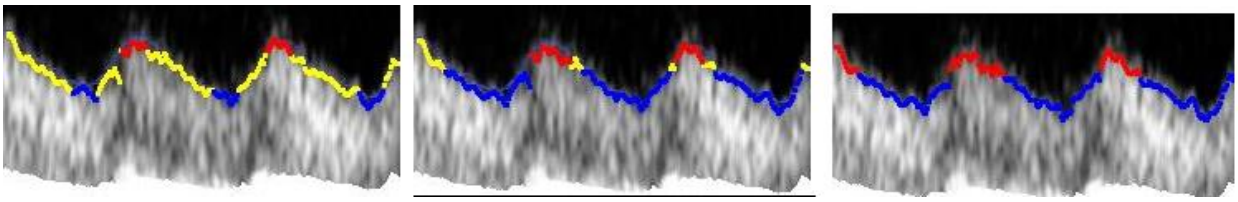

Figure 18. The peaks (red) and troughs (blue) of the wave (yellow) are first found (left). The peaks and troughs are then extended to their neighbors (middle). Finally the remaining points are clustered to classify the remaining points (right). This contextualizes the waveform.

Features are then extracted from the waveform. For the stenosis index to be calculated, a discrete Fourier transform is performed on the waveform envelope. The ratio of the high frequency components to low frequency components is calculated as shown in [9]. For the other features, the ratio indices (such as the resistive index and pulsatility index) can be calculated from the raw values as the pixel scaling cancels. The other features, time features and the peak systolic velocity, must be rescaled by the scaling factors on the image. This scaling must be read in through object character recognition (OCR), as the software assumes access to only the Doppler ultrasound. Tesseract, through the wrapper package Tess4J, was utilized for OCR. The image location is automatically cropped to only include the y-axis and x-axis of the image as to remove as much of the other visual data as possible. The OCR system can misread dashes and numbers, thus, a regular expression was created to weed out any non-numerical characters generated. The scale was then self-validated by iterating through the entire scale to validate the y-scaling chosen. That is, the difference between each successive pair of the y-axis generated from OCR is found and the most agreed on number is the scaling. For the x-axis, only the number near the bottom of the screen needs to be read. Finally, the distances between tick marks in both the x and y directions are used to produce the scaling factor. The features are then converted into appropriate unites and saved to a text file for further analysis or use.

Asymmetric Gaussian Blurring

One dimensional Gaussian blurring can be used to refine the image to negate random noise by averaging over several pixels as well as smoothen the waveform to a degree. To compute the blur of a given pixel, an interval of a set radius about it has Gaussian proportions assigned to each point and then averaged using those proportions. Mathematically, the weighting mechanism for a pixel is: $W(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{x^2}{2\sigma^2}}$ for pixel x and standard deviation,σ. This technique, alone, has some undesirable features. Troughs and peaks of the waveform shouldn't be blurred with the same standard deviation and shouldn't be considered when with blurring the other type. That is, when calculating the blur of a trough near the end of a wave, the rising peak to its right shouldn't blur into it, affecting the true trough. Also, regions of high change and discontinuity should be blurred differently than low change. To allow for this, a hyperparameter should be utilized to account for the slope and continuity as well as the wave contextualization of peak and trough. Currently, asymmetric Gaussian blurring only utilizes wave contextualization. Two different standard deviations are chosen for the baseline blurring effect for peak and trough separately, these values are further augmented based on the slope and concavity at the point. This is illustrated in Figure 19.

Machine Learning

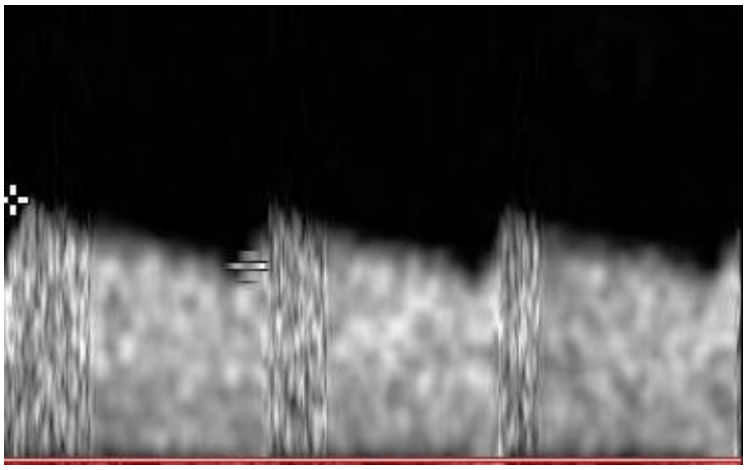

Figure 15. Asymmetric Gaussian blurring where the peaks and troughs are blurred differently.

Machine learning models were trained to classify stenosis. From the waveform extraction method detailed above, the resistive index, stenosis index, acceleration time, global frequency, and right and left half-lives were extracted as features. The right and left half-lives are the time it takes the blood flow velocity to reduce to half the distance from the peak systolic velocity to the diastolic velocity on the right and left side of the peak systolic velocity. The ratio of these half-lives was also computed and used in model building. Naïve Bayes, decision trees, and random forests were investigated as models. Naïve Bayes estimates the probability of stenosis under the assumption of conditional independence of features and was utilized as a baseline to determine how much predictive power could be gained from machine learning models. Decision trees greedily partition the training space to either minimize misclassification, with the Gini impurity metric, or gain the most information, based on entropy. Decision trees benefit from highly interpretability but are not guaranteed to generate an optimal model. Furthermore, they can perform poorly if the true relationship between features and stenosis is complex and can overfit to the training set, which pruning can help negate. Random forests are a modified ensemble of decision trees, each with its own data set generated by random subset of the feature set. The use of sampled features, and no pruning, constructs a tree that,taken

individually is a high variance model. Random forests tend to overfit less to the training set as they are an average of these high variance unpruned trees; however, as more trees are utilized, the interpretability of the model decreases. Spark and its MLlib library was used to calculate these models, except decision trees, as it utilizes the Java Virtual Machine, making future integration of it and the software easier if needed. Being built for parallelization, Spark performs quickly when cross validating many tree random forests. Also, it is highly scalable, allowing future work to utilize as many patients and images as possible for learning, quickly. Decision trees were created in R using the Rpart package as the Sparks MLlib library doesn't include pruning. Cross validation was utilized with all models at 10-fold cross validation to decide the best model before moving onto a final test set. Currently, the number of trees and depth of trees is fixed and set to 1000 trees and a maximum depth of 30 for Random forests.

CHAPTER 4

RESULTS

Percentile based extraction outperformed edge detection extraction. The result of wave envelope detection with an hourglass horizontal edge filter is shown in Fig 20. The speckled nature of some Doppler waveforms causes the edge filter to frequently underestimate the velocities, finding edges inside the waveform. In contrast, the 95th percentile of the instantaneous velocity distribution performs quite well on these images as seen in Figure 21. When the sonogram is sparse or the noise above the form vast and bright, the diastolic minimum velocity can be overestimated as noise above the waveform might correspond to the 95th percentile. Here, "Go Fish" methods outperform the basic 95th percentile method as seen in Figure 22. This is due

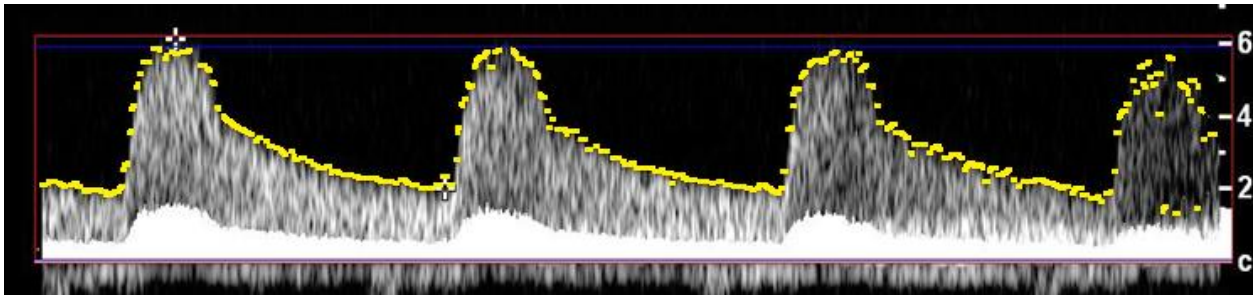

Figure 20. Edge detection extraction. The algorithm works well on nice, dense areas of the waveform but finds inner edges of the waveform near end.

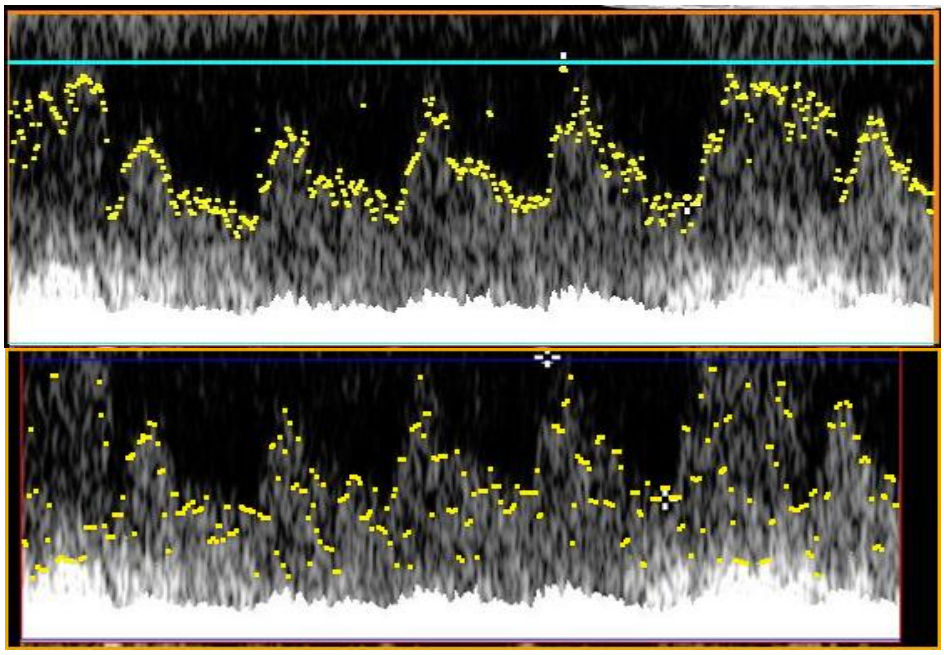

Figure 21. Comparison of edge detection (bottom) and percentile methods (top). Edge detection tends to get detect edges within the waveform even when weighted by distance while percentile tends to do better across the entire form.

to regarding the distances changed due to oscillation of a percentile, essentially trying to find the best percentile for each column instead of a flat value for the entire waveform.

To formalize how well the percentile algorithms performed, human generated waveforms were drawn on top of the images for 50 randomly selected testing images. Then, the difference between the human determined waveform and the machine generated waveform were found. This was used to calculate the absolute mean percent difference for the images, normalized by the number of points per image, for different refinement schemes and the "Go Fish" 3form extraction shown in Figure 23. These extractions were performed on images that had undergone horizontal Gaussian blurring of standard deviation 3.The horizontal 3form refinement finds better regions of the waveform, without gaps or malformations and lower noise. The addition of pseudo counts reduces the error and variance of the error for this method as well. The mean square error for the 3form refinement scheme in Figure 23 also shows that for the majority of the

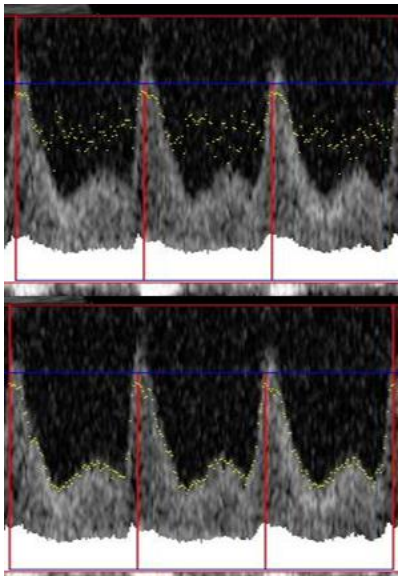

Figure 22. Comparison of 95th percentile (top) and "Go Fish" extraction (bottom). As the distance between noise and waveform is great, fishing from above yields a good estimate; similarly, as the waveform is mostly dense, fishing from below also works well.

images, the extraction performs close to human generated waveform, where they are off by 7-10 pixels for the most part as the waveform height ranges between 150 to 250 pixels. "Go Fish" estimation had an average percent error of 6.70 percent compared to the 7.44 percent difference of just 95th percentile. However, the means are not to be significantly different using the Wilcoxon paired signed ranks test ($p = .1645$). Asymmetric Gaussian blurring reduced the error of Go Fish estimation to 6.50 and 95th percentile estimation to 7.3 and almost a statistically significant difference using a Wilcoxon paired signed ranks test ($p = .06$) between the estimation techniques. While "Go Fish" tended to outperform on mildly noised images, it performed similarly on no noised images and worse on heavily noised images, causing the overall non-significant difference. The images with larger error are the highly noised images where current extraction and refinement methods fail to delineate the waveform from noise.

On some noisy images, this error could be reduced by finding a better inner rectangle; however, on the majority of images where percentile based extraction does poorly, it is due to the

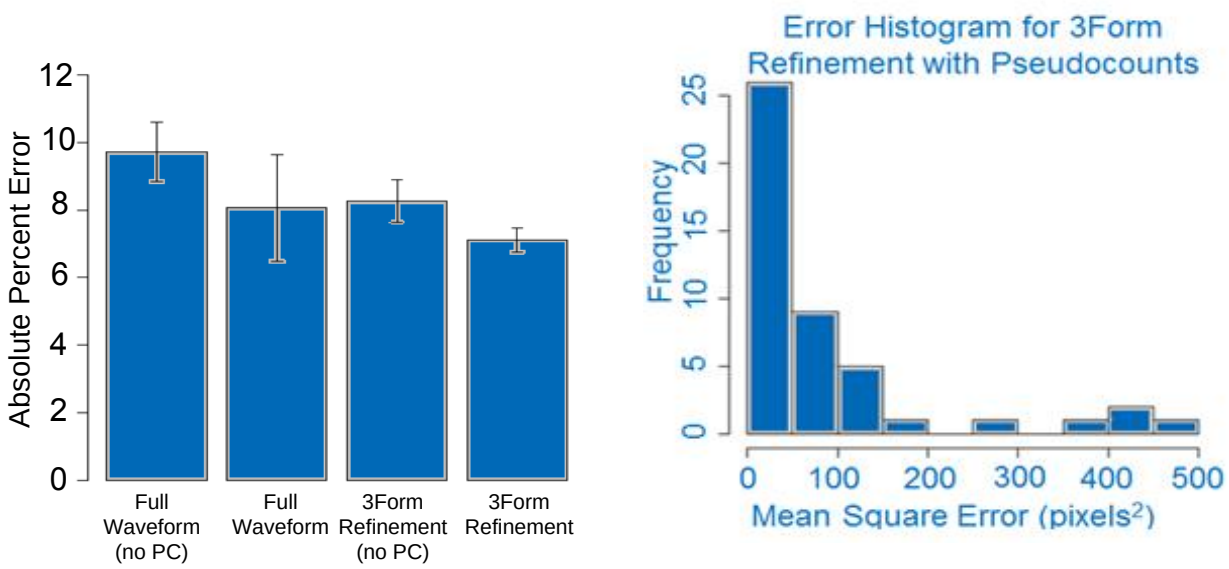

Figure 23. Absolute percent error, normalized by number of points, for refinement schemes displayed (left). Mean square error for 3form refinement (right). These are computed using 95th percentile extraction.

high noise above and within the waveform, washing it out almost entirely. For these images, a different extraction scheme would have to be investigated.

The data set for feature extraction and stenosis prediction contains 77 patients, where the patients either are from a healthy control group having no stenosis, or are liver transplant patients who underwent angiography and had stenosis evaluated. 40 patients had no stenosis, either undergoing angiography and being evaluated as having no stenosis or were part of the healthy control group. 37 patients had varying degrees of mild to severe stenosis. Seventy percent of the data was used to form a training set, which models used 10 fold cross validation to determine the best model, where this model was then tested again on the remaining thirty percent of the original data.

For feature extraction, the automated stenosis index was compared with the Matlab and human derived stenosis index. The absolute percent average error is a 33 percent difference, but dropping to 22 percent difference when images that contain stenosis index further than 3

standard deviations away from the mean index are removed. While a pixel for an automated wave is only 5 to 7 pixels away from a human determined wave on average, if the estimation varies alternatingly where the estimation should be smooth, the high frequency power of the frequency envelope may be inflated, affecting the calculation of the stenosis index. This occurs often on images where the top of the waveform is sparse, causing extraction techniques to over fit to the machines output and can be seen in Figure 24. Demanding continuity and smoothening the waveform itself could lessen this and is being explored through contextual cubic splining. Both the automated stenosis index and the Matlab human guided stenosis index had a poor time differentiating between patients with stenosis and no stenosis alone. The two patient groups, stenosis and no stenosis, were found to not be statistically significantly different for either metric ($p = .745$ automated, $p = .58$ human guided Matlab) using a Wilcoxon rank-sum test.

For stenosis prediction, single feature models and machine learning models utilized 10 fold cross validation on seventy percent of the total data. Single feature thresholds were found using linear discriminant analysis, where the thresholds found for the automated resistive index and Matlab guided stenosis index agreed with previous studies. Sensitivity, specificity, and ROC curves of how the models performed on the remaining thirty percent of the data were used to

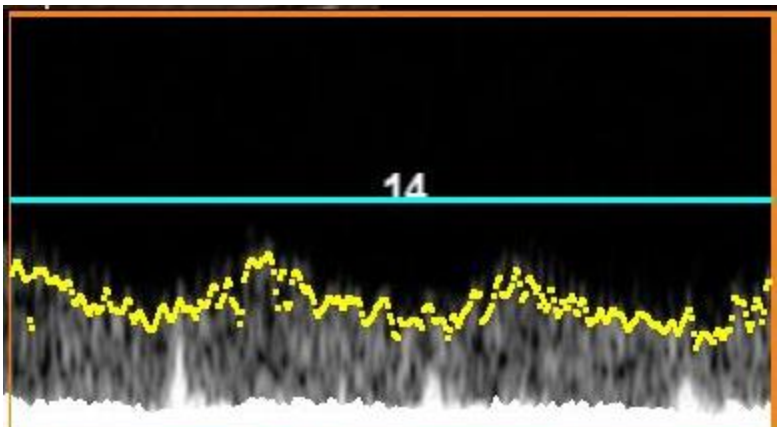

Figure 24. Variation of extracted waveform causes the high frequency components of the waveform to be inflated, affecting the calculation of the stenosis index.

describe their classification abilities. These results and thresholds for classification were also found for single feature models as shown in Figure 25. Neither the automated stenosis index nor the Matlab human guided stenosis index performed well as classifiers and gave poor ROC curves as seen in Figure 26. The left half-life seemed to a better predictor of stenosis than the acceleration time as seen in Figure 25 and the ROC curve in Figure 26. This is probably due to the left half-life being easier to estimate automatically as it has a more precise definition that is more resilient to noise and variation. Thus, the left half-life is probably capturing the temporal information the acceleration time aims to capture when calculated by technologist pinpointing. Finally, the ratio of the left half-life to the right half-life is a better indicator of stenosis than the right half-life. Thus, for model building purposes, the ratio was preferred to the right half-life. The resistive index and left half-life had the best sensitivity and specificity and ROC curves for all single predictor features as seen in Figure 27. This can be explained in part in that the information they measure has been shown in the literature review to be informative when making a classification of stenosis and also that these features are easy to extract from the waveform envelope and aren't affected by noise greatly.

Model	Threshold	Sensitivity	Specificity	AUC
RI	0.5	0.8	0.67	0.77
SI-automated	1.41	0.88	0.08	0.49
SI-matlab	1.19	0.5	0.45	0.45
Frequency	1.16 Hz	0.3	0.92	0.595
Acceleration time	0.172 sec	0.5	0.58	0.54
Left Half Life	0.208 sec	0.9	0.75	0.813
Right Half Life	0.27 sec	0.8	0.083	0.57
Half Life Ratio	4.72	0.9	0.33	0.756
Naïve Bayes		0.54	0.67	0.725
Random Forest		0.8	0.67	0.833

Figure 25. Thresholds and prediction of single feature and machine learned models at determining stenosis.

For machine learning models, Naïve Bayes performed at 0.66 sensitivity and 0.60 specificity. After pruning, the decision tree model became a single node classifier using the resistive index, whose threshold agreed with the LDA threshold. Finally, the random forest model performed at 0.80 sensitivity and 0.66 specificity. The ROC curves for the machine learning models can be seen in Figure 27. Not only does the random forest model outperform Naïve Bayes in terms of sensitivity and specificity, its ROC curve reaches high sensitivity much faster, making it a better screening method that could reduce the amount of false positives and unnecessary angiography. The ROC performance of machine learning models and top single feature classifiers, left half-life and resistive index, is shown in Figure 28. The left half-life reaches high sensitivity (90%) very fast (at 83% sensitivity), but requires a massive drop in

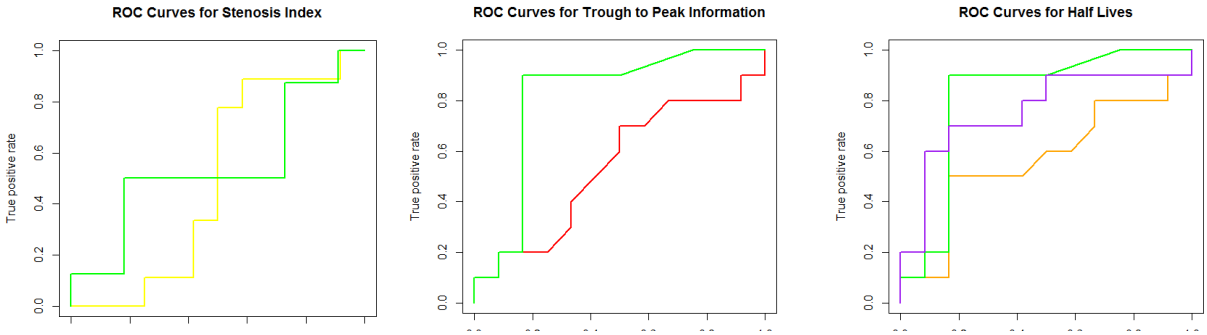

Figure 26. ROC curves for stenosis index (left), automated (yellow) and human guided (green), for trough to peak estimators (middle), left half-life (green) and acceleration time (red), and for half-lives(right), left half-life (green), right half-life (orange) and ratio of right to left (blue).

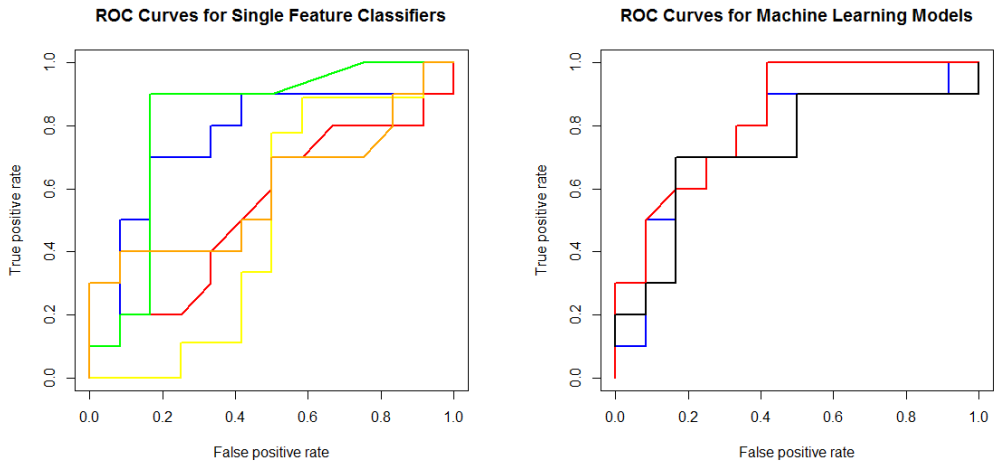

Figure 27. ROC curve for single feature classifiers (left), left half-life (green), resistive index (blue), stenosis index (yellow), acceleration time (red), and frequency (orange). ROC curve for machine learning models (right), Naïve Bayes (black), decision tree (blue), and random forest (red).

specificity to improve it. The resistive index and random forest performed the same in terms of sensitivity and specificity, but the random forest model reaches high sensitivity (greater than 90%) much faster than both it and the left half-life in the ROC curve. Thus, it could be thresholded to reduce false negatives, increasing sensitivity, while keeping a higher specificity than the other models. The random forest model uses the square root of number of features, 3, to sample the feature space to build trees; this could lead to trees being formed with uninformative or correlated features which could muddy classification, such as trees made with the stenosis index, frequency or acceleration time or trees with left half-life and the ratio of half-lives. Better feature bagging, feature removal, or better feature extraction could improve performance. Also, there w some mixture of features that is allowing the forest to learn how to differentiate stenosis where the resistive index and left half-life were unable to alone.

It is important to underscore that the random forest model achieves greater than 90% sensitivity at the lowest cost of specificity and false positives. As a screening method, false negatives need to be minimized to ensure failing liver transplant are detected. With the current

test set, it outperforms the left-half life by over 30% specificity and the resistive index by almost 50%. So while there isn't a statistically significant difference between the AUC for the random forest model and left half-life (p =0.30) or the resistive index (p =0 .23) models for the entire curve, the random forest combination of high sensitivity with acceptable specificity make it a model scheme worth exploring.

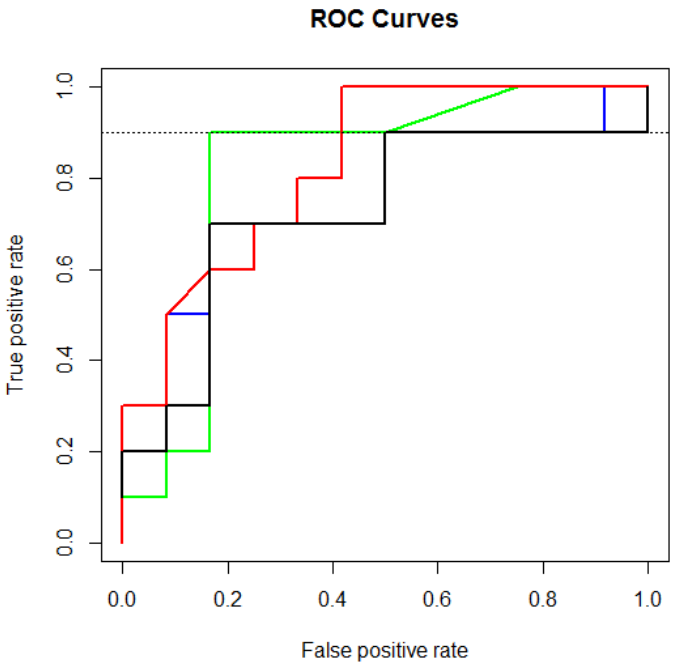

Figure 28. ROC curves for top performing models. Left half-life (green), Naïve Bayes (black), resistive index (blue), and random forest (red). The dotted line corresponds to 90% sensitivity threshold.

CHAPTER 5

FUTURE WORK

The future of this research can have many directions. When a user guides the machine's refinement process, the software logs what changes were made to what it had decided and what estimations were left unchanged. This information can guide distribution based extraction by identifying if a similar distribution or region has been identified accurately before. Wave averaging, as seen through wave apoptosis, has promise in reducing noise above the waveform.

The trough region tends to be the nosiest region due to the large space between it and the top of the inner rectangle. Images that contain light to moderate noise or deep gaps in the waveform due to sparsity can have major variation within small frames of the trough. Using the variance of the waveform as a scoring mechanism, fitting cubic splines to the random trough points within the same trough has shown some promise at being able to fix regions of noise or Doppler ultrasound machine error where there are black gaps in the top of the waveform. Currently, the choice of knots is done stochastically, but needs better guidance to have consistent results.

Finally, machine learning algorithms and feature selection can be further explored. Currently, only one wave within the waveform is used to calculate all features except for stenosis index which uses the entire refined waveform. Incorporating the variability of these single wave features within a patient's waveform could be useful. Also, better extraction could make features such as the stenosis index more informative in the model.

REFERENCES

1. Abu-Wasel, B., Renfrew, P.D., & Molinari, M. (2013). Liver transplantation and endoscopic management of bile duct complications, endoscopy of GI tract. *InTech.* DOI: 10.5772/52630. Retrieved from: https://www.intechopen.com/books/endoscopy-of-gi-tract/liver-transplantation-and-endoscopic-management-of-bile-duct-complications.

2.. Alexandrov, A.V., Brodie, D.S., McLean, A., Hamilton, P., Murphy, J., & Burns, P.N. (1997). Correlation of peak systolic velocity and angiographic measurement of carotid stenosis revisited. *Stroke, 28*(2): 339-342.

3. Baxter, G.M., Ireland, H., Moss, J.G., Harden, P.N., Junor, B.J., Rodger, R.S., & Briggs, J.D. (1995). Colour doppler ultrasound in renal transplant artery stenosis: Which doppler index? *Clinical Radiology, 50*(9): 618-622.

4. Chan, S., McNeely, M.F., Le, T.X., Hippie, D.S., Dighe, M., & Dubinsky, T.J. (2013). The Sonographic Stenosis Index: Computer simulation of a novel method for detecting and quantifying arterial narrowing. *Ultrasound Quarterly 29*(3): 155-60.

5. Dodd, G.D., Memel, D.S., Z., A.B., Baron, R.L., & Santaguida, L.A. (1994) Hepatic artery stenosis and thrombosis in transplant recipients. Doppler diagnosis with resistive index and systolic acceleration time. *Radiology 192*(3): 657-61.

6. Gottlieb, R.H., Lieberman, J.L., Pabico, R.C., & Waldman, D.L. (1995). Diagnosis of renal artery stenosis in transplanted kidneys: Value of doppler waveform analysis of the intrarenal artery. *American Journal of Roentgenolgy 165*(6): 1441-1446.

7. Krishnamoorthy, P, Ravindra B, & Ravi, V. (2015). Temporal and spectral analysis of internal carotid artery doppler signal for normal and abnormal flow detection. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. 6122-5.

8. Latifoğlu, F., Polat, K., Kara, S. & Güneş, S. (2008). Medical diagnosis of atherosclerosis from carotid artery Doppler signals using principal component analysis (PCA), k-NN based weighting pre-processing and artificial immune recognition systems (AIRS). *Journal of Biomedical Informatics 41*(10):15-23.

9. Le, T., Hippie, S. McNeely, M., Dighe, M. Dubinsky, T, & Chan, S. (2017). The Sonographic Stenosis Index: A new specific quantitative measure of transplant hepatic arterial stenosis. *Journal of Ultrasound in Medicine 36* (4): 809-819.

10. Liver Foundation. *Liver transplants*. Retrieved January 1, 2017 from https://liverfoundation.org/downloads/alf_download_134.pdf.

11. National Center for Health Statistics. *Chronic liver disease and cirrhosis*. Retrieved January 1, 2017 from https://cdc.gov/nchs/fastats/liver-disease.htm.

12. Park, Y., Kim, K.W., Lee, S.J., Lee, J., Jung, D., Song, G., Ha, T., Moon, D., Kim, K., Ahn, C., Hwang, S., & Lee, S. (2011) Hepatic arterial stenosis assessed with doppler US after liver transplant: Frequency false-positive diagnoses with tardus parvus waveform and value of adding optimal peak systolic velocity cutoff. *Radiology 260*(3): 884-891.

13. Patel U., Khaw, K.K., & Hughes, N.C. (2003). Doppler ultrasound for detection of renal transplant artery stenosis-threshold peak systolic velocity needs to be higher in a low-risk or surveillance population. *Clinical Radiology 58*(10): 772-777.

14. Paulson, E.K., Kliewer, M.A., Frederick, M.G., Keogan, M.T., Delong, D.M., & Nelson, R.C. (1996). Hepatic artery: Variability in measurement of resistive index and systolic acceleration in healthy volunteers. *Radiology 200*(3): 725-729.

15. Petersen, L.J., Petersen, J.R., Talleruphuus, U., Ladefoged, S.D., Mehlsen, J., & Jensen, H.A. (1997). The pulsatility index and the resistive index in renal arteries. Associations with long term progression in chronic renal failure. *Nephrology Dialysis Transplant 12*(7): 1376-1308.

16. Platt, J.F., Yutzy, G.G., Bude, R.O., Ellis, J.H., & Rubin, J.M. (1997) Use of Doppler sonography for revealing hepatic artery stenosis in liver transplant recipients. *American Journal Roentgenology 168*(2): 473-476.

17. Polat, K., Sadik, K., Latifoğlu, F, & Güneş, S. (2007). Pattern detection of atherosclerosis from carotid artery doppler signals using fuzzy weighted pre-processing and Least Square Support Vector Machine (LSSVM). *Annals of Biomedical Engineering* 35(5):724-732.

18. Prewitt, J.M.S. (1970). *Picture processing and pscyhopictorics*. Cambridge, MA: Academic Press.

19. Rifkin, M.D., Needleman, L., Pasto, M.E.,Kurtz, A.B., Foy, P.M.,McGlynn, E., Canino, C.,Baltarowich, O.H., Pennell, R.G., & Goldberg, B.B. (1987). Evaluation of renal transplant rejection by duplex doppler examination: Value of the resistive index. *AJR American Journal of Roentgenology 148*(4): 759-762.

20. Saarinen O, Salmela, K, & Edgren, J. (1994). Doppler ultrasound in the diagnosis of renal transplant artery stenosis- value of the resistive index. *Acta Radiologica 35*(6). 586-589.

21. Sidhu, P.S., Ellis, S.M., Karani, J.B., & Ryan, S.M. (2002) Hepatic artery stenosis following liver transplantation: significance of the tardus parvus waveform and the role of microbubble contrast media in the detection of a focal stenosis. *Clinical Radiology 57*(9):.789-799.

22. Tublin, M., Ronald, B., and Platt. J. (2003). The resistive index in renal doppler sonography: Where do we stand? *American Journal of Roentgeneology 180*(4): 885-892.

23 Übeyli, E.D., & Güler, I. (2003). Neural network analysis of internal carotid arterial doppler signals: Predictions of stenosis and occlusion. *Expert Systems with Applications 25*(1): 1-13.

24. Übeyli, E.D., & Güler, I. (2005). Feature extraction from doppler ultrasound signals for automated diagnostic systems. *Computers in Biology and Medicine 35*(9):735-764.

25. Vit, A., De Candia, A., Como, G., Del Frate, C., Marzio, A., & Bazzocchi, M. (2003). Doppler evaluation of arterial complications of adult orthotopic liver transplantation. *Journal of Clinicial Ultrasound 31*(7): 339-345.

VITA

Justin Jay Louis Baraboo was born on August 16th, 1990, in Little Rock, Arkansas. He was educated in public schools and graduated from Jefferson City High School in 2008. He received a degree in physics and a degree in mathematics from Truman State University, cum laude and a member of Sigma Pi Sigma, in 2015.

Upon graduating, he began his master's program in computer science at the University of Missouri-Columbia in 2016. Upon completion of his degree requirements, Mr. Baraboo plans to continue his career of research and secondary education pursuing a Ph.D.