

RICE UNIVERSITY

Using Deep Learning and Macroscopic Imaging to Predict the Uniaxial Stress-Strain
Response of Glutaraldehyde-Fixed Porcine Aortic Heart Valve Leaflets

By

Luis Hector Victor

A THESIS SUBMITTED
IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE

Master of Science

APPROVED, THESIS COMMITTEE

K Jane Grande-Allen

Jane Grande-Allen
Professor of Bioengineering

Antonios Mikos

Antonios G. Mikos
Professor of Bioengineering & Chemical and
Biomolecular Engineering

Rafael Verduzco

Rafael Verduzco
Professor of Chemical and Biomolecular
Engineering, Materials Science and
Nanoengineering

Christopher Jermaine

[Christopher Jermaine \(Aug 22, 2022 11:56 CDT\)](#)

Christopher Jermaine
Professor of Computer Science

HOUSTON, TEXAS

August 2022

RICE UNIVERSITY

**A Deep Learning Approach for the Prediction of
Fixed Porcine Aortic Heart Valve Uniaxial
Stress-Strain Response from Macroscopic Imaging**

by

Luis Hector Victor

A THESIS SUBMITTED
IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE

Master of Science

APPROVED, THESIS COMMITTEE:

K. Jane Grande-Allen
Isabel C. Cameron Professor of
Bioengineering

Christopher Jermaine, Chair
Professor of Computer Science

Antonios Mikos
Louis Calder Professor of Bioengineering
and Chemical and Biomolecular
Engineering

Rafael Verduzco
Professor of Chemical and Biomolecular
Engineering, Materials Science and
Nanoengineering

Houston, Texas

August, 2022

ABSTRACT

A Deep Learning Approach for the Prediction of Fixed Porcine Aortic Heart Valve Uniaxial Stress-Strain Response from Macroscopic Imaging

by

Luis Hector Victor

Advancements in computational tools and manufacturing processes have catapulted the capabilities of device development in various fields and industries. For example, pairing artificial intelligence optimization, classification, and discovery tasks with additive manufacturing, one can rapidly prototype and optimize device designs for biomedical applications.

One biomedical device that can benefit from the synergy of the tools previously mentioned at the design and development stage are prosthetic heart valves (HVs). HVs facilitate the flow of blood in the heart, thus enabling the delivery of oxygen and nutrients to the body. Native HVs consist of leaflets that normally open and shut, guaranteeing unidirectional flow; however, their proper function is limited by their degradation due to a range of disease processes. For this reason, the study of leaflet mechanics is important for understanding the effect of cardiovascular diseases and designing prostheses.

Although traditional mechanical testing of heart valve leaflets (HVLs) is the standard for evaluating mechanical behavior, it is time-intensive, tedious, technical, and requires specialized and expensive equipment. On the other hand, computational resources and imaging are readily available and are more cost-effective. However, de-

spite accessibility to imaging and deep learning (DL) networks, such as convolutional neural networks (CNNs), investigators have not leveraged these tools to study the mechanics of HVLs.

In this work, a dataset consisting of uniaxial tensile tests on porcine aortic valve (PAV) leaflets and imaging of their aortic surface, was curated to determine the ability of a CNN to predict the stress-strain response of the leaflets. Due to inherent limitations of using a CNN to predict continuous data, a third-degree polynomial was used to fit the data. From these fits, the coefficients are extracted and used as the ground truth. The findings indicate that the presented framework can accurately predict the polynomial coefficients needed for reconstructing the toe and linear regions from macroscopic imaging of HVLs, which can significantly reduce the amount of time that it takes to evaluate the mechanics of HVLs and devices, effectively speeding up the decision-making and design processes. Additionally, by making the curated data publicly available, other researchers will be able to explore and refine deep learning architectures to form better predictions and find other applications.

Contents

Abstract	ii
List of Illustrations	vi
List of Tables	xii
1 Introduction	1
1.1 Goals and Objectives	1
1.2 Motivation	3
2 Literature Review	4
2.1 The Role of HVLs in the Heart	4
2.2 Morphological Features and Mechanics of HVLs	6
2.3 Acquisition of Mechanical Properties Using Mechanical Testing	11
2.4 Deep Learning and its Application in the Prediction of Physical Properties	13
2.5 A Brief, Closer Look at Convolutional Neural Networks	15
3 Methods	19
3.1 Preparation of HVL Samples	19
3.2 Macroscopic Imaging	20
3.3 Mechanical Testing	21
3.4 Data Preprocessing and CNN Framework	23
3.5 Experiments	26
4 Results	29

4.1	Curve Reconstruction for Ground Truth Collection	29
4.2	Coefficient Prediction	31
5	Discussion	39
5.1	Preliminary Findings	39
5.2	Curve Reconstruction	40
5.3	Prediction Results	41
5.4	Conclusion	42
	Bibliography	44

Illustrations

- 1.1 Example sets of stress-strain data to failure per HVL, with stress (MPa) plotted on the y-axis and strain on the x-axis. Note the similarities and differences between the HVLs and their respective stress-strain responses. 2

- 2.1 Flow of deoxygenated blood in blue arrows and oxygenated blood in red arrows through the heart shown on the left. On the right, the aortic valve's leaflets in the open and closed positions during systole and diastole, respectively, under normal conditions. 4
- 2.2 Illustration of the flow of blood through the aortic valve during systole and diastole under normal conditions. 5
- 2.3 One of three leaflets of an aortic heart valve. The nodule and coaptation regions on the free end come together with those of the other two heart valve leaflets, while the basal attachment of each leaflet is attached to the walls of the heart. The axes shown indicate the radial and circumferential directions of the leaflet with an 'R' and 'C', respectively. The leaflet shown here is approximately 14 mm long. 6
- 2.4 The aortic valve's leaflets mostly consist of three layers: fibrosa, spongiosa, and ventricularis. Each layer has a characteristic component as shown above. 7

2.5	Collagen fiber arrangement in a heart valve leaflet. The collagen fibers collect into thick bundles that are aligned with the circumferential axis of the leaflets.	8
2.6	During systole, collagen in the fibrosa is crimped and unaligned, producing wrinkles along the arterial surface, while elastin in the ventricularis is relaxed as blood pushes past the valve. During diastole, as blood attempts to rush back through the valve, the leaflets shut close and the collagen becomes uncrimped, while the elastin is also stretched.	10
2.7	The toe region occurs during the collagen uncrimping phase, while the linear transition region occurs during the collagen stretching portion.	11
2.8	The in-house CNN, by Liang L et al. used for predicting the biaxial properties of glutaraldehyde-treated bovine pericardium from second harmonic generation images.	15
2.9	A general CNN architecture showing the input, convolutional layers, fully connected layers, and output. The flow of data is from left to right through forward propagation until the loss is computed, at which point back propagation occurs and the kernels and weights of the CNN layers are updated. The cycle then repeats for a defined number of times and a minimum in the loss is determined.	16
2.10	Alexnet architecture, consisting of 5 convolutional layers and 3 fully-connected layers. Here the changes in the shape of data is clear, as the input image of size 227x227x3 is transformed to a shape of 1000x1 at the output.	17

2.11	a) Resnet building block showing the direct mapping between the input of two weighted layers and the output. The resulting residual learning block shown then consists of an identity mapping and a residual that is learned. b) Comparison of a VGG, plain CNN, and Resnet architectures. The arched arrows in the Resnet represent the residual learning blocks as shown in a)	18
3.1	An HVL that has been sectioned into a rectangular shape with the long axis parallel to the circumferential axis of the leaflet. The ends are secured under two wooden tabs on either side with Loctite 495. The gauge length is about 8 mm	20
3.2	The Leica stereomicroscope used to acquire macroscopic imaging of HVLs shown on the right, while the left shows an HVL image that has not been processed.	21
3.3	A sample is placed between the grip heads of the EFPB. The lock represents the grip head that is kept in-place during mechanical testing, while the 'T' represents the grip head that is moving and placing load on the sample. The sample is in solution kept at 37°C.	22
3.4	The entire procedure for ground truth collection involves extraction of the HVLs, sample preparation, and mechanical testing. The alignment of the sample in the EFPB is critical so as to not impose unwanted forces on the sample.	23
3.5	The stress-strain of a sample is shown above, displaying up to 60% strain. This sample experienced its first macro failure at 30% strain, as marked by the dashed vertical red line, and later a second macro failure slightly after 50% strain. The boxed region represents the portion of the data of interest for the polynomial fit, which seems to be just before the yield point.	24

- 3.6 During **training**, the input image, which is the arterial surface of a single sample, is transformed into many smaller images through data augmentation. Then, these smaller images are the input to the CNN and the labels $[a_1, a_2, a_3]$ for each of these are the same. 26
- 3.7 During **testing**, the input image undergoes the same data augmentation as in the training phase. However, the CNN will predict unique $[a_1, a_2, a_3]$ for each of the smaller images and then the average of these are taken as the overall sample's coefficients. 27
- 4.1 Prior to the extraction of the ground truth coefficient values, the stress-strain data was fitted with reconstruction curves. Stress-strain data of two different samples are shown in blue, while the reconstructed curves from a three-degree polynomial fit are shown in red, dashed lines. These plots illustrate the ability to reconstruct the data with a polynomial alone up to the yield point. However, the yield point can vary greatly from sample to sample, as can be seen for the samples above. 30
- 4.2 The evolution of the polynomial fit, or the reconstructed curve, as higher strain was considered. On the left, the strain was capped to the physiological strain, while on the right, the strain was capped to 30% strain. In both cases, the RMSE was relatively low, however, the reconstructed curve on the right does not correctly describe the behavior of the actual stress-strain data. 31
- 4.3 The RMSE between the reconstructed curve and the actual stress-strain data of the same sample shown in Figure 4.2 at various strain caps. Although the RMSE was higher at 30% strain than at 10% strain, the RMSE was still relatively low. 32

4.4 Two predicted curves, one resulting from the small sample set, 8, in orange and the other from the larger sample set, 51, in yellow from the first experiment. The ground truth is shown in blue, which is closely aligned with the predicted curve from the large sample set. The physiological strain is marked by a vertical red, dashed line. Note that the RMSE and R^2 is lower for the predicted curve from the larger sample set. 33

4.5 A close up of the small sample set, 8, predicted curve in orange and the ground truth shown in blue from the first experiment. The physiological strain is marked by a vertical red, dashed line. Note that the RMSE and R^2 was lower for the predicted curve since less strain is considered than it was in Figure 4.4. 34

4.6 Two predicted curves, one from setting a strict strain cap to 15% strain while the other was from a loose strain cap, which were hand-picked values near yield points, as seen from the second experiment. The physiological strain is marked by a vertical red, dashed line. In this plot, data was considered up to 14% strain and it is clear that the predicted curve from a loose strain cap not only has a higher RMSE and R^2 , but also predicts incorrect behavior of the stress-strain relationship. 36

4.7 Current capabilities allow for the prediction of the toe and linear transition regions as illustrated above. The high variation in the yield point from sample to sample makes it more difficult to predict the stress-strain curve closer and past the yield point. However, to physiological strain, the predictions are good as seen from Table 4.2. 37

4.8	Two predicted curves, one from excluding the cubic term and the other from including all three terms of the polynomial, as seen from the third experiment. Exclusion of the cubic term gives a lower RMSE and R^2 . The physiological strain is given by the vertical red, dashed line.	38
-----	---	----

Tables

4.1	In the first experiment, a lower MAE when using a larger set of samples, 51, was observed as compared to when using a fewer set of samples, 8.	32
4.2	In the second experiment, there was higher stress-strain variation at higher strain thresholds, which resulted in higher MAE regardless of architecture used as the strain cap increased.	35
4.3	In the third experiment, there was a higher accuracy when excluding the cubic term, since the cubic term has high variations, especially at higher strain.	37

Chapter 1

Introduction

1.1 Goals and Objectives

HVLs are arguably one of the most important structures in the heart, acting as one-way valves to facilitate the flow of blood and prevent oxygenated blood from mixing with deoxygenated blood [1]. However, diseases such as stenosis can impact the leaflets, altering their mechanical properties and ability to function properly [2]. Thus, to evaluate and approach the functionality and durability of natural HVLs, studying their mechanics is critical. However, the standard for acquiring the mechanics of HVLs relies on mechanical testing - a time-intensive process that requires specialized equipment, technical expertise, and, often, the destruction of the samples being tested.

On the other hand, despite the availability of accessible imaging methods and DL networks, no work has considered leveraging these two tools to study and approximate the mechanics of HVLs. CNNs are one subset of DL networks that would be advantageous to use because they are especially designed to extract features from images [3].

Therefore, due to the heterogeneous, macroscopically-visible morphology of HVLs and the variation in individual leaflets' stress-strain response, as shown in Figure 1.1, macroscopic imaging and deep learning tools could be used to predict mechanical behavior - minimizing the need for advanced imaging tools and traditional mechanical testing. Thus, in this work, CNNs and macroscopic imaging of the HVLs were used

to predict the stress-strain response of the leaflets under uniaxial tension. Since there was no existing dataset for the investigation proposed in this work, priority was given to the curation of a dataset that could then be used to explore and optimize a CNN architecture for stress-strain prediction. Additionally, since a relatively low number of samples' data was collected - as far as traditional computer vision problems are concerned - pre-trained CNNs were used to address this issue earlier on in the work.

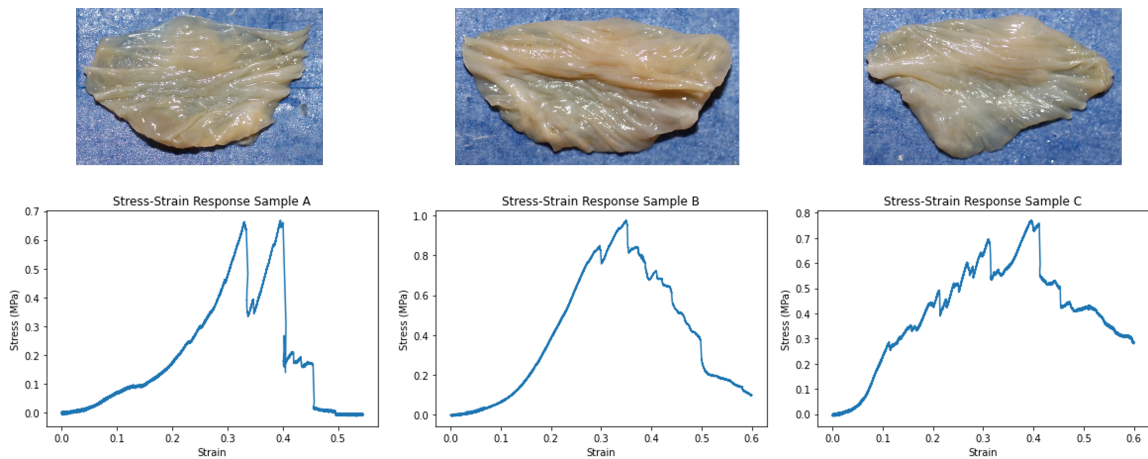


Figure 1.1 : Example sets of stress-strain data to failure per HVL, with stress (MPa) plotted on the y-axis and strain on the x-axis. Note the similarities and differences between the HVLs and their respective stress-strain responses.

In this thesis, **CNNs were used to predict the mechanical properties of fixed porcine aortic heart valves** as an example of how this approach can be used in cases where a tissue's morphological features are macroscopically visible. The overall objective is to demonstrate the capability and accessibility of the proposed framework to quickly and accurately predict what the stress-strain curves are for different HVLs.

1.2 Motivation

Convolutional Neural Networks (CNN) have shown excellent capability in optimization, classification, regression, and other similar tasks in various fields from computer vision to bioengineering [4, 5, 6]. Their architecture allows them to excel at working with images, which not only place significance on the value, such as the color or 'intensity', of a pixel, but also its location in the image.

Within biomechanics, the morphology of tissue plays a major role in the overall mechanical response of a sample. In the case of HVLs, despite the complexities present within their tissue, the mechanics are mostly dominated by a component that produces folds along the aortic surface which are visible to the human eye. This observation produced the hypothesis that a CNN supplied with macroscopic imaging of an HVL's surface can predict the mechanical response of the sample.

This work could provide the first steps in being able to deploy a mobile resource capable of predicting the most critical component of the mechanical response of HVLs, prior to yield and failure, and other tissue whose mechanical properties are exposed through macroscopically visible morphology. This is valuable for researchers investigating native tissue and prosthetic designs, while also offering pathological relevance.

To fully appreciate the value of this work, a brief review on the mechanics of HVLs and the acquisition of their mechanical properties using mechanical testing is provided. This information is also necessary for understanding the collection of the ground truth in this work. Then, CNNs are briefly covered and their application in similar works. Then, the methods and experiments used in this investigation are explained, followed by the results which show the capability of the proposed framework.

Chapter 2

Literature Review

2.1 The Role of HVLS in the Heart

The heart consists of four valves that control the inflow and outflow of blood from the heart and the basic operation can be summarized in four steps: 1) oxygen-poor blood enters the right ventricle through the tricuspid valve, 2) the tricuspid valve closes and the heart pumps the blood from the right ventricle to the lungs through the pulmonary valve, 3) the blood then becomes oxygen-rich and enters the left atrium of the heart through the mitral valve, and finally 4) the heart pumps the blood from the left ventricle to the body through the aortic valve [7]. This process is shown in Figure 2.1.

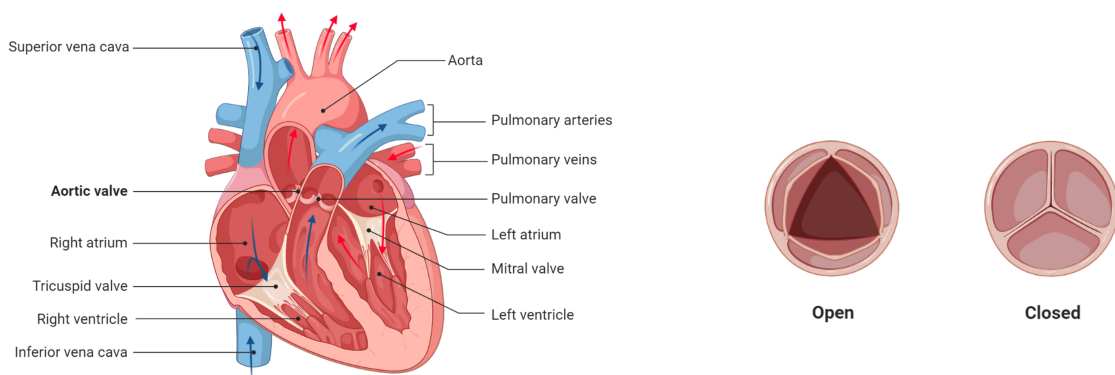


Figure 2.1 : Flow of deoxygenated blood in blue arrows and oxygenated blood in red arrows through the heart shown on the left. On the right, the aortic valve's leaflets in the open and closed positions during systole and diastole, respectively, under normal conditions.

The load that leaflets experience in vivo is complex – experiencing tensile, shear, bending, and compressive stresses [8] – yet, the peak load under normal conditions is well below the load that would be required for the HVLs to experience failure, or tearing of the tissue.




Figure 2.2 : Illustration of the flow of blood through the aortic valve during systole and diastole under normal conditions.

During systole, the heart contracts and blood rushes past the valve. Then, during diastole, the heart relaxes and the aortic valve closes tightly to prevent any blood from reentering the left ventricle, as shown in Figure 2.2. It is then during diastole that the aortic valve experiences peak load, requiring a load transfer from the leaflets to the walls of the heart. Through those cyclical deformations, the leaflets undergo changes in area by up to 50% in the radial direction with stresses typically in the range of 200-300 kPa [9, 10].

2.2 Morphological Features and Mechanics of HVLs

The aortic valve is a complicated structure, separating the left ventricle from the aorta, and the valves can be thought of as multi-layered, strong, thin strips of tissue, about 1 mm in thickness [11, 12].

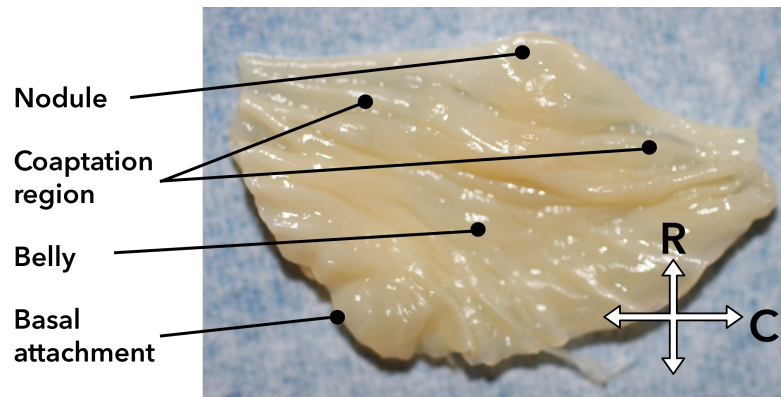


Figure 2.3 : One of three leaflets of an aortic heart valve. The nodule and coaptation regions on the free end come together with those of the other two heart valve leaflets, while the basal attachment of each leaflet is attached to the walls of the heart. The axes shown indicate the radial and circumferential directions of the leaflet with an 'R' and 'C', respectively. The leaflet shown here is approximately 14 mm long.

The aortic valve has three leaflets, or cusps, which are the valve's central feature. These leaflets consist of four parts each: a hinge, belly, coaptation surface, and lannula (free end) with a nodule of Arantius [13], see Figure 2.3. The nodulus of each leaflet can be found at the center of its free end toward the exterior along the coaptation surface, or where the three leaflets meet upon closing. On either side of the nodule, a thin semi-circular-shaped portion forms the lannula, which has a thin margin at its free end and continues in the coaptation region to about one third of the leaflet height. The three leaflets meet at the lannula upon valve closure. Continuing down the leaflet is the main portion of the leaflet, the belly. The leaflets' arrangement of collagen is macroscopically visible throughout these regions and the structure seems almost

transparent. Finally, the leaflets are attached to the annulus, which is connected to the surrounding heart muscle or wall, through the hinge area. This area is composed of thick collagenous bundles that cross into the aortic wall and are responsible for load transfer from the leaflets to the wall.




Figure 2.4 : The aortic valve’s leaflets mostly consist of three layers: fibrosa, spongiosa, and ventricularis. Each layer has a characteristic component as shown above.

Then, taking a cross-section of the leaflet reveals three overall layers: the lamina radialis (or ventricularis), spongiosa, and fibrosa, where the lamina radialis faces the ventricular side, and the lamina fibrosa faces the arterial side, see Figure 2.4. Further zooming into the layers, it is clear that the collagen and elastin fibers have a particular orientation and arrangement, which result in a leaflet’s highly anisotropic behavior, see Figure 2.5, from Fig. 2 in [14].




Figure 2.5 : Collagen fiber arrangement in a heart valve leaflet. The collagen fibers collect into thick bundles that are aligned with the circumferential axis of the leaflets.

These fibers are in mechanical contact with each other, through a sponge-like configuration which maintains collagen fiber orientation and geometry after load has been released [15, 16, 17]. The arterial surface (fibrosa) contains thick bundles of circumferentially-aligned type I collagen fibers [11], which result in the macroscopically visible folds or wrinkles that run parallel to the free end of the leaflets, and it is this configuration that transfers load from the leaflets to the annulus and wall [13].

All three layers contribute to the overall mechanics of the leaflet, due to the major components that compose each one. The radialis/ventricularis is mostly collagenous with elastin fibers running along the radial axis, the spongiosa layer contains loosely arranged collagen and has an abundance of glycosaminoglycans (GAGs), and the fibrosa contains a majority of circumferentially-aligned collagen type I fibers. Thus, the three layers' dominant features come together to optimize performance of the valve, where the radialis' elastin fibers extend and contract in diastole and systole,

respectively. The spongiosa's GAGs and loose collagen provide shock absorption between its adjacent layers during the cardiac cycle. Finally, the fibrosa's thick collagen bundles provide strength and rigidity to maintain coaptation during diastole [9].

Furthermore, the structures of the leaflets change during physiological function. In the relaxed state, the valve leaflets are flexible, but when fully extended in the stressed state, they are inextensible to the peak loads experienced during systole and diastole. That is, the complex arrangement of the fibers allow the leaflets to undergo elastic deformation during the cardiac cycle in the radial direction more than in the circumferential direction, and, at full extension, the fibers are allowed to take on high loads and transfer it to the walls.

Due to the elastic deformation from the relaxed state, the internal fibers of the leaflets experience rearrangement, an important note to consider when one is to conduct mechanical tests on leaflets, which prepare the leaflet for withstanding peak loads while experiencing minimal stress during this initial extension. Figure 2.6 provides a representation of the leaflet's structural changes during systole and diastole and the interplay between collagen and elastin managing load, from [18].

Most importantly, experimental data of leaflets will experience a 'toe' region, where the collagen fibers uncrimp, causing a nonlinear stress-strain response. Then, there is a transition region, in which the stress-strain response becomes more linear and, finally, a collagen region. In the final region, microstructure failure starts to occur, where individual fibers start to fail and tear, followed by macroscopic failure, where tears at the macro level start to occur. This characteristic stress-strain behavior is shown in 2.7, from Fig. 1 in [9]

Although the HVLs have mechanosensitive cells, which are valvular endothelial cells (VECs) and valve interstitial cells (VICs), the leaflets' mechanical properties

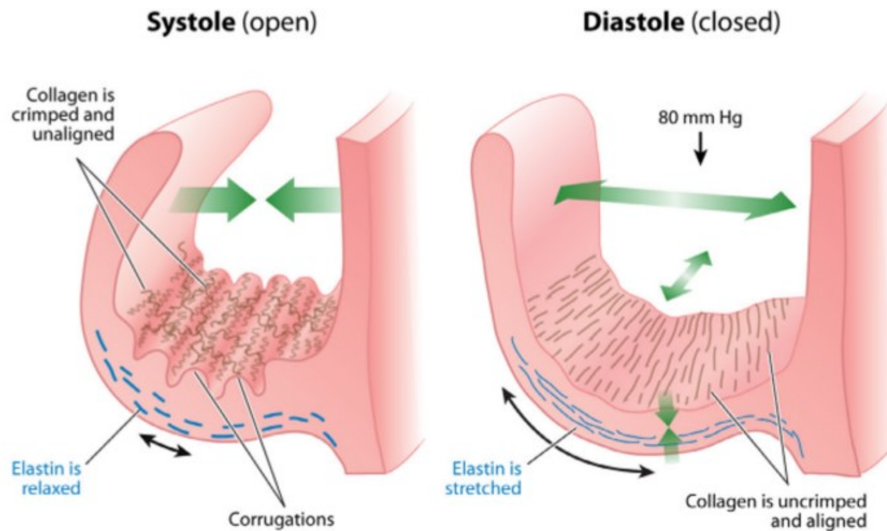


Figure 2.6 : During systole, collagen in the fibrosa is crimped and unaligned, producing wrinkles along the arterial surface, while elastin in the ventricularis is relaxed as blood pushes past the valve. During diastole, as blood attempts to rush back through the valve, the leaflets shut close and the collagen becomes uncrimped, while the elastin is also stretched.

are still largely dominated by the multi-layer structure that they are composed of. In the case of uniaxial or biaxial testing, the fibrosa layer is largely responsible for upholding the mechanical response outputted by the leaflet [19, 20]. Thus, uniaxial tensile mechanical tests along the circumferential axis of the leaflets highlights the leaflet's most mechanically robust axis.

Finally, porcine aortic valves have been used as bioprosthetic valves due to their similar dimensions and operation when compared to human aortic valves [21]. Porcine hearts are typically treated with chemicals, such as glutaraldehyde, to avoid immune rejection from the host [22] and to preserve the leaflets [10, 23]. The leaflets that were used in this work were also fixed with glutaraldehyde, since these have a longer shelf life, allowing for the development of procedures and methods for fresh hearts or




Figure 2.7 : The toe region occurs during the collagen uncrimping phase, while the linear transition region occurs during the collagen stretching portion.

other time-sensitive studies.

When HVLs are fixed with glutaraldehyde, this chemical effectively crosslinks the present collagen. This causes the mechanics of the leaflets to change, essentially degrading the leaflet's microstructure's ability to extend itself due to load and, thus, become less compliant [9], though this also depends on the fixation method [23].

2.3 Acquisition of Mechanical Properties Using Mechanical Testing

Extracting the mechanical properties from HVLs is a costly endeavor because of a number of factors including sample and instrument preparation. The mechanical testing instrument available dictates what type of mechanical test can be performed, as different tests will reveal different mechanical properties. Then, experience is needed in extracting valves from the heart, sectioning off the leaflets, and applying a grip

method that is compatible with the mechanical testing instrument. Preparation of the samples must be done carefully since damage to the leaflets can lead to compromised results during mechanical testing.

Then, instrument preparation involves the attachment of the appropriate grip heads, loading mechanical test and data collection protocols/files, media preparation, and ensuring operation of the load cell (i.e. the mechanical load probe).

In general, preparation of the samples is difficult because aortic HVLs are small and like to fold/curl up due to their morphological structure. Adult aortic valves, not the leaflets, have a diameter, circumference, and area of about 2.3 cm, 7.3 cm, and 4.5 cm^2 , respectively [24]. As mentioned previously, the thickness is taken to be about 1 mm.

In literature, there is difficulty in finding a standard method for sample preparation and testing. This problem mostly arises because there is difficulty in gripping the samples without causing damage. Three methods are commonly seen in literature in which samples are prepared and grasped with sutures, clamps, or rakes. Sutures and rakes can more easily apply unwanted stresses and even tear the sample, since both methods must pierce through the sample [25, 26]. Clamps are typically used in uniaxial testing, and sutures and rakes are most typically used in biaxial testing [27, 28, 10], which is reasonable since clamps, unless in a cruciform shape [25, 26], do not allow the sample to extend appropriately in biaxial testing. Clamping methods may allow for a higher deformation range when performing uniaxial testing as opposed to using sutures, while, there are also indications that there are higher stress values at yield when using sutures [26, 10, 23, 20]. Thus, a standard uniaxial tensile testing protocol using a clamping method is used in this work, as it is an accepted method and one that has been used extensively in the past by the Grande-Allen Integrative

Matrix Mechanics Lab [29, 30].

2.4 Deep Learning and its Application in the Prediction of Physical Properties

DL and data-driven methods have been increasing in demand across disciplines as more data is becoming available [31, 32]. In particular, using these approaches have contributed to a wide range of topics such as optimization, classification, and discovery tasks. For example, in bioengineering applications, advancements in computational tools for modeling and manufacturing processes have catapulted available datasets. Pairing these technologies with data-driven methods aimed at optimization, one can rapidly prototype and optimize device designs and tools. Furthermore, one way to evaluate devices and materials is through the characterization of their mechanical properties, so it is not surprising that many works have already explored related questions to the one at hand. However, there is no open dataset that is currently available to investigate the proposed work presented here.

As the main objective is to determine if a DL network paired with imaging of tissue structure can be used to predict stress-strain data, there is an additional effort to contribute to the field by collecting the first set of data to explore such a task. This could be useful for many applications in biomedical engineering or materials science.

In brief, DL networks are function approximators. They are special because they are universal function approximators that guarantee an output from any input, assuming the right network is employed [33]. This means that DL networks can be useful tools to use when the data is constrained and when there are unknowns, even unknown functions, in the framework to be used. In the present work, the limited

available morphological data is suspected to be related to the stress-strain response. One drawback, however, is that typical problems require a large amount of training data. However, there are methods for data augmentation and obtaining surrogate data from computation methods that can aid in this, placing data-driven modeling as a candidate for a new paradigm [34]. In this work, data augmentation is used as a way to handle the relatively low number of samples available for training, when compared to other DL problems such as in computer vision. Furthermore, although physics-informed models or other more sophisticated models are not used in this work, the integration of these are readily available with the collection of ground truth completed and these will be explored in the future.

As mentioned previously, there are other works that have used DL to explore physical sciences. For example, Yixuan S. et al. [35] use a DL network to predict the mechanical properties from microstructure images in fiber-reinforced polymers. In that work, the authors specifically use a CNN, which has a rich history in computer vision problems and tasks involving images [36]. The CNN is trained and evaluated from finite element simulation data that was generated from exact microstructures of the composite materials. Their network, after being fully trained, is able to make predictions within seconds in comparison to almost a hundred hours for a finite element simulation. The use of CNNs for this type of application is not new, since previous researchers had seen success in predicting several material properties from material topology because of their ability to learn nonlinear relations and account for spatial and intensity information from the images [37, 38].

Liang L et al. [39] took a similar approach as the one developed in this work. Liang used an in-house CNN, see Figure 2.8, and images of collagenous tissue to predict the biaxial stress-strain response of a given sample. However, the images that

Liang used were taken from second harmonic generation (SHG) imaging. Although using this imaging method gives access to microscopic details of the collagen fibers as well as through the thickness of the sample, it does not account for macroscopic information and is less accessible than using a camera or a stereomicroscope.




Figure 2.8 : The in-house CNN, by Liang L et al. used for predicting the biaxial properties of glutaraldehyde-treated bovine pericardium from second harmonic generation images.

Additionally, Liang et al. performed biaxial testing, whereas uniaxial testing was used in this work. Liang argues that, with the collection of biaxial data, principal component analysis can be used for the reconstruction of the two stress-strain curves. This meant that they would predict the coefficients associated with individual samples, while the modes of variation were the same for all samples.

2.5 A Brief, Closer Look at Convolutional Neural Networks

CNNs are standard DL networks employed when images are used as input data, which is due to the importance being placed not only pixel values, but also on spatial information. The characteristic property of CNNs is that at a given convolutional layer, convolution filters or kernels produce feature maps that are normalized using activation functions and resized using pooling before data is taken to the next layer through forward propagation. The final output of a CNN is a single vector, achieved through a fully-connected layer. During training, the loss between the predicted

output and the ground truth is computed and used to improve the predictions through the back propagation. Figure 2.9, from Fig. 1 in [40], shows what a general CNN looks like as well as the flow of the data, from the input images to the output which is the target to predict.

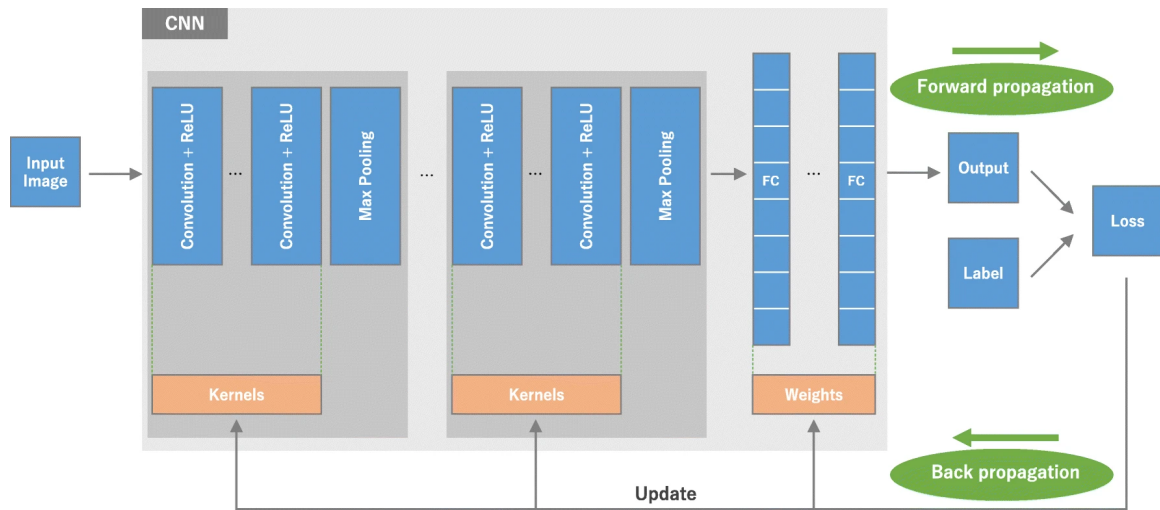


Figure 2.9 : A general CNN architecture showing the input, convolutional layers, fully connected layers, and output. The flow of data is from left to right through forward propagation until the loss is computed, at which point back propagation occurs and the kernels and weights of the CNN layers are updated. The cycle then repeats for a defined number of times and a minimum in the loss is determined.

CNNs are typically designed and tuned to specific applications, despite their ability to act as universal function approximators. The selection of types of layers, sizes, quantity, and other architectural features and their combinations are typically explored to tune the network in the specific application. Three different CNN architectures that have been used in literature include Alexnet [41, 42], VGG11 [43], and Resnet18 [44]. These can be pretrained on ImageNet [45], which means that the CNNs have already undergone the training process described previously with the images and classes included with the ImageNet dataset.

Despite the similarities shared between the previously mentioned CNN architectures, there are a few differences that are worth noting. Figure 2.10 shows an Alexnet architecture, as explored in [46], and it is clear that it mainly consists of 5 convolutional layers and 3 fully-connected layers, making it 8 layers deep.




Figure 2.10 : Alexnet architecture, consisting of 5 convolutional layers and 3 fully-connected layers. Here the changes in the shape of data is clear, as the input image of size $227 \times 227 \times 3$ is transformed to a shape of 1000×1 at the output.

A VGG11 network is very similar to an Alexnet, however in the VGG11 network, there are 11 layers total, 8 convolutional layers and 3 fully-connected layers. Additionally, the size of the kernels in VGG are fixed (3×3), whereas in Alexnet, the kernels vary in size (11×11 , 5×5 , 3×3), which ends up reducing the total number of parameters.

Resnet18 on the other hand, consists of 17 convolutional layers and 1 fully-connected layer. Although Resnet also has a fixed kernel of size (3×3) throughout the architecture, it also makes use of residual learning through shortcut connections [44]. In short, the shortcut connections make direct connections from an input to an output, which implies that there is an identity mapping x and the block needs only to learn the residual features from the input, as shown in 2.11. These direct shortcut

connections arose as a solution to a degradation problem that occurred as more layers were added to a deep CNN - as additional layers are added, the earlier layers become negligible since the information required to change the weights diminishes.

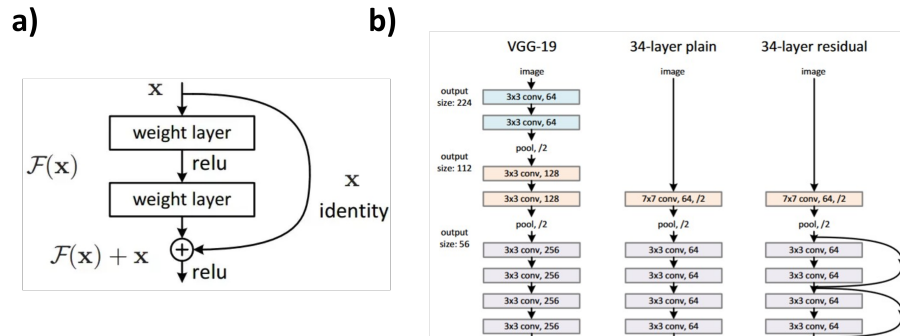


Figure 2.11 : **a)** Resnet building block showing the direct mapping between the input of two weighted layers and the output. The resulting residual learning block shown then consists of an identity mapping and a residual that is learned. **b)** Comparison of a VGG, plain CNN, and Resnet architectures. The arched arrows in the Resnet represent the residual learning blocks as shown in **a)**

Chapter 3

Methods

3.1 Preparation of HVL Samples

Prior to imaging or to mechanical testing, the porcine aortic HVLs fixed in 0.2% glutaraldehyde HVLs must be prepared. Although fixing tissue with glutaraldehyde affects the mechanical properties of the tissue, all of the samples used were fixed, thus eliminating this as a potential variable impacting the prediction results. These samples were chosen due to the strong similarities between porcine and human hearts, accessibility in obtaining porcine hearts, and the cost-effective method of fixing samples. Furthermore, fixing tissue allows for an increased viability for study over long periods of time.

In this work, a total of 51 samples were prepared, imaged, and tested. The samples had been previously extracted from porcine hearts and were continuously stored at 7°C until testing. All samples were then sectioned along the circumferential direction using a scalpel and wood handles, made from balsa wood strips of 1 x 5.3 x 10 mm, were glued at the edges for mechanical testing using Loctite 495. Figure 3.1 shows an image of a fully prepared sample, with wooden tabs attached to the ends of the sample. The gauge length and thickness are recorded for each sample and then they are stored in phosphate buffered saline (PBS) until they are ready for imaging and testing, which was no more than 24 hrs. after sample preparation. The average gauge length, width, and cross-sectional area was 7.178 ± 1.477 mm, 4.769 ± 0.690 mm,

and $4.769 \pm 0.690 \text{ mm}^2$.

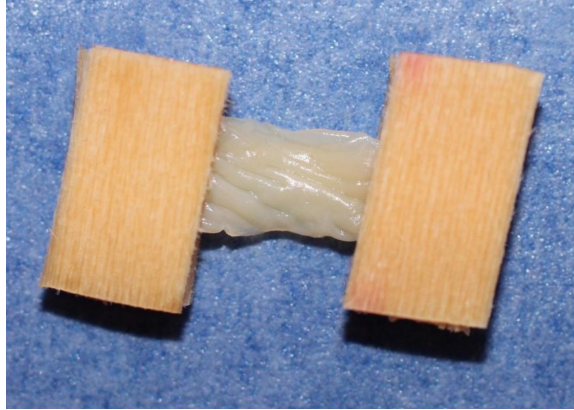


Figure 3.1 : An HVL that has been sectioned into a rectangular shape with the long axis parallel to the circumferential axis of the leaflet. The ends are secured under two wooden tabs on either side with Loctite 495. The gauge length is about 8 *mm*

3.2 Macroscopic Imaging

Imaging of the samples was done on a Leica MC170 HD mounted on a Leica MZ6 stereomicroscope, producing images at the millimeter scale in RGB. The images of sectioned leaflets here have a size of 2592x1944, however, the final images used consisted of cropped images of the sectioned HVLs. The images are cropped to minimize imaging of the wooden tab and the masking tape background. Figure 3.2 shows the imaging setup and an image that has not been processed for the CNN framework. The imaging setup remained consistent throughout samples.

Images of the leaflets were always taken on the arterial side of the leaflet (i.e. lamina fibrosa). As described in the previous section, samples were sectioned from whole leaflets on a cover slip with masking tape to hold the sample down, which is visible in the images taken on the stereomicroscope.




Figure 3.2 : The Leica stereomicroscope used to acquire macroscopic imaging of HVLs shown on the right, while the left shows an HVL image that has not been processed.

3.3 Mechanical Testing

Prior to testing, the gauge length, measured along the sample from one inner end of a wooden tab to the other, and thickness were recorded for all samples. Uniaxial mechanical tests were carried out on an ElectroForce Planar Biaxial (EFPB) instrument, which consists of two LM1 motors. The testing protocol consisted of a block test with two phases: a preconditioning phase with a 10% strain at 3mm/s for 10 cycles followed by a ramp/pull-to-failure phase at a rate of 1mm/s along the circumferential direction.

This testing is carried out by holding one LM1 motor and its grip head in-place, while the other LM1 motor and its grip head carry out the displacement as dictated by the block waveform. The load is measured by a 22N (1N = 100g) load cell that is fixed between the stationary LM1 motor and its grip head.

As described previously, small wooden handles were carefully secured onto the sample edges using Loctite Superglue to prevent the adhesive from affecting the sam-

ple along the gauge length. The sample edges within the handles remain at least at or less than half the width of the handles so as to avoid crushing the tissue upon placement in the EFPB's grips, which function similarly to C-clamps. The fully prepared samples were then placed in a bath of phosphate-buffered saline solution (PBS) maintained at 37°C by fastening each handle at the grips of the EFPB as shown in Figure 3.3.




Figure 3.3 : A sample is placed between the grip heads of the EFPB. The lock represents the grip head that is kept in-place during mechanical testing, while the 'T' represents the grip head that is moving and placing load on the sample. The sample is in solution kept at 37°C.

The raw data outputs of the mechanical testing consisted of five variable readouts: time (s), temperature (°C), displacement of grip head 1 (mm), displacement of grip head 3 (mm), and load (N), which is measured across a sensor mounted on one of the grip heads, see Figure 3. In the mechanical testing protocol, grip head 1 remains fixed in position while actuation occurs on grip head 3, which is constrained to move along a single axis as described above. For traditional mechanical analysis, two variables,

strain and stress, were computed by the displacement of grip head 3 and by the load and dimensions of each sample, respectively. Stress is calculated by using $\sigma = F/A$, where σ is the stress, F is the force, and A is the cross-sectional area. Strain is calculated by using $\epsilon = \Delta D/l_g$, where ϵ is the strain, ΔD is the displacement of the moving grip head, and l_g is the gauge length of the sample. However, rather than using this continuous data directly with the CNN, the curves are approximated with a three-degree polynomial and the fits are evaluated through the root-mean-squared error (RMSE), which is further described in the next section. Figure 3.4 shows the setup and procedure for mechanical testing.




Figure 3.4 : The entire procedure for ground truth collection involves extraction of the HVLs, sample preparation, and mechanical testing. The alignment of the sample in the EFPB is critical so as to not impose unwanted forces on the sample.

3.4 Data Preprocessing and CNN Framework

Stress-strain curves are calculated as described in the previous section, but rather than predicting the continuous data, a three-degree polynomial of the form $Y = a_1x + a_2x^2 + a_3x^3$ is used to reconstruct the curves to minimize the predictions that must be made by the CNN framework. That is, the CNN framework is first tasked

with learning the associations between the set of a_1, a_2, a_3 for an individual sample and the image of its arterial surface. Then, the CNN is tasked with predicting the three coefficients for a sample given only the image of its arterial surface. Figure 3.5 shows how a portion of the stress-strain data is selected for the polynomial fit.

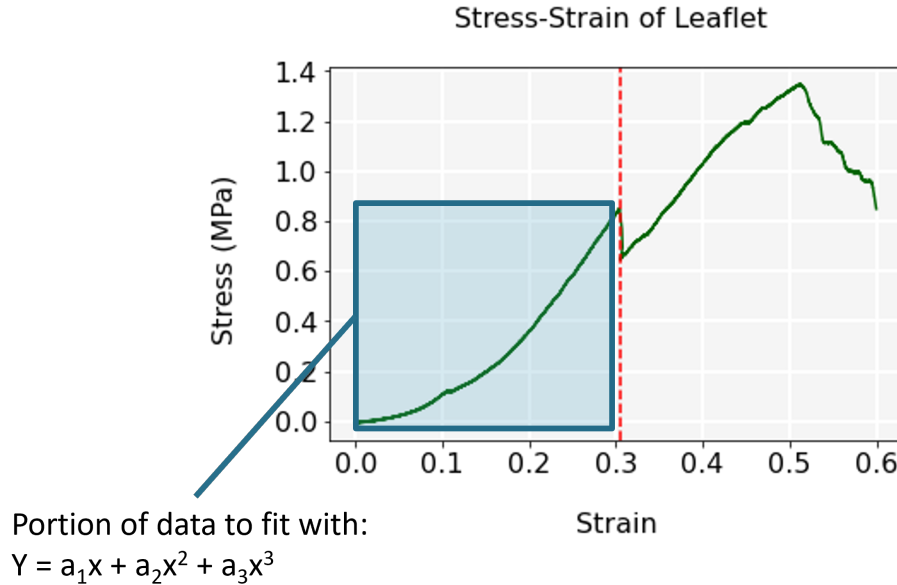


Figure 3.5 : The stress-strain of a sample is shown above, displaying up to 60% strain. This sample experienced its first macro failure at 30% strain, as marked by the dashed vertical red line, and later a second macro failure slightly after 50% strain. The boxed region represents the portion of the data of interest for the polynomial fit, which seems to be just before the yield point.

Since the CNN must predict the coefficient values, it is a regression task. The coefficients predicted from the CNN are evaluated with the mean absolute error (MAE), given by:

$$MAE = \frac{\sum_i^n |y_p - y_i|}{n},$$

where y_p is the predicted value and y_i is the actual value. Then, the goodness of fit produced by the predicted coefficients per sample is evaluated by measuring the R^2 ,

given by:

$$R^2 = \frac{\sum_i^n (y_i - y_p)^2}{\sum_i^n (y_i - \bar{y})^2},$$

and root-mean-squared error (RMSE) value, given by:

$$RMSE = \sqrt{\frac{\sum_i^n (y_i - y_p)^2}{n}},$$

where \bar{y} is the mean.

As with any DL application, there are many architectures that can be used, and the final choice depends on various elements such as state-of-the-art computational equipment and algorithms. The importance of this work rests on the ground truth collection and the validation of a CNN as a tool to quickly predict the mechanical properties of HVLs and allow for the continuation and progression of using related tools to this end. As such, three pretrained CNNs were selected for evaluating how well predictions could be made from the data collected and the overall framework. The CNN architectures that were pretrained on Imagenet [45] were: Alexnet [41, 42], VGG11 [43], and Resnet18 [44]. These were selected because they have been extensively used for a variety of tasks in bioengineering, from classification tasks [47, 48] to composite design [49], allowing focus to be given to the development and preparation of the data.

The input to the CNN consists of RGB HVL sample images, while the ground truth consists of the samples' stress-strain reconstructed curve coefficients. Additionally, a data augmentation technique that consisted of creating image patches (224x224) from the original image was used; each image patch derived from the same image had the same set of coefficients. These image patches were created by specifying the size of a window that served as an image patch's template and then translating it across the original image. After these transformations, the data was ready to be used

for training and testing. Figure 3.6 shows the proposed CNN framework during the training phase. This work was inspired by that of [39], however, in this work, using more accessible imaging was an objective, which provided imaging at the macroscale, and ultimately required a different framework.




Figure 3.6 : During **training**, the input image, which is the arterial surface of a single sample, is transformed into many smaller images through data augmentation. Then, these smaller images are the input to the CNN and the labels $[a_1, a_2, a_3]$ for each of these are the same.

Testing also involved image patch transformations on a test sample image, yielding three coefficient values, $[a_1, a_2, a_3]$, for these image patches. The test sample's overall coefficient values were then taken as the mean of each coefficient's value, $[a_{avg1}, a_{avg2}, a_{avg3}]$. The coefficients are predicted by passing the data from the fully-connected layers to a linear layer and these are then evaluated using the mean absolute error (MAE). Figure 3.7 shows how the overall coefficients for a testing sample are derived.

3.5 Experiments

In this work, three experiments were performed to validate the framework. It is important to note that the three experiments were done at various stages, as more data was curated and made available, however, only the final results are shown here as they reflect the final total number of samples.

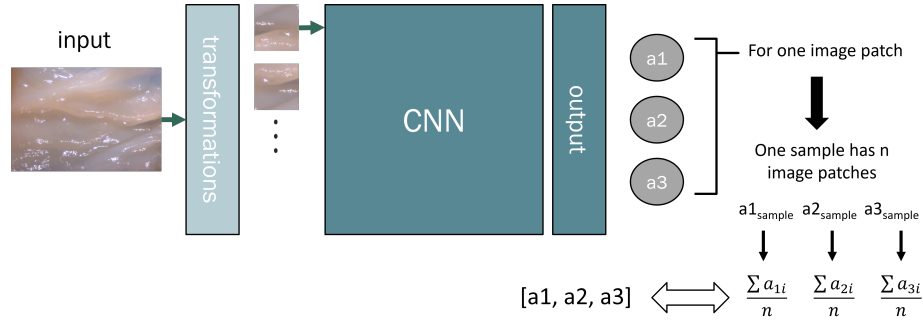


Figure 3.7 : During **testing**, the input image undergoes the same data augmentation as in the training phase. However, the CNN will predict unique $[a_1, a_2, a_3]$ for each of the smaller images and then the average of these are taken as the overall sample's coefficients.

In the first experiment, the performance capability is compared for when using a few samples' data, 8 samples, and when using a large samples' data, 51 samples. Here, 7 samples were used for training and 1 was used for testing in the smaller dataset, while 41 samples were used for training and 10 were used for testing in the larger dataset. All of the samples are capped to 10% strain. This was done because of the limited ability to collect many samples' data at a time, yet the CNN framework could be applied to the amount of samples' data available at any given time. This is an important metric to investigate because it can give insight as to how many samples are needed to be tested through mechanical testing - that is, ground truth collection - to obtain good predictions. Additionally, only the linear and quadratic terms were predicted for in the final investigations because of the results obtained in the third experiment. Furthermore, for the smaller set of data, a stride of 10 pixels in the x and y direction of the original images was used for data augmentation resulting in 22,000 images for training and 5,000 images used for testing. All of the images were 224x224. However, for the larger dataset, there was a stride of 25 pixels in the x and in the y direction of the original images, which resulted in 99,251 images for training

and 21,165 images used for testing. The image patching was different for the larger dataset due to limited computational resources as well as time.

In the next experiment, the investigation focused on how the coefficients' values differed with different strain values. The change in the coefficients' values for different strain thresholds, or caps, is due to the stress-strain data behavior. As the strain reaches and goes past the yield point, plastic deformation and microstructure failure begins to occur. It was suspected that higher variation in the coefficient values would ultimately affect error in the predictions, requiring investigation. Thus, the strain was capped at 0.05, 0.1, 0.15, and hand-picked values near an individual sample's yield point to observe the evolution of the reconstructed curves. This included a fit to the physiological strain (10% strain) [50, 51]. For this experiment, 41 samples were used for training and 10 samples were used for testing in all of the strain cap groups. Furthermore, predictions were made for the linear and quadratic coefficients.

The third and final experiment performed was to confirm the findings from the previous paragraph: variations in the coefficient values after yield increase greatly, so by reducing these variations, better predictions can be obtained. To do so, only the linear and quadratic terms' coefficients were predicted for because, prior to yield, the data seems to be linear and quadratic. Then, the resulting predictions are compared to when all of the terms' coefficients are predicted for. In this experiment, 41 samples were used for training and 10 samples were used for testing for the two cases. All of the samples were capped to 10% strain.

Chapter 4

Results

4.1 Curve Reconstruction for Ground Truth Collection

Stress-strain responses from HVLs are commonly reported up to the physiological strain of 10% in literature [50, 51]. In alignment with the second experiment, the strain at which to limit the curve reconstruction using a three-degree polynomial to fit the data was varied. Up to and past the physiological strain, the curve reconstruction scheme had low RMSE, however, incorrect behavior was observed when including data past or near the yield, which was defined as a "loose" threshold through manually picked strain values. Since the RMSE increased the closer and further away the strain was to the yield point, it was hypothesized and confirmed through experiment 2 and 3 that the current framework would perform better when limiting the data to the toe and transition regions.

In general, there is a good fit between the ground truth and the reconstructed curves, especially when constraining the data to or below the yield point. For example, in Figure 4.1, it is clear that excellent fits are achieved as the observed stress-strain data are shown in blue, while the fitted three-degree polynomials are plotted as a red, dashed line. Figure 4.1 also shows that the two samples reach yield at different strain values. For sample Z23 and Z29, the coefficients that reconstruct their stress-strain response, when considering the three term polynomial, are about $[0.9200, -7.4410, 24.9681]$ and $[0.7670, -17.5070, 224.6300]$, respectively. It is then obvious




Figure 4.1 : Prior to the extraction of the ground truth coefficient values, the stress-strain data was fitted with reconstruction curves. Stress-strain data of two different samples are shown in blue, while the reconstructed curves from a three-degree polynomial fit are shown in red, dashed lines. These plots illustrate the ability to reconstruct the data with a polynomial alone up to the yield point. However, the yield point can vary greatly from sample to sample, as can be seen for the samples above.

that the cubic term in these coefficients experience a very large amount of variation, which was hypothesized would lead to problems in the prediction performance of the CNNs and, thus, tested for this by excluding the cubic term in one set of predictions. Furthermore, since the yield point can vary greatly from sample to sample, it was decided to investigate the prediction performance by selecting various strain caps, including physiological strain.

Figure 4.2 shows the evolution of the fitted curve as different strain values are considered. Although the RMSE for both reconstructed curves are relatively low, as shown in Figure 4.3, only the reconstructed curve considering physiological strain, 10% strain, is an acceptable representation of the stress-strain data. This is because it is evident that including data equal to or greater than 30% strain leads to an unwanted trend in the data – the reconstructed curve for 0.3 strain is concave down, even when

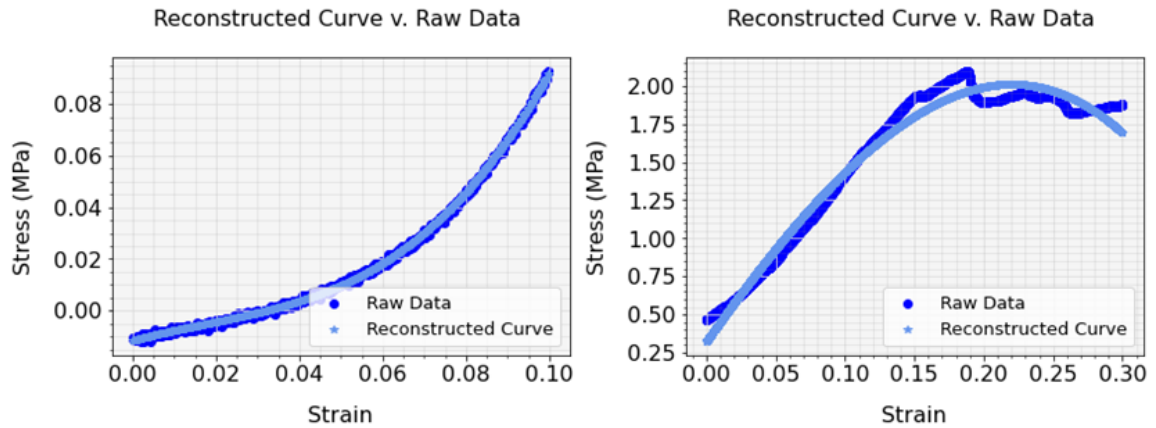


Figure 4.2 : The evolution of the polynomial fit, or the reconstructed curve, as higher strain was considered. On the left, the strain was capped to the physiological strain, while on the right, the strain was capped to 30% strain. In both cases, the RMSE was relatively low, however, the reconstructed curve on the right does not correctly describe the behavior of the actual stress-strain data.

considering strain up to the physiological strain, whereas the correct behavior is concave up, as can be seen in the raw data shown in dark blue.

4.2 Coefficient Prediction

In the first experiment, there was an increase in the accuracy of the predicted coefficients when using more samples, which was expected. For the small set, the resulting MAE were 6.88, 4.40, and 6.73 for Alexnet, VGG11, and Resnet18, respectively, while these had an MAE of 3.99, 3.66, and 3.81 for the larger set, see 4.1. To understand how this impacts individual sample predictions, the predicted reconstructed curves are observed and compared to ground truth and then described through the RMSE and the R^2 value.

For example, 4.4 shows the results for one sample when the data included extends to 25% strain, with an RMSE and R^2 value of 0.001 and 0.995 for the large dataset,

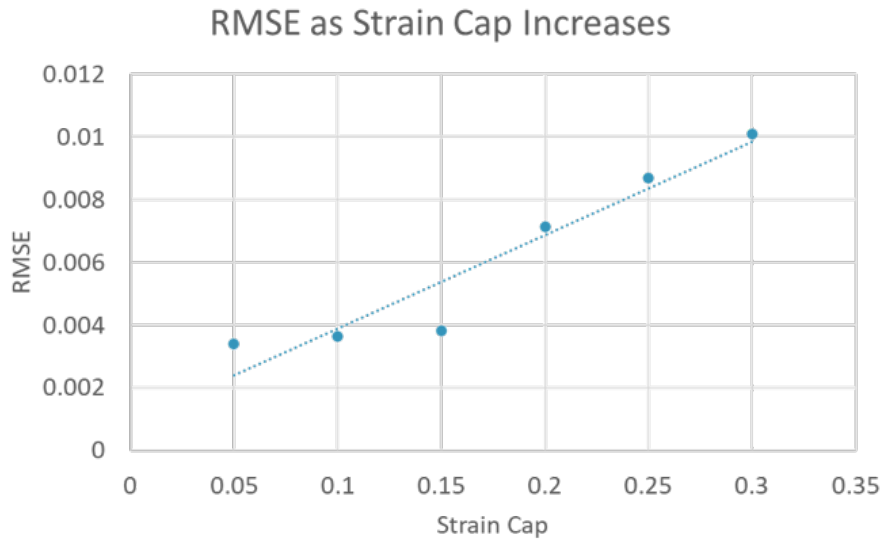


Figure 4.3 : The RMSE between the reconstructed curve and the actual stress-strain data of the same sample shown in Figure 4.2 at various strain caps. Although the RMSE was higher at 30% strain than at 10% strain, the RMSE was still relatively low.

Table 4.1 : In the first experiment, a lower MAE when using a larger set of samples, 51, was observed as compared to when using a fewer set of samples, 8.

Architecture	Large sample set (51) MAE	Small sample set MAE
Alexnet	3.99	6.88
VGG11	3.66	4.40
Resnet18	3.81	6.73

respectively, and 0.071 and 0.904 for the small dataset, respectively. Although the error was larger for the smaller dataset's prediction, it is clear that closer to the physiological strain, the error was much smaller. Figure 4.5 shows the results for the same sample when instead including data only up to 14% strain, and this results in

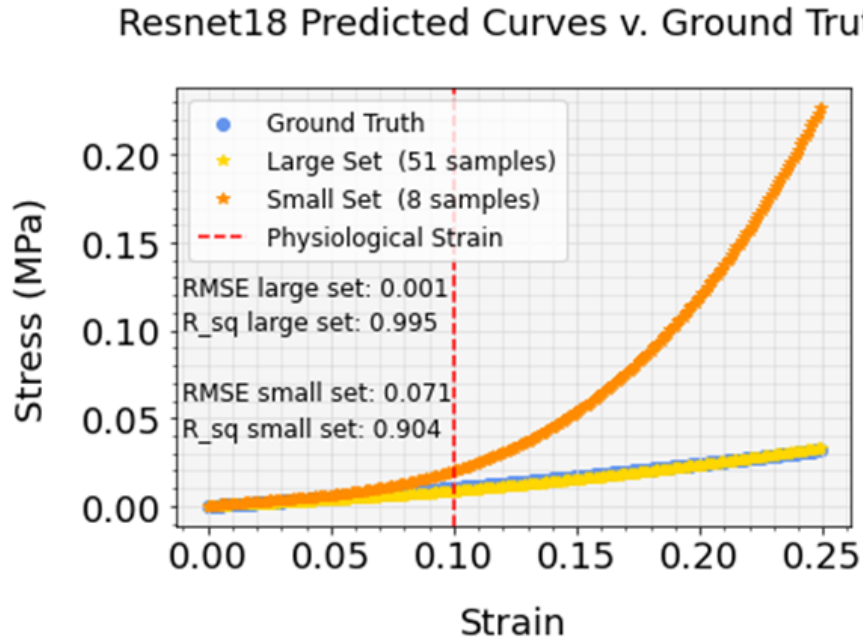


Figure 4.4 : Two predicted curves, one resulting from the small sample set, 8, in orange and the other from the larger sample set, 51, in yellow from the first experiment. The ground truth is shown in blue, which is closely aligned with the predicted curve from the large sample set. The physiological strain is marked by a vertical red, dashed line. Note that the RMSE and R^2 is lower for the predicted curve from the larger sample set.

an RMSE and R^2 value of 0.011 and 0.933, respectively, for the small sample dataset. Although having more data is better, mechanical testing is an expensive procedure and the ability to leverage data augmentation methods to use the relatively fewer number of samples in the small set to acquire better prediction results is optimistic for future work.

In the second experiment, it was confirmed that limiting the ground truth data to the toe and linear transition regions improves the predictions of the coefficients as shown in Table 4.2. Similar to what was seen when performing the curve reconstruction, as more data is included, through a greater strain cap value, the higher variations

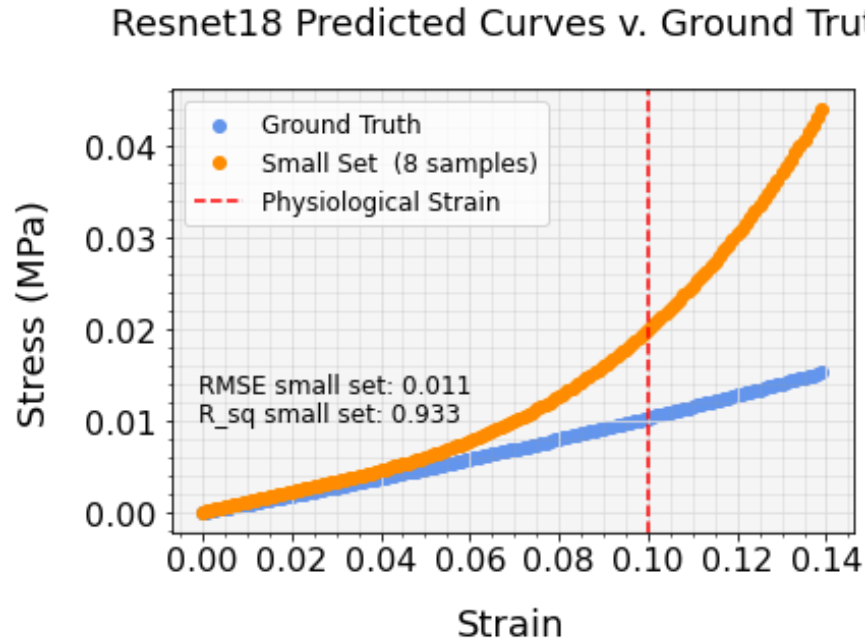


Figure 4.5 : A close up of the small sample set, 8, predicted curve in orange and the ground truth shown in blue from the first experiment. The physiological strain is marked by a vertical red, dashed line. Note that the RMSE and R^2 was lower for the predicted curve since less strain is considered than it was in Figure 4.4.

in the stress-strain response from sample to sample result in increased variations in the coefficient values. Furthermore, it becomes apparent that the yield point varies from sample to sample and this impacts the reconstructed curve and, ultimately, the coefficients to the point where the correct behavior of the leaflets can be incorrectly predicted for. For example, this can be seen for the sample shown in Figure 4.6, where the strict strain cap was set to 15% strain while the loose strain cap consisted of hand-picked values near the yield points.

However, to the physiological strain, this was not an issue from sample to sample, which is critical because of the physiological and clinical relevance - leaflets that have failed must be replaced. This means that, at minimum, predictions for up to the

physiological strain must be accurate, as illustrated in Figure 4.7. This would further mean being able to predict the toe and linear transition regions.

Table 4.2 : In the second experiment, there was higher stress-strain variation at higher strain thresholds, which resulted in higher MAE regardless of architecture used as the strain cap increased.

Architecture	Ground Truth Threshold - Strain %	Test MAE
Alexnet	5	2.40
VGG11	5	1.98
Resnet18	5	1.83
Alexnet	10	3.99
VGG11	10	3.66
Resnet18	10	3.81
Alexnet	15	6.97
VGG11	15	5.34
Resnet18	15	6.64
Alexnet	Loose	8.37
VGG11	Loose	8.36
Resnet18	Loose	7.90

This was further confirmed in experiment 3, where only the linear and quadratic terms were predicted for since these terms are capable of capturing the variance from the toe and linear transition regions in the reconstructed curves. For the case in which 3 terms were predicted, the MAE was 65.60, 59.05, and 66.75 for Alexnet, VGG11, and Resnet18, respectively, while these were 3.99, 3.66, and 3.81 when only

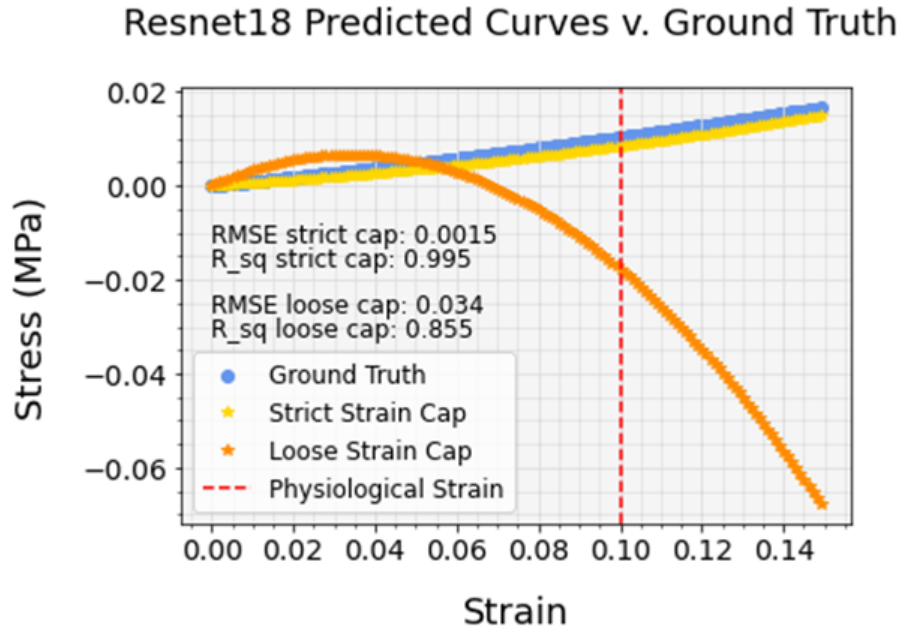


Figure 4.6 : Two predicted curves, one from setting a strict strain cap to 15% strain while the other was from a loose strain cap, which were hand-picked values near yield points, as seen from the second experiment. The physiological strain is marked by a vertical red, dashed line. In this plot, data was considered up to 14% strain and it is clear that the predicted curve from a loose strain cap not only has a higher RMSE and R^2 , but also predicts incorrect behavior of the stress-strain relationship.

2 terms were predicted, see Table 4.3. Although the error in the predicted coefficients was high when including the cubic term, this did not result in drastically incorrect stress-strain curves. For example, the RMSE and R^2 values for the sample shown in Figure 4.8 when including data up to 15% strain are 0.014 and 0.980, respectively, for the predicted curve when including the cubic term and 0.0016 and 0.995, respectively, when excluding the cubic term. However, it can be seen that excluding the cubic term results in improved reconstruction of the ground truth, which was the reason why only the linear and quadratic terms were used in the first and second experiments.

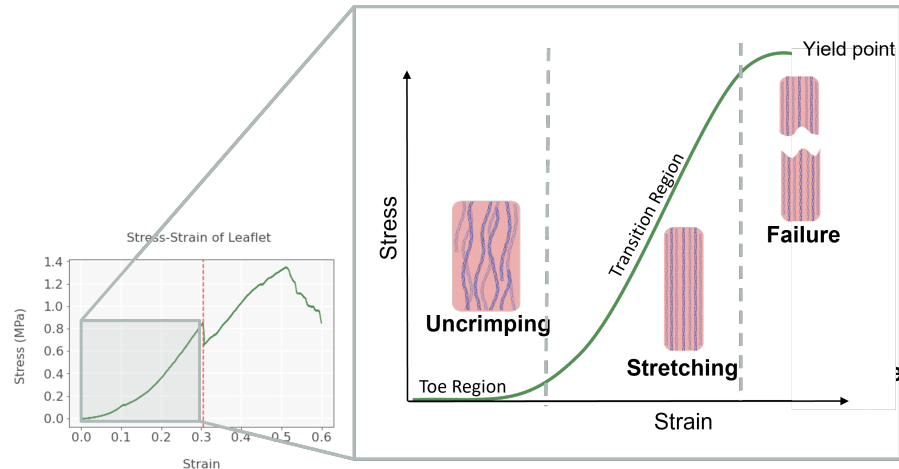


Figure 4.7 : Current capabilities allow for the prediction of the toe and linear transition regions as illustrated above. The high variation in the yield point from sample to sample makes it more difficult to predict the stress-strain curve closer and past the yield point. However, to physiological strain, the predictions are good as seen from Table 4.2.

Table 4.3 : In the third experiment, there was a higher accuracy when excluding the cubic term, since the cubic term has high variations, especially at higher strain.

Architecture	Excluding Cubic MAE	Including Cubic MAE
Alexnet	3.99	65.60
VGG11	3.66	59.05
Resnet18	3.81	66.75

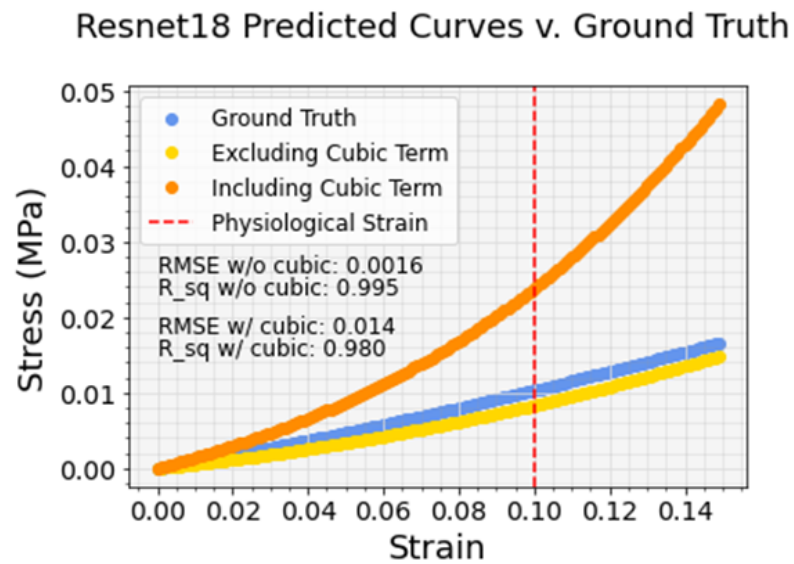


Figure 4.8 : Two predicted curves, one from excluding the cubic term and the other from including all three terms of the polynomial, as seen from the third experiment. Exclusion of the cubic term gives a lower RMSE and R^2 . The physiological strain is given by the vertical red, dashed line.

Chapter 5

Discussion

5.1 Preliminary Findings

The efforts of this work have contributed towards the overall goal, accessible prediction of the mechanical properties of HVLs, by curating a dataset that consisted of macroscopic imaging of the HVLs and their respective stress-strain response along the circumferential axis. The current dataset consists of 51 samples and will be made available through open-access platforms. This is critical because it will allow for continued addition of appropriate data to the current dataset for future and continuing works. Towards that end, early in the investigations the methods that were ultimately adopted were refined for the preparation of the leaflets and the data collection.

Preparation of the leaflets can be difficult because they are small pieces of tissue. Masking tape on the surface of a cover slip was used because this allowed for placement of a sample while constraining it in place without damaging the tissue. However, these cover slip and masking tape platforms were also used in images of the leaflets taken through the stereomicroscope. Initially, the resulting images of the sectioned leaflets with the wooden tabs were used as-is, but the performance of the CNN was not good. The images were then cropped down to minimize the background displayed as well as the wooden tabs, which increased the performance of the CNN. Therefore, segmentation of the leaflet images would be of high value to consider in future works – an additional step to the presented framework would consist of segmentation of the

leaflet area in the image prior to the application of data augmentation.

Additionally, whole-leaflet imaging was considered early on in this work, using a Canon EOS Rebel T1i with an 18 Mp CMOS sensor and a Canon EF 50mm f/2.5 Compact Macro Lens, mounted on a tripod and set to macrophotographic settings. However, this setup was not as consistent as the Leica stereomicroscope setup that was ultimately used. Despite the Canon camera being more accessible than the stereomicroscope, the curation of a consistent dataset was more important at the time.

Additionally, the stress-strain data that was collected from each sample was not used as the ground truth because of its continuous nature. A selection of discrete points along the stress-strain curve was not sufficient for correct reconstruction. However, a data-driven approach could be used in the current framework because variations in the stress-strain curve are manageable up to the physiological strain. The approach presented here is quick and validates the use of the current framework, however, physics-informed models [52] will be explored in the near future now that the necessary data has been collected.

5.2 Curve Reconstruction

Although a simple polynomial was used for reconstruction of the stress-strain response, low RMSE and R^2 values are observed when using the linear and quadratic coefficients, and when predicting up to the physiological strain. This results from the low variations in the stress-strain curve from sample to sample up to the physiological strain. Once the strain reaches the yield and extends past it, the variations become too high for the polynomial to correctly describe the toe and linear transition regions, where the tissue suffers micro and macro level tearing of the collagen fibers.

There are various strategies that can be considered and explored with the current dataset. Rather than reconstructing the stress-strain response to fixed strains, the yield point per sample could be predicted first. Then, the stress-strain curve could be constructed up to this point. Another strategy could be to use a physics-informed model as mentioned in the previous section.

5.3 Prediction Results

In the first experiment, it was observed that more samples led to better results. However, it was also seen that the results from having used very few samples were promising due to the use of data augmentation. In the work presented here, a smaller image patch was used for the smaller set of samples, 8, than was used in the final set of samples, 51, because of time constraints and limited computational power. Due to the results that were obtained with the smaller sample set and the smaller image patch, it is suspected that for the larger sample set a smaller image patch would further improve performance, but will be investigated in future works. This could provide insight as to what minimum number of samples and what size image patch are needed for achieving good predictions, since ground truth collection is expensive.

In the second experiment, strain caps past the yield produced poor results, while strain caps to the physiological strain or less resulted in good predicted stress-strain curve reconstructions. This was expected since the variations prior to the yield point are low. However, as mentioned previously, it would be of interest to determine the yield points for the samples to reconstruct the stress-strain response per sample as accurately as possible. The challenge would then be the strategy adopted to determine the yield point, since this isn't easily defined in literature. Experimentally, one can observe dips in the stress response when there are microscopic tears in the

collagen fibers, yet the stress response will then increase linearly as more fibers become aligned and stretched until macroscopic tearing is observed. One approach could be to instead consider various points as potential yield points and then treat the problem as a piecewise function. The determination of the yield points could then be more strict, since the entire stress-strain curve can be considered without the fear of one part of the curve heavily influencing another.

Finally, in the third experiment, it was observed that the predicted curves were more accurate when the cubic term was excluded. Again, this was due to the high variations in the cubic term. Initially, the cubic term was considered in an attempt to capture the yield point. However, from the preliminary findings, it was evident that the variations in the stress-strain response from sample to sample were too great for this term to effectively predict the point. There are many other approaches that can be considered now that it has been shown that the stress-strain response can be predicted well up to the physiological strain from only a polynomial.

5.4 Conclusion

Through the experiments performed, it is evident that CNNs can be used to predict the circumferential stress-strain response of HVLs from a single image acquired from the sample. This is because macroscopic imaging of the leaflet captures sufficient morphological details related to the mechanics of the leaflet.

Using a larger training set, a strain cap to physiological strain or less, and limiting predictions to the linear and quadratic terms, the MAE of the predicted coefficients can be minimized. These results motivate further investigation of the architecture used to form predictions as well as preparation of the data, since this work used a polynomial fit to reconstruct the stress-strain curves but could explore other models.

As previously noted, this dataset will be made publicly available, making it the first openly available dataset of macroscopic images and raw mechanical data for porcine HVLs. This will allow for other researchers to not only explore different architectures and schemes for improving the prediction results, but also address a wider set of issues. It will allow for researchers to be aware of what data can be collected to enhance the study and design of devices that relate to pathological topics. Additionally, other researchers can contribute to the collected dataset by adding their own data, increasing the publicly available dataset just as it is done for ImageNet. It is further exciting that the additional data does not need to consist of the same HVL samples that were used in this work, but rather can include diseased valves. Furthermore, researchers can compare how architectures trained on the data collected in this work performs on other collagenous tissue, as open datasets can serve as benchmarking tools. Thus, there is excitement for continued efforts to improve and accelerate the device design process by effectively integrating data that is already or can be easily collected during current protocols.

Bibliography

- [1] F. J. Schoen, “Mechanisms of function and disease of natural and replacement heart valves,” *Annual Review of Pathology: Mechanisms of Disease*, vol. 7, pp. 161–183, 2012.
- [2] A. Hasan, K. Ragaert, W. Swieszkowski, Š. Selimović, A. Paul, G. Camci-Unal, M. R. Mofrad, and A. Khademhosseini, “Biomechanical properties of native and tissue engineered heart valve constructs,” *Journal of biomechanics*, vol. 47, no. 9, pp. 1949–1963, 2014.
- [3] Y. Chen, H. Jiang, C. Li, X. Jia, and P. Ghamisi, “Deep feature extraction and classification of hyperspectral images based on convolutional neural networks,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 10, pp. 6232–6251, 2016.
- [4] A. Dhillon and G. K. Verma, “Convolutional neural network: a review of models, methodologies and applications to object detection,” *Progress in Artificial Intelligence*, vol. 9, no. 2, pp. 85–112, 2020.
- [5] W. Rawat and Z. Wang, “Deep convolutional neural networks for image classification: A comprehensive review,” *Neural computation*, vol. 29, no. 9, pp. 2352–2449, 2017.
- [6] Z. Li, F. Liu, W. Yang, S. Peng, and J. Zhou, “A survey of convolutional neural networks: analysis, applications, and prospects,” *IEEE transactions on neural*

networks and learning systems, 2021.

- [7] A. M. Katz, *Physiology of the Heart*. Lippincott Williams & Wilkins, 2010.
- [8] J. T. Butcher, C. A. Simmons, and J. N. Warnock, “Mechanobiology of the aortic heart valve,” 2008.
- [9] F. J. Schoen and R. J. Levy, “Tissue heart valves: current challenges and future research perspectives,” *Journal of Biomedical Materials Research: An Official Journal of The Society for Biomaterials, The Japanese Society for Biomaterials, and The Australian Society for Biomaterials and the Korean Society for Biomaterials*, vol. 47, no. 4, pp. 439–465, 1999.
- [10] J. M. Lee, D. W. Courtman, and D. R. Boughner, “The glutaraldehyde-stabilized porcine aortic valve xenograft. i. tensile viscoelastic properties of the fresh leaflet material,” *Journal of biomedical materials research*, vol. 18, no. 1, pp. 61–77, 1984.
- [11] A. Zhong and C. A. Simmons, “Heart valve mechanobiology in development and disease,” in *Molecular and Cellular Mechanobiology*, pp. 255–276, Springer, 2016.
- [12] C. M. Otto, J. Kuusisto, D. D. Reichenbach, A. M. Gown, and K. D. O’Brien, “Characterization of the early lesion of ‘degenerative’ valvular aortic stenosis. histological and immunohistochemical studies.,” *Circulation*, vol. 90, no. 2, pp. 844–853, 1994.
- [13] M. Misfeld and H.-H. Sievers, “Heart valve macro-and microstructure,” *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 362, no. 1484, pp. 1421–1436, 2007.

- [14] C. S. Peskin and D. M. McQUEEN, “Mechanical equilibrium determines the fractal fiber architecture of aortic heart valve leaflets,” *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 266, no. 1, pp. H319–H328, 1994.
- [15] M. Adamczyk, T. Lee, and I. Vesely, “Biaxial strain properties of elastase-digested porcine aortic valves.,” *The Journal of Heart Valve Disease*, vol. 9, no. 3, pp. 445–453, 2000.
- [16] M. Scott and I. Vesely, “Aortic valve cusp microstructure: the role of elastin,” *The Annals of thoracic surgery*, vol. 60, pp. S391–S394, 1995.
- [17] I. Vesely, “The role of elastin in aortic valve mechanics,” *Journal of biomechanics*, vol. 31, no. 2, pp. 115–123, 1997.
- [18] E. Aikawa and F. J. Schoen, “Calcific and degenerative heart valve disease,” in *Cellular and molecular pathobiology of cardiovascular disease*, pp. 161–180, Elsevier, 2014.
- [19] K. L. Billiar and M. S. Sacks, “Biaxial mechanical properties of the native and glutaraldehyde-treated aortic valve cusp: part ii—a structural constitutive model,” *Journal of biomechanical engineering*, vol. 122, no. 4, pp. 327–335, 2000.
- [20] I. Vesely and R. Noseworthy, “Micromechanics of the fibrosa and the ventricularis in aortic valve leaflets,” *Journal of biomechanics*, vol. 25, no. 1, pp. 101–113, 1992.
- [21] W. Vongpatanasin, L. D. Hillis, and R. A. Lange, “Prosthetic heart valves,” *New England Journal of Medicine*, vol. 335, no. 6, pp. 407–416, 1996.

- [22] R. A. Manji, B. Ekser, A. H. Menkis, and D. K. Cooper, “Bioprosthetic heart valves of the future,” *Xenotransplantation*, vol. 21, no. 1, pp. 1–10, 2014.
- [23] J. M. Lee, D. R. Boughner, and D. W. Courtman, “The glutaraldehyde-stabilized porcine aortic valve xenograft. ii. effect of fixation with or without pressure on the tensile viscoelastic properties of the leaflet material,” *Journal of biomedical materials research*, vol. 18, no. 1, pp. 79–98, 1984.
- [24] S. Westaby, R. B. Karp, E. H. Blackstone, and S. P. Bishop, “Adult human valve dimensions and their surgical significance,” *The American journal of cardiology*, vol. 53, no. 4, pp. 552–556, 1984.
- [25] J. O. V. Delgadillo, S. Delorme, F. Thibault, R. DiRaddo, S. G. Hatzikiriakos, *et al.*, “Large deformation characterization of porcine thoracic aortas: inverse modeling fitting of uniaxial and biaxial tests,” *Journal of Biomedical Science and Engineering*, vol. 8, no. 10, p. 717, 2015.
- [26] A. Avanzini and D. Battini, “Integrated experimental and numerical comparison of different approaches for planar biaxial testing of a hyperelastic material,” *Advances in Materials Science and Engineering*, vol. 2016, 2016.
- [27] A. Pokutta-Paskaleva, F. Sulejmani, M. DelRocini, and W. Sun, “Comparative mechanical, morphological, and microstructural characterization of porcine mitral and tricuspid leaflets and chordae tendineae,” *Acta Biomaterialia*, vol. 85, pp. 241–252, 2019.
- [28] C. J. Ross, D. W. Laurence, M.-C. Hsu, R. Baumwart, Y. D. Zhao, A. Mir, H. M. Burkhart, G. A. Holzapfel, Y. Wu, and C.-H. Lee, “Mechanics of porcine heart valves’ strut chordae tendineae investigated as a leaflet–chordae–papillary

- muscle entity,” *Annals of biomedical engineering*, vol. 48, no. 5, pp. 1463–1474, 2020.
- [29] S. M. Mehta, T. Jin, I. Stanciulescu, and K. J. Grande-Allen, “Engineering biologically extensible hydrogels using photolithographic printing,” *Acta Biomaterialia*, vol. 75, pp. 52–62, 2018.
- [30] D. S. Puperi, A. Kishan, Z. E. Punske, Y. Wu, E. Cosgriff-Hernandez, J. L. West, and K. J. Grande-Allen, “Electrospun polyurethane and hydrogel composite scaffolds as biomechanical mimics for aortic valve tissue engineering,” *ACS Biomaterials Science & Engineering*, vol. 2, no. 9, pp. 1546–1558, 2016.
- [31] P. Mishra, S. Biswas, and S. Dash, “Deep learning based biomedical named entity recognition systems,” in *Deep Learning Techniques for Biomedical and Health Informatics*, pp. 23–40, Springer, 2020.
- [32] P. Baldi, “Deep learning in biomedical data science,” *Annual review of biomedical data science*, vol. 1, pp. 181–205, 2018.
- [33] M. A. Nielsen, *Neural networks and deep learning*, vol. 25. Determination press San Francisco, CA, USA, 2015.
- [34] Y. Amamoto, “Data-driven approaches for structure-property relationships in polymer science for prediction and understanding,” *Polymer Journal*, pp. 1–11, 2022.
- [35] Y. Sun, I. Hanhan, M. D. Sangid, and G. Lin, “Predicting mechanical properties from microstructure images in fiber-reinforced polymers using convolutional neural networks,” *arXiv preprint arXiv:2010.03675*, 2020.

- [36] W. Wang, Y. Yang, X. Wang, W. Wang, and J. Li, “Development of convolutional neural network and its application in image classification: a survey,” *Optical Engineering*, vol. 58, no. 4, p. 040901, 2019.
- [37] P. Z. Hanakata, E. D. Cubuk, D. K. Campbell, and H. S. Park, “Accelerated search and design of stretchable graphene kirigami using machine learning,” *Physical review letters*, vol. 121, no. 25, p. 255304, 2018.
- [38] Z. Yang, Y. C. Yabansu, R. Al-Bahrani, W.-k. Liao, A. N. Choudhary, S. R. Kallidindi, and A. Agrawal, “Deep learning approaches for mining structure-property linkages in high contrast composites from simulation datasets,” *Computational Materials Science*, vol. 151, pp. 278–287, 2018.
- [39] L. Liang, M. Liu, and W. Sun, “A deep learning approach to estimate chemically-treated collagenous tissue nonlinear anisotropic stress-strain responses from microscopy images,” *Acta biomaterialia*, vol. 63, pp. 227–235, 2017.
- [40] R. Yamashita, M. Nishio, R. K. G. Do, and K. Togashi, “Convolutional neural networks: an overview and application in radiology,” *Insights into imaging*, vol. 9, no. 4, pp. 611–629, 2018.
- [41] A. Krizhevsky, “One weird trick for parallelizing convolutional neural networks,” *arXiv preprint arXiv:1404.5997*, 2014.
- [42] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” *Communications of the ACM*, vol. 60, no. 6, pp. 84–90, 2017.
- [43] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.

- [44] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- [45] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255, Ieee, 2009.
- [46] X. Han, Y. Zhong, L. Cao, and L. Zhang, “Pre-trained alexnet architecture with pyramid pooling and supervision for high spatial resolution remote sensing image scene classification,” *Remote Sensing*, vol. 9, no. 8, p. 848, 2017.
- [47] D.-X. Xue, R. Zhang, H. Feng, and Y.-L. Wang, “Cnn-svm for microvascular morphological type recognition with data augmentation,” *Journal of medical and biological engineering*, vol. 36, no. 6, pp. 755–764, 2016.
- [48] L. Chato and S. Latifi, “Machine learning and deep learning techniques to predict overall survival of brain tumor patients using mri images,” in *2017 IEEE 17th international conference on bioinformatics and bioengineering (BIBE)*, pp. 9–14, IEEE, 2017.
- [49] R. Cang, H. Li, H. Yao, Y. Jiao, and Y. Ren, “Improving direct physical properties prediction of heterogeneous materials from imaging data via convolutional neural network and a morphology-aware generative model,” *Computational Materials Science*, vol. 150, pp. 212–221, 2018.
- [50] A. Sauren, M. Van Hout, A. Van Steenhoven, F. Veldpaus, and J. Janssen, “The mechanical properties of porcine aortic valve tissues,” *Journal of biomechanics*, vol. 16, no. 5, pp. 327–337, 1983.

- [51] S. Arjunon, S. Rathan, H. Jo, and A. P. Yoganathan, “Aortic valve: mechanical environment and mechanobiology,” *Annals of biomedical engineering*, vol. 41, no. 7, pp. 1331–1346, 2013.
- [52] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, “Physics-informed machine learning,” *Nature Reviews Physics*, vol. 3, no. 6, pp. 422–440, 2021.