Characterization of a piezoelectric-actuated, customizable-stiffness device for vibration suppression

by

Nathan Allen Featherstone

B.S., Kansas State University, 2021

A THESIS

submitted in partial fulfillment of the requirements for the degree

MASTER OF SCIENCE

Alan Levin Department of Mechanical and Nuclear Engineering
Carl R. Ice College of Engineering

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2023

Approved by:

Major Professor
Dr. Raj Kumar Pal

Copyright

© Nathan Featherstone 2023.

Abstract

Vibration causes harmful effects for machines including total resonance failure, fatigue failure, user discomfort, and decreased performance. Because of these harmful effects creating effective vibration control measures is essential, especially for light weight structures used in aerospace applications. This research seeks to develop a customizable-stiffness device that could be usable in aerospace, vibration-control applications. The device is designed around a slender post-buckled beam with intermediate supports to influence its stiffness. Piezoelectric materials are used to actuate between two different support scenarios. There are three phases of this research a nonlinear, dynamic simulation; experimental testing; and a neural network analysis for design optimization.

In the first phase, the primary concern was to obtain results for the post-buckling stiffness of the device. The first result was a nonlinear, dynamic simulation that provided consistent post-buckling results. Initially, a static simulation was expected to be sufficient, but it would not provide consistent results. From the dynamic simulation, results for the compressive force and displacement could be obtained to create a force-displacement diagram. By taking a numeric derivative, the stiffness results could be obtained. Secondly, these results suggested it was possible to affect the devices stiffness by changing the position of the supports.

In the second phase, the primary concerns were proving physical efficacy of actuating the beam between two support scenarios and corroborating evidence for simulation results. Compression testing was completed in a custom buckling rig to obtain experimental force-displacement and stiffness-displacement curves. For the first two experiments, Macro Fiber Composites (MFCs) were bonded to the steel test specimen. The first experiment proved that an MFC was able to influence the direction of buckling. The second experiment demonstrated

significant differences in the displacement-stiffness plot based on support positioning. The third experiment was useful in supporting the simulated results for loads at buckling and contacts. Additionally, the shapes of the simulated and real beam were compared.

In the final phase, a neural network is created to predict results for the different support locations. A neural network was created to predict the stiffness of the device given a displacement and the location of the two supports. Using the stiffness predictions from the network, good force displacement diagrams. A root mean square error (RMSE) of 0.54 N was obtained from the best network. The stiffness-displacement diagrams did not align as well with a RMSE of 230 N/mm.

Table of Contents

List of Figures

List of Tables

Acknowledgements

I would like to thank Dr. Jared Hobeck for helping me begin and progress through this academic journey and Dr. Raj Kumar Pal for helping me to finish it. I would also like to thank the numerous professors at Kansas State University who helped me to learn material I sought to learn and some that I did not want to learn. Additionally, I would like to thank the friends and family who pushed me and encouraged me throughout the process. Finally, I would like to thank my Father in Heaven, the giver of all good gifts, and Jesus Christ His son in whom, I put my hope.

Chapter 1 - Introduction and Literature Review

1.1 Vibration and Vibration Suppression

The study of vibrations and their suppression is important because vibration can cause catastrophic failure, causes fatigue failure, decreases the ability to perform its intended purpose, and it reduces the comfort of the user. A failure to understand the bridges dynamics led to a failed bridge within four months. Besides catastrophic failure, fatigue is another source of costly failure. Dynamic loading from undesired vibrations reduces the life of machine components. Vibration can wreck the effectiveness of a machine; a little vibration ruins the ability of a camera to take a good image. Finally, the comfort of the machine's user is often diminished because of vibration.

Because of vibration's significant effect, it is important to identify its source. Vibration is oscillating motion about a nominal point often an equilibrium point [1]. Vibration requires three components a mass, a spring, and a constraint. Mass is an inherent property of matter, and all real materials have some elasticity. The only remaining requirement is a constraint. While these three factors are required, real systems also have the following components. No real material is perfectly elastic meaning it has some damping effect. Finally, a disturbance from equilibrium is required to cause motion.

Because of vibrations negative effects, it should be minimized through the process of vibration suppression. This process has three main objectives: vibration elimination, vibration isolation, and vibration absorption. The first goal is obviously the best, but it is not always possible. The main source of vibration of a chain saw is from its engine removing the source (the engine) removes the source of power. The goal of vibration isolation is to isolate (separate) one part from vibrations in the rest. The shocks of a car isolate the passenger from vibrations caused by bumps in a road. Finally, in vibration absorption the goal is to target specific frequencies of vibration and minimize their effects [1]. Noise cancellation headphones target specific frequencies protecting the wearer from harmful noises. The goal of vibration can be achieved by stopping the source of vibration, by separating the object from the source of vibration, or by targeting and eliminating the most impactful frequencies of vibration.

1.2 Vibration control methods.

To accomplish the vibration objectives, there are passive, active, and semi-active methods of vibration suppression. Passive methods are best for either isolation or absorption for specific and consistent frequencies [2]. Passive methods are not a good choice for controlling a large range of frequencies. For better broadband control, an active approach can be used [3]. Active approaches consist of a controller, sensors, and actuation to create a signal that destructively interferes with the undesirable vibration. The main two difficulties for active vibration control are instabilities and energy cost [2], [3]. Semi-active control is meant to have a broader range of frequency control while minimizing the power requirements and eliminating instabilities [2]. A semi-active vibration control system includes the same components as an active system, but the actuation system does not add energy to the system. Instead, the actuator may change the mass, stiffness, or damping of the system [4]. By changing these characteristics, the system's natural frequency can change to minimize the systems response to the excitation frequency.

1.2.1 Examples of Passive Vibration Control

The essential example of passive vibration control is the tuned mass damper (TMD). A TMD is small vibrating system (containing a mass, spring, and damper) added onto a larger vibrating system to absorb specific vibration frequencies. TMDs have been used to counteract vibrations in buildings since 1971, but they are used in a variety of applications. Tuned Liquid Column

Dampers, or more simply Tuned Liquid Dampers, use a similar concept but use water as the mass and its random flow as the damping mechanism. Beneficially, the water level can be adjusted as the natural frequencies of the building change with construction [5].

1.2.2 Examples of Semi-Active Vibration Control

Magnetorheological elastomers (MREs) are commonly used in semi-active devices. Magnetorheological materials change their stiffness based on the presence of magnetic field. This magnetic field can easily be produced by an electromagnet controlled by a controller. Using this control, a magnetorheological device can have a continuous range of custom stiffness. In Wang et. al., magnetorheological devices in a semi-active TMD control the vibrations of a bridge tower model during construction. The customizable stiffness allows the TMD to adapt its frequency based on the changing frequency of the bridge tower while minimizing actuation energy [2]. In a paper by Jabarri and Bobrow, a MRE is used in a resetting mode to maintain maximum stiffness and get greater energy dissipation under impact conditions [6].

1.3 Vibration Control Approaches Using Piezoelectric Materials

Piezoelectric materials are useful in both active and passive vibration suppression methods. In an active approach, the piezoelectric material can be used as an actuator (according to inverse piezoelectric effect). In a passive approach, the piezoelectric material converts mechanical motion into an electric signal (direct piezoelectric effect) which can be dissipated in a shunt circuit. The main advantage of the active approach is the wide range of frequencies which the active system can control, but the main drawbacks are the required power for input signals to the piezoelectric actuators and instabilities from the control model not being able to account for all factors. The main advantage of the passive approach is inherent stability, but its drawback is that it can only control specific frequencies of vibration [3].

1.3.1 Active Methods

Piezoelectric materials develop an electric potential when they experience strain. These materials will inversely experience strain when a voltage is applied to the material. Many useful piezoelectric materials are ceramics. By using the second phenomena, a voltage signal can be used to cause physical change in a system. This is useful for actuation. The advantages of piezoelectric actuation include relatively high force, fast response time, and easily implemented actuation.

1.3.2 Passive Methods

A main method of passive piezoelectric is shunt damping. Piezoelectric material converts mechanical energy of the vibrations it experiences into electrical energy which is dissipated in a shunt circuit. The shunt circuit is designed to maximize the energy dissipation. One architecture of the shunt circuit is the RL-series shunt circuit. The circuit shown includes the piezoelectric patch which acts like a capacitor, an inductor, and a resistor. The damping is maximized when the reactance from the capacitor and inductor is zero. To control multiple frequencies of vibration, it is proposed that multiple RL circuits could be arranged in parallel. The challenge to be solved is how to design electrical components for different frequencies when each parallel branch affects the other [3]. Two ways to simplify the design of the branches is the current blocking shunt or especially the current flowing shunt. For a system designed to control frequencies ω_1, ω_2, and ω_3, the branch of a current blocking shunt for ω_1 will have additional components, so it will have a large reactance at ω_2 and ω_3 (infinite) so no current will flow at those frequencies. This means that each branch must add circuitry for every additional frequency that should be controlled. The branch, for ω_1 for a current flowing shunt, is designed to have low reactance (a short) when the mechanical system has frequency ω_1. This allows for the superposition of RL circuits separately designed for their individual frequency [7]. While both these methods make the design of the individual RL branches easier, it appears that near resonance conditions are required for control. Piezoelectric shunt damping is an effective way to suppress vibration in a system, but the electrical components should be designed to maximize energy dissipation.

1.4 Metastructures

Metastructures are macroscale structures that have properties that depend on the geometry of periodically repeated simpler cells. The design of the unit cells geometry is the most significant factor in the properties of the metastructure [8]. Metastructures are designed to have unique properties that do not occur in structures built from bulk materials including significantly higher specific strength (ratio of strength to weight), thermal insulation, and significant energy absorption (at least for a band gap of frequencies) [9], [10]. Because metastructures have high specific strength, they are great for building aerospace structures. However, structures with less mass are more susceptible to low frequency excitations. Metastructures can also be used in vibration suppression because of their energy absorption properties. Since both ideas are relevant to this research, the work of selected papers will be reviewed in the following subsections.

1.4.1 Metastructures as a Structural Material

Architected materials have been adopted for a variety of applications in construction of weight optimized structures. The materials' specific strength and specific stiffness are of primary importance, but so is their tuneability for other characteristics. Bici *et. al.* illustrate their usefulness in a design for a leading edge of an aircraft [11]. "The [leading edge] must with stand aerodynamic forces and possible bird-strikes, [while also integrating] anti-ice system functions," [11]. In the design, the deicing components are combined with a lattice sandwich panel to simplify the assembly of the wing tip. In summary, an additively manufactured metastructure simplifies assembly, maintains strength requirements, while also decreasing the weight [11]. Due to their light weight and low thermal expansion, composite honeycomb structures are used in rockets and space antenna [12]. These structures are expensive to make, and their complex geometry compounds the problem, so Saito *et. al.* proposed forming them using kirigami [12]. With the potential for weight optimization and maintaining high strength, metastructures are desirable for aerospace structures.

1.4.2 Metastructures as a Vibration Suppression Material

While a honeycomb structure may be more common for structural materials, it is not complicated enough for designing a metastructure's vibration absorption bandwidth [8], [9]. For example, Xu *et. al.* designed a 2D metastructure with a unit cell that included a conventional hexagon and an auxetic hexagon [9]. With this geometry, the metastructure combines the stiffness advantage of the conventional hexagon with the ability to design for unique mechanical traits from the auxetic hexagon [9]. In a paper by Fan *et. al.*, a 2D quasi-zero stiffness metastructure is produced for vibration isolation applications. Besides the good performance as a vibration isolator, a main benefit of this metastructure is the scalability as compared to conventional quasi-zero stiffness isolators [13]. While the previous metastructures consisted solely of one material, Matlack *et. al.* creates a metastructure with embedded masses with in a polymer matrix [10]. By adding the resonant masses, the metastructure can absorb lower frequencies with a larger bandwidth. By removing axial members from the matrix, a further lowering of the band gap is possible [10]. Gupta *et. al.* further complicate their metastructure, by using 3D hour glass shaped unit cells [8]. With the added complexity, this metastructure has a customizable stiffness and Poisson's ratio. With more design parameters, researchers can design metastructures with greater customizability in terms of its mechanical properties such as its Poisson’s ratio and vibrational absorption bandwidth.

1.5 Outline for Remaining Chapters

The remaining portion of this thesis will be dedicated to the characterization of a non-linear variable stiffness device that could be useful for semi-active vibration control, especially for applications that require weight minimization.

Chapter 2 will begin the characterization process by introducing the objectives for the device and the design concepts of the device. These objectives include building actively controlled vibration control into the structure of the device, a semi-active metastructure. The main design concept is using buckling to amplify the force and displacement achievable with a piezoelectric actuator.

Chapter 3 discusses the Finite Element Analysis (FEA) characterization of the device with the goal of predicting the post buckling behavior of a slender beam with supports. Additionally, data is collected for the understanding the key design parameter (the positioning of these supports).

Chapter 4 develops the experimental testing of a device prototype to provide corroboration to the simulation methods. Data was collected using a load cell, deflectometer, and digital camera. The goal is to validate numerical methods for their future use.

Chapter 5 develops a neural network approach for designing the support locations. To be able to customize stiffness, there needs to be a reliable, quick way for testing a support location combination.

Chapter 6 will discuss conclusions from this research along with recommendations for future work. The first recommendation to continue this work is to develop a metastructure that can effectively use this device. While the problem is posed with slender beams in mind, many metastructures are made up of plates. It appears that a truss structure is the clearest application, but does that mean extra truss elements need to be carried, which defeats the purpose of weight minimization?

Chapter 2 - Design Concepts Behind Variable Stiffness Device

2.1 Importance of Force

The need for larger displacements is often cited as a problem for piezoelectric actuated devices, but the need for larger force is also an important factor. With limited force, the piezo electric actuated system cannot change system characteristics for the suppression of vibration.

To demonstrate the importance of force, a two-beam cantilever example is introduced. In Figure 2.1, two cantilever beams are shown. In both cases, a smaller beam is mechanically fastened to the large beam. The smaller beam is compressed, so it experiences a preload greater than the buckling load. In Case A, a MFC is applied directly to the large beam. In Case B, the MFC is bonded to the smaller beam. When the voltage is applied to the first case, the MFC will cause a small strain to develop on the surface of the large beam, and it may experience slight bending. For the second case, the same voltage applied will cause a large displacement change in the small beam. Since the beam has already buckled, this is either accomplished by increasing the bending in one direction or by changing the direction. Whichever case, the large beam will experience a significant change in bending. The two-beam cantilever system shows the importance of force in changing the characteristics of the system, and it demonstrates a method to magnify the force of the piezoelectric actuator.

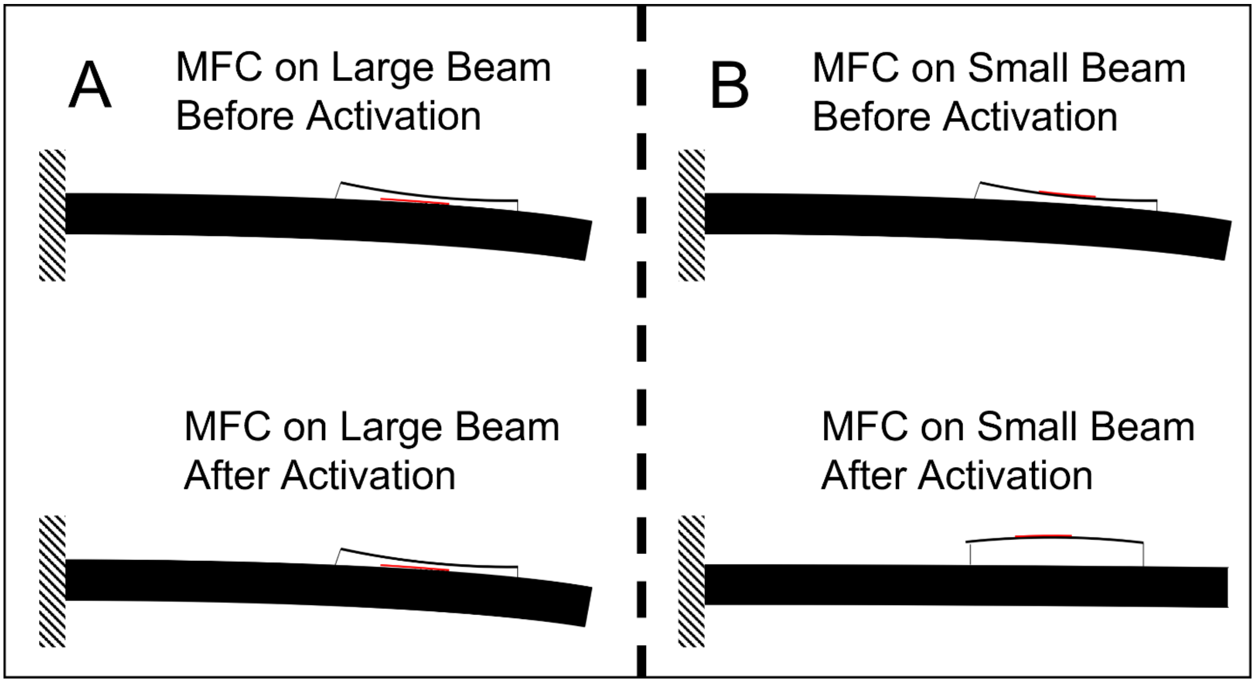

Figure 2.1 Two systems with two beams each that compares methods of influencing the larger beam's stiffness.

2.2 Redefining the System

In the previous section a two-beam cantilever system was introduced, but the systems of interest are aerospace metastructures, specifically deployable truss structures. These structures are optimized to reduce their weight, to minimize the cost of sending them to space, so they are vulnerable to low frequency induced vibrations [14]. However, these induced vibrations can make the space vehicle useless for its intended purpose. With the importance of weight minimization, space structures are more sensitive to vibration, so vibration control is necessary.

With a redefined structure, the way the buckled small beam is introduced is different as well. Now, the small-buckled beam from the previous example is introduced into the truss. The truss could be made up of some elements which have MFCs (active elements) and some that do not (passive elements). A drawing of a possible metastructure is shown in Figure 2.2. If the loading force has a frequency nearly equal to the natural frequency of the system, the natural frequency of the structure can be raised (away from resonance) by stiffening the active elements. The method for increasing the stiffness of these elements will be discussed in the next section, but the device should be able to affect the structure's stiffness by a change in its own stiffness.

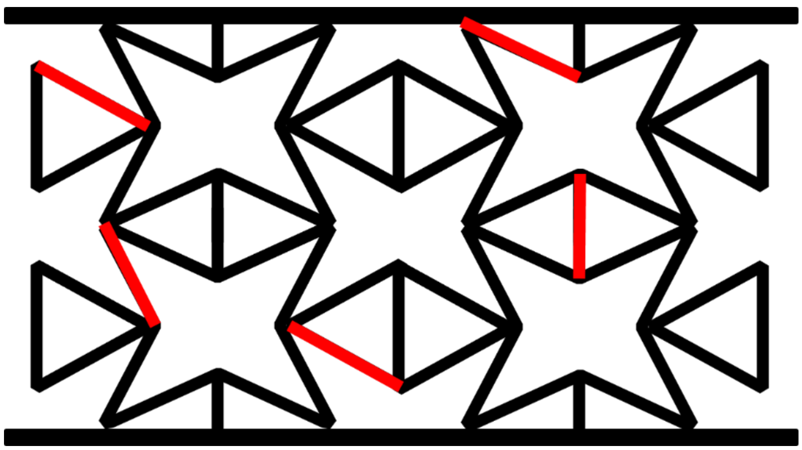

Figure 2.2: A metastructure with a combination of active and passive elements, adapted from [15] . Active elements are shown in red (extended).

2.3 Device Geometry

To change the active element device's stiffness, the beam will be actuated between two different sets of supports. By positioning the supports differently, the device will have different stiffness characteristics. Figure 2.3 contains a diagram of the device with the key parameters labeled.

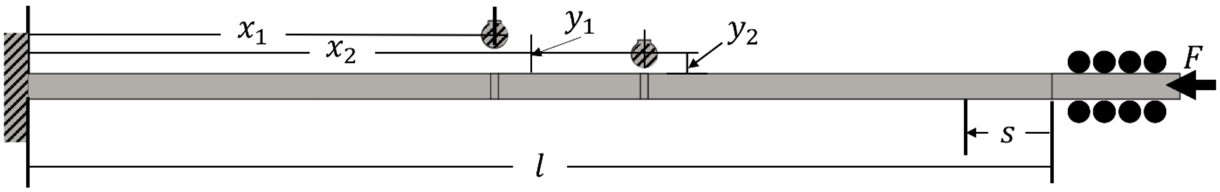

Figure 2.3: Conceptual diagram of the variable stiffness device with important variables labeled.

Of first importance is the constraints, on the left side of the device there is a fixed device, and on the right side there is a sliding fixed constraint. The length of the beam is measured from the fixed constraint to the inside edge of the sliding constraint as shown in the figure. The sliding constraint has a defined but variable position as a function of time. The displacement (s) of the sliding constraint is measured as shown in the figure. The force (F) results from the compressive displacement s of the support and is measured at one of the constraints. The stiffness (k) of the device is calculated by taking the derivative of F with respect to s as seen in Eq. 2.1. This stiffness is varied to help control the vibrations of the structure.

$$k = \frac{dF}{ds} \qquad \text{Eq. 2.1}$$

There are two intermediate supports shown in the figure, they have positions (x_1, y_1) and (x_2, y_2) respectively. The supports positions are fixed. These positions are modified to change the stiffness of the device. The physical geometry is important for understanding future chapters, but it is also important to understand the characteristics of the force displacement and stiffness displacement curves.

2.4 Force Displacement Curve

Understanding the force displacement curves and stiffness displacement curves is important to understanding the work in future chapters. An example force deflection curve is shown in Figure 2.4. The plot has s on the x-axis and F on the y-axis. While compressive stresses (forces) and strains (displacements) are often included a negative sign, a positive sign is given to both since everything is negative. The plotted curve can be divided into several distinct regions. First there is a linear relationship between the Force and Displacement in the red section of the plot. This is the pre-buckling portion of the graph where the stress in the beam is proportional to strain according to Hooke's Law (shown in Eq. 2.2).

$$\sigma = E\varepsilon \qquad \text{Eq. 2.2}$$

After the linear portion, the beam buckles when F_i is equal to the critical buckling load P_{cr}. The critical buckling load can be solved for using the Euler-column buckling formula in Eq. 2.3.

$$P_{cr} = \frac{\pi^2 EI}{l^2} \qquad \text{Eq. 2.3}$$

After the beam has buckled, it experiences a brief portion where the force is essentially constant. In this constant section the stiffness is greatly decreased and is nearly zero. Next, the force begins to start rising again. This change in behavior is caused by a contact with a support. The support constrains the beam from continuing to bend outward which causes the beam to stiffen. After this stiffened portion, the beam will continue deforming until it reaches another constant force portion. When there is another support, the force will once again experience a rise in force followed by an essentially constant force portion. The force displacement plot is important when understanding buckling, especially when designing for maximum compressive load, but for this analysis, the stiffness displacement plot is of more importance.

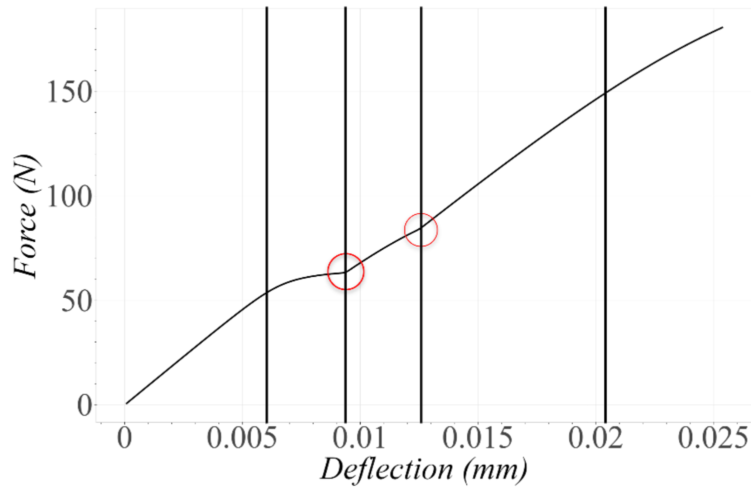

Figure 2.4: A force deflection plot. Contacts are circled. The lines divide the plot into sections.

2.5 Stiffness Displacement Plot

The force displacement behavior is helpful, but since stiffness is critical, the plot of stiffness should also be understood. A stiffness plot is shown in Figure 2.5. Again, s is on the x-axis of the plot, and k is shown on the y-axis of the plot. During the pre-buckling portion, the stiffness will be nearly constant until the beam buckles. The buckling point is identified in the stiffness plot by the large finite jump discontinuity, immediately after the constant portion. After buckling, the stiffness plummets to a low constant value. Once the beam contacts the first support, there is a sudden rise in stiffness. After the jump in stiffness the beam will maintain a constant stiffness. If the beam is allowed to deform unhindered by another support, there will be a jump in stiffness down to a lower constant value. When there is another support contact, there will once again be a discontinuous jump up to a period of constant stiffness, followed by a relaxation to a lower stiffness. It is important to understand the stiffness plot, and how to interpret its response considering the physical system.

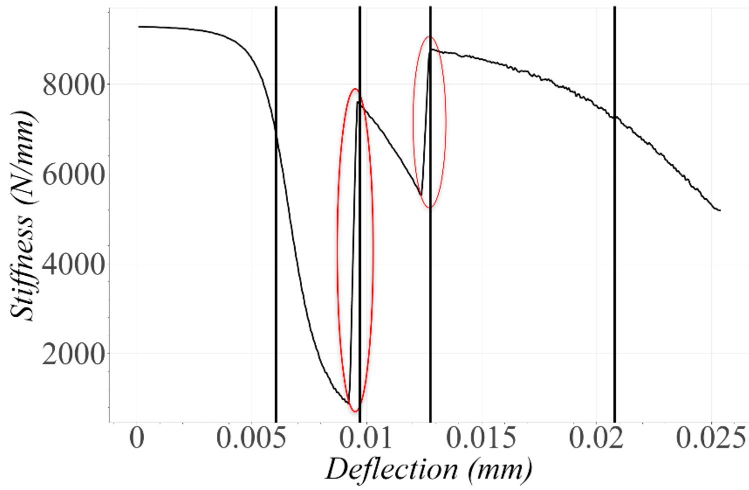

Figure 2.5: A stiffness-displacement plot. The contacts are circled. This is the stiffness-displacement representation of the same behavior as in Figure 2.4.

2.6 Key Features of the Stiffness Deflection Plot

From the stiffness deflection plot, there are several key features that improve the performance of the active element device. First, the device should be within a range where its stiffness retains most of its original stiffness. When the device's stiffness is in the portion of the stiffness plot immediately after buckling, it will not be able to carry a load. If a compressive force is applied, the member must deflect minimally to maintain rigidity in the structure. Second, if the supports are positioned so there are two levels of acceptable stiffness, there will be a greater ability to respond to multiple loading scenarios. A third factor to consider is the displacement, the displacement conditions must be able to be met without actuation. The space structure will not be including heavy equipment to actuate between displacement conditions, so the displacements must be met during normal operation. From the stiffness plot, several important characteristics can be found including how much of original stiffness the member will retain, what stiffnesses the beam will reach for a given support scenario, and the working range for the displacement of the beam.

2.7 Paper Scope

While the previous sections introduced a whole system and application for a variable stiffness device, the goal of this research is to characterize the behavior of a single active element. No multiple beam structures are produced. No beam is loaded dynamically. No active control is implemented. The beam is first modeled using Finite Element Analysis (FEA) to provide data for post-buckled beam behavior. Secondly, laboratory experimentation provides evidence for the efficacy of actuation and corroboration of the FEA results. Finally, a machine learning model is developed for faster iteration through beam designs.

Chapter 3 - Finite Element Analysis Study

3.1 Purpose: Develop Understanding of Post Buckling Behavior

From the previous chapter, the goal of this research is to design a device that can modify a structure's stiffness by controlling the post buckled stiffness of the device. To understand the post buckling behavior of this device, a FEA study was completed.

3.2 Simulation Study Description

To obtain results through simulation, decisions must be made in setting up the simulation. The validity of these decisions are important for the validity of the results, so they should be thoroughly described.

3.2.1 High Level Details: Nonlinear Dynamic Simulation Study

SOLIDWORKS Simulation was used to perform a dynamic nonlinear FEA study to understand the post buckling behavior of a slender beam with intermediate supports. A nonlinear simulation is required because the goal is to understand post buckled behavior. A linear simulation assumes stress and strain are proportional according to Hooke's law, but this is not the case once the beam has buckled. Since the loading condition of interest is quasistatic, a dynamic simulation should be unnecessary. The simulations began as static nonlinear simulation, but there were problems. First, data had to be scaled by the input force (a commonly required step in Euler buckling simulations). However, Figure 3.1 shows the results from the static simulation are still dependent on the value of the input force. Based on further research, a dynamic nonlinear simulation was implemented to circumvent this issue [16]. In the dynamic nonlinear simulation, the compression force was provided by applying a prescribed displacement to the end of the beam. This is beneficial because the results no longer are dependent on the value of the input force (fixing the problem from the static simulation). Additionally, prescribing a displacement is exactly the method of testing in the lab. Using the nonlinear dynamic simulation, consistent results for the post buckling behavior were obtained.

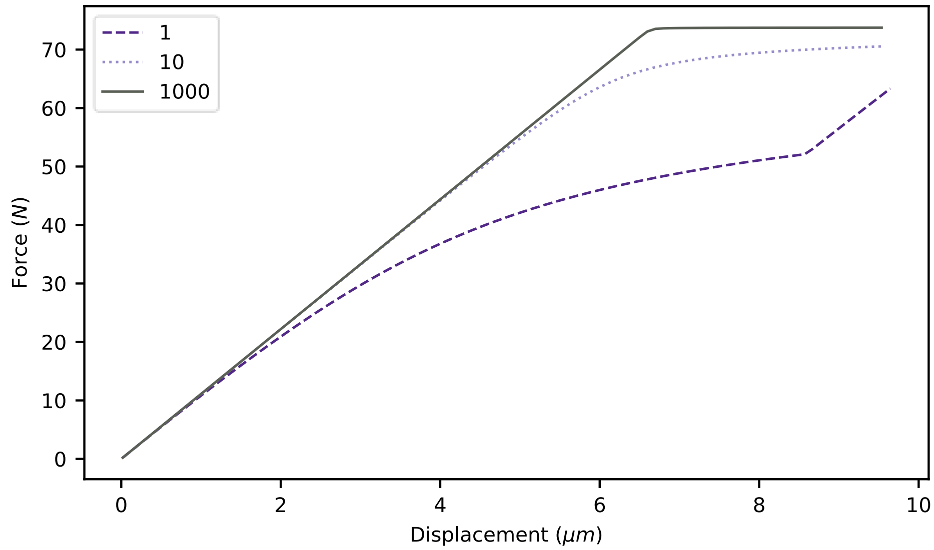

Figure 3.1: Load displacement plot based on SOLIDWORKS static simulation demonstrating the dependence on input load.

The second major high-level decision is the 2D simplification used in this study. Figure 3.2 contains the free body representation of the beam with its loads and reactions. Since the beam, experiences no loading in the z-direction, the stress of the beam will be independent of z and the problem can be simplified to 2D by making a plane stress assumption. In using a 2D simplification, the mesh is simplified to 2D using shell elements. In many cases, shell elements are used to represent the thickness of thin materials that experience out of plane bending. Consequently, shell elements are modeled using high order beam theories, and have rotational degrees of freedom. Since the motion of the beam is in plane, two rotations and one translation are not used. This leaves three degrees of freedom. For a 2D mesh, the elements would usually have two degrees of freedom because bending can be modeled with elements above the neutral axis in tension and elements below the neutral axis in compression. It is not clear how much the additional degree of freedom affects results.

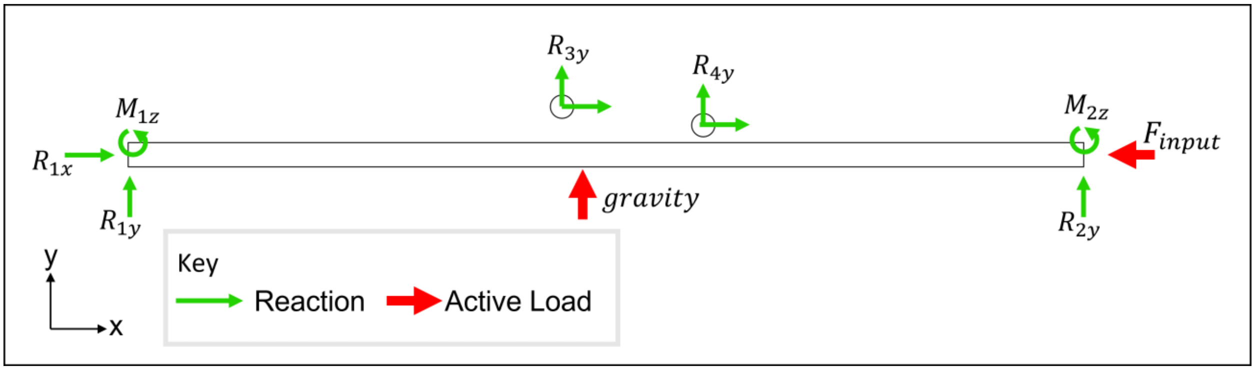

Figure 3.2: Free body diagram of the analyzed beam. No loads in the z-direction mean that the simulation can be simplified to 2D.

3.2.2 Description of Parts

In the simulation study, two parts involved in the simulation: the beam and the supports.

The finite element studies geometry is important to define. The slender beam dimensions are shown in Table 3.1. The total length is equal to the analysis length (l) together with a half inch extension required for the constraints. The rounded portion of the supports have a radius of 0.005 inches which is their critical dimension. Both the supports and the beam were modeled with AISI 1020 steel with material properties shown in Table 3.2.

Table 3.1: Geometric specifications of the beam analyzed with FEA.

Analysis Length (l) [in.]	Total Length [in.]	Width (w) [in.]	Thickness (t) [in.]
5	5.5	0.5	0.02

Table 3.2: Material properties of the AISI 1020 steel used for the FEA study.

Elastic Modulus (psi)	Poisson's Ratio	Mass Density (lb/in^3)	Tensile Strength (psi)	Yield Strength (psi)
29007547.53	0.29	0.285405526	60989.38395	50991.06

3.2.3 Loads and Constraints

The definition of the loads and constraints is critical to the results of the simulation; because the constraints are clearer, they will be described first. Figure 3.3 contains a graphic that shows the loads and constraints. The left edge of the beam is constrained with a fix support constraining rotation and translation. The two supports have a fixed support applied on the entire round edge of the support. The supports are treated as infinitely stiff in this analysis; the goal is to understand the behavior of the beam not the supports. On the right side of the beam, there is a half inch section divided by a split line and has roller supports on the top and bottom edge. The purpose of the extended roller supports is to constrain the moment on that side of the beam as necessary for a sliding fixed support. A normal fixed support could not be used, as that would prevent sliding. The final constraint is a prescribe constraint in the negative x-direction, but it will be addressed in the next paragraph. The constraints are important for creating an accurate loading scenario, so how they are defined is important for the results of the simulation.

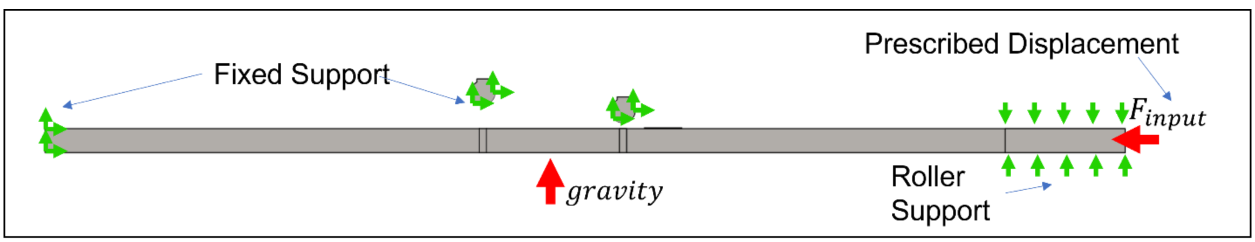

Figure 3.3: Loads and constraints applied to the beam. The portion of the beam where the roller support is applied is an extension to prevent a change in angle at the beginning of the beam.

To load the beam, the beam experiences a prescribed compressive displacement of 0.001 inches. This value was refined from a previously larger value. This displacement should be small since the beam is designed to support loads on the structure. Additionally, it was selected based on previous simulation runs. The goal was to include a lot of simulations where the beam contacts both supports. However, it is not desirable to have one contact followed by a much later contact. The beam will have low stiffness and will not be able to support load between contacts.

The final load the beam experiences is gravitational loading. The purpose of this loading is to provide a disturbance from equilibrium. Without the gravity, there are two equally likely scenarios, buckling in the positive y-direction or buckling in the negative y-direction. This bifurcation will cause the simulation to fail. Adding gravity eliminates the bifurcation, making only one result correct. In these simulations, the gravity constant was defined as 2g. With the definition of the gravity and the prescribed displacement, all loads and constraints are defined; the next section will discuss the properties needed for the dynamic part of the simulation.

3.2.4 Dynamic Properties: Loading Schedule and Damping Coefficients

Since the simulation is dynamic, two additional requirements should be defined the loading curve with respect to time and the Raleigh coefficients. This displacement is applied linearly with respect to time over 200 seconds. The length of time is not of great significance except that it needs to be long to satisfy the quasi-static loading condition. The simulation will by default calculate simulation steps for every half second, giving 400 data points in a solution plot, but with auto stepping there may be additional steps near difficult places. The gravity also needed a loading curve defined, so its magnitude was applied immediately and remained constant for the remainder of the test. There is more flexibility in defining the loading curve because the goal is simply a quasistatic test, but not defining it could cause problems.

Damping must be added to the simulation. The Rayleigh damping coefficients were defined as shown in the following equations.

$$\alpha = 0 \qquad \text{Eq. 3.1}$$

$$\beta = 0.2 \qquad \text{Eq. 3.2}$$

To define these parameters, the gravity was applied, but not the prescribed displacement. The displacement in the y-direction (transverse) is measured. β was increased until the y displacement decayed to zero in two seconds. One could just as easily have used α. With the quasistatic loading condition, the response should be nearly time independent, so the value of damping is not critical. It should attenuate any dynamics that may occur though. Like the loading curve, the damping coefficients are not of great interest, but they should be defined to limit dynamic effects on the results.

3.3 Determining location of supports

One of the main goals of this simulation is to gain information on how the location of the supports affects the post buckling behavior. The location of the supports was the only variable between all the simulations.

The locations of the beam supports were chosen based off the first mode shape of the clamped-clamped beam. The mode shape was defined based on the support location 0.03 inch from the top of the beam. This base distance was determined, so that the beam would contact the supports within the 0.001-inch prescribed displacement. Once the main mode shape was defined, support locations were randomly selected from a range around this curve. The x-coordinate of the supports were selected from a range that include 60% of the beams length but did not include a range in the middle of the beam equal to the diameter of the support. This was to prevent intersecting geometry. The y-coordinate of support locations were randomly selected from a range above and below the main mode shape. This range is defined by Eq. 3.3.

$$y_i = \frac{1}{2} f(x_i) + a_i\, f(x_i) \qquad \text{Eq. 3.3}$$

Where y_i is the randomly assigned y-coordinate, x_i is the randomly assigned x-coordinate, $f(x)$ is the function that defines the main mode shape, and a_i is a random coefficient between zero and one. Randomly assigning the support position allows for a larger coverage of the design space. From the random data, patterns can be drawn about how the supports should be positioned.

3.3.1 One contact Two contact

Based on the location of the randomly selected support locations, the beam may contact one or two times during the prescribed displacement. While having two contacts is beneficial, for greater customizability of the stiffness, both results were retained.

3.3.2 Discussion on unstable simulations

One problem, in the process of picking location supports, was that a significant number of simulations would go unstable. It appears that these simulations often occurred when one of the supports was at the bottom of the acceptable range. The unstable simulations were not included in further analysis, so limiting them was important. The lower y-limit was increased to try to avoid these solutions.

3.4 Automation

To gain a lot of simulation data, which will be used in the Machine Learning section, the simulation was created to minimize the required operator input. In the model, equations were used to define the x_1, x_2, y_1, and y_2 of the supports. The values for these variables are set the equation values from a text file. A Python script would randomly select the coordinates for a simulation and write the corresponding text file. Additionally, it would record the locations for that simulation in a Comma-separated Values (CSV) file. Simulations were ran using the SOLIDWORKS Task Scheduler. The 2022 version would automatically update the equations, from the text file (this did not seem to work for the previous version). Once a set of simulations had been run, the results were obtained by manually opening each one and saving the stress versus time plot data as a CSV. No automatic method of saving these results was found. One benefit of manually saving the data plots, is that the number of contacts made by the beam was manually recorded at the same time.

3.5 Post Processing

After obtaining the stress versus time results, further data inclusions and manipulations were completed. For example, data columns for the x and y positions of each support and the number of contacts the beam made with the support. Additionally, different simulation runs were combined, so a column with the file name was included. The data manipulations included, creating a displacement column from the time data by applying Eq. 3.4.

$$s_i = \frac{s_n(t_i - t_0)}{t_n - t_0} \qquad \text{Eq. 3.4}$$

Where s_i is the displacement at the time step calculated, s_n is the final displacement (0.001 inch), t_i is the time at that time step, t_0 is the initial time (two seconds), and t_n is the final time (202 seconds). A force data column is created by multiplying the stress data from the simulation (σ) by the cross-sectional area of the beam (A_{beam}) as shown in Eq. 3.5.

$$F = \sigma(A_{beam}) \qquad \text{Eq. 3.5}$$

The stiffness column is created by taking a numerical derivative of force with respect to displacement as shown in Eq. 3.6.

$$k_i = (F_i - F_{i-1})/(s_i - s_{i-1}) \qquad \text{Eq. 3.6}$$

After postprocessing, the important variables of stiffness, displacement, and force are defined, so the resulting force-displacement and stiffness-displacement diagrams can be analyzed.

3.6 Results

The finite element study gave three key results. First, it provided evidence that the variable stiffness device was feasible. Second, general trends in behavior based on the positions of the supports could be gathered. Third, data was collected for the machine learning models of a future section.

3.6.1 Variable Stiffness Device Feasibility

For the variable stiffness device to be feasible, it is required that the slender beam be able to carry a load even in the post-buckled state, and for the device to be useful there needs to be a difference between two support scenarios. If the device is not able to carry load, it can only be judged on its merits as a vibration suppression device and does not provide any weight minimization advantages to the system. If there is not a difference between different support scenarios, the system cannot be optimized to effectively counteract a variety of forcing frequencies. In these two ways, the finite element study can provide evidence supporting the devices feasibility.

The simulation study supports the idea that the active elements will be able to support load because most of the stiffness can be maintained. There are several examples of two-contact support scenarios that after the second contact have stiffnesses similar to the pre-buckling stiffness. To see examples, see Figure 3.4 (b). While high stiffness is important for structural rigidity, variable stiffness is important for vibration control.

The device should be able to be optimized for different frequencies. To be optimized, there needs to be variable inputs that correspond to variable outputs. The data finite element study demonstrates that there is variability in the resulting stiffness based on the positions of the supports. In Figure 3.4 (a), the support locations of several different locations are shown versus the reference mode shape while the stiffness displacement plots are shown in the (b) plot. These specific support scenarios were chosen to show a variety of stiffness levels. The differences between the stiffness curves of different support scenarios provide evidence to the feasibility to the device.

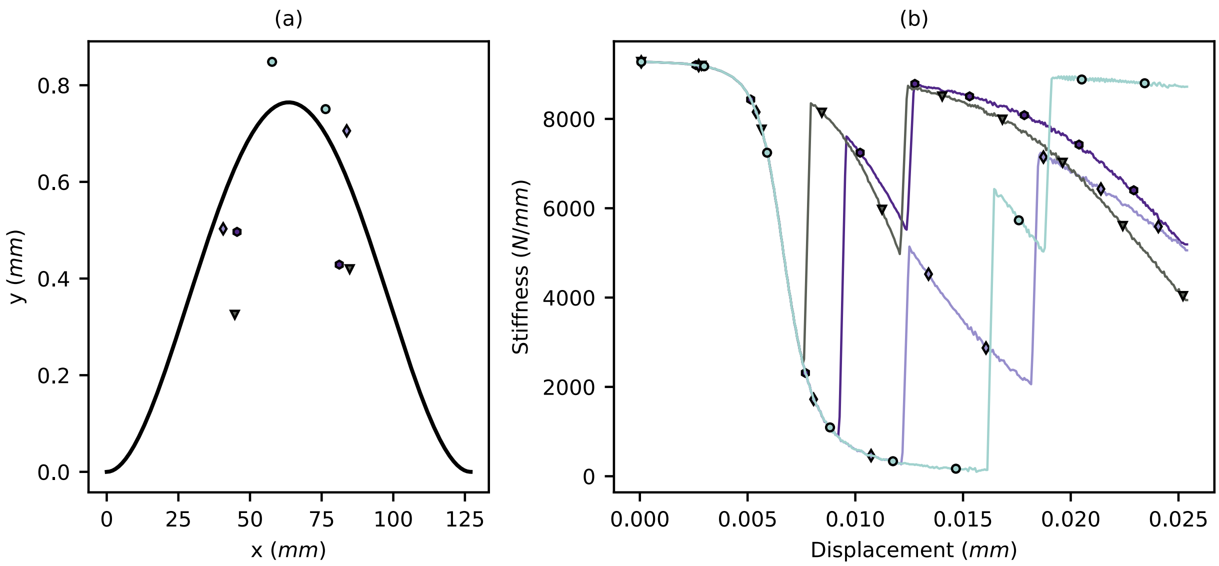

Figure 3.4: (a) Shows the positions of the supports for five different scenarios. (b) Compares the corresponding stiffness displacement plots.

Chapter 4 - Experimental Investigation of Nonlinear Stiffness Device

Experimental investigation provides corroborating evidence for the numerical simulation work in the last chapter. Since the objective is to create a device that suppresses vibration in the real world; it is necessary to support it with physical evidence. For these experiments, the goal is not to test vibration damping in a dynamic situation, but to test the stiffness modification on a real device and validate previous simulation work.

4.1 Experimental Equipment and Testing Samples

The equipment and samples used in experimentation will be described to support the results. A description of the buckling rig, MFCs, high Voltage amplifier, and beam samples will be included in this section.

4.1.1 Buckling Rig

A custom buckling rig was used in this experimentation. The basic components of the buckling rig are shown in Figure 4.1. The low friction rail and cart acts as a sliding fixed support on the left side of the test beam. It is actuated by the hand crank. On the right side, the load is measured at the fixed support by a 661.18E-02 MTS load cell. The entire structure is lifted off the table, to allow for load cell clearance. Now that the basic components have been introduced, the design objectives will be discussed.

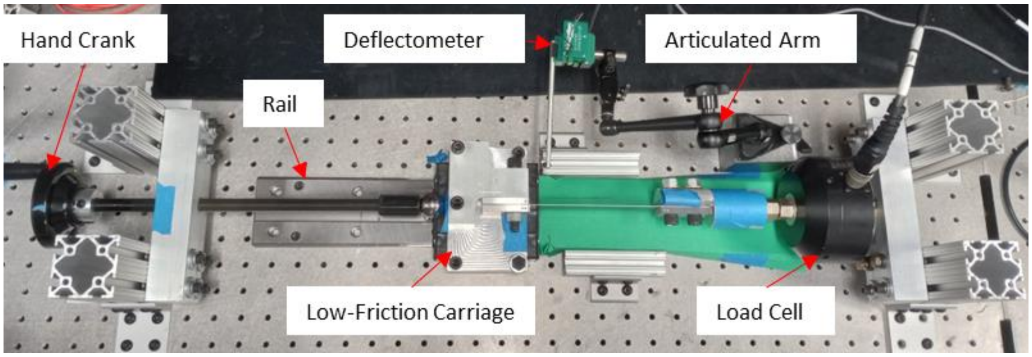

Figure 4.1: Diagram of the experimental buckling rig.

The buckling rig was designed to meet three requirements: approximate ideal loading conditions, conveniently position equipment, and sustain the buckling loads. First, the buckling rig is designed to approximate the ideal loading conditions. An ideal buckling test loads a uniform beam in pure compression, as such bending loading should be minimized. By orienting the beam to buckle in the horizontal plane, gravitationally induced bending is minimized. Additionally, the loading should be directly in line. Besides affecting results, too large of moments could damage the load cell. Another requirement for an ideal buckling test is rigid supports. Rigid supports are important for both load and deflection measurements. To increase the rigidity, larger mass and stiffness is required. The whole loading assembly must be raised off the table to allow clearance for the load cell which makes the assembly less stiff. Additionally, cost and other factors limit the mass of the design. Two important requirements for the design are to increase rigidity and decrease bending loads, to approximate an ideal buckling test.

The second important requirement for the design is equipment positioning convenience. Besides minimizing the effects of gravity, a horizontal layout is convenient for positioning equipment. For instance, the deflectometer is easily positioned using the articulated arm which is directly supported by the table. Similar flexibility in positioning would require at least one, but probably several structures to support the instrumentation if the buckling rig was designed for the vertical plane. For taking images, it would be easier if the buckling rig was oriented vertically, but this advantage is small compared to the positioning of measuring equipment and supports.

Finally, the buckling rig is designed for required buckling loads. Since slender beams are being analyzed, the buckling load is on the order of tens of pounds. A several thousand-pound load frame is not necessary and is not helpful. If something were to go wrong using a load frame that large, there would be a lot of damage. Similarly, this device requires less rigidity than a device designed for larger loads.

4.1.2 Macro Fiber Composites

The MFC was chosen for its flexibility and its ability to withstand large out of plane displacements. Additionally, it is convenient to apply to the host structure. Both the MFCs and the high voltage amplifier used were from Smart Material. The MFC was a d33 type MFC designed for actuation. The notation d33 communicates the direction of the voltage differential and the direction of the strain. For a d33 MFC, the strain develops in the same direction as the voltage differential as opposed to a transverse direction (see Figure 4.2). Since the goal is maximal displacement, a Smart Material Type-1 MFC was chosen. It allows for the largest voltage and therefore the largest actuation. Finally, the model, M2814-P1, was chosen to maximize the area over a one-inch-wide beam. The voltage source chosen for the necessary high voltage was the EEL1500-1 from Smart Material. Since this amplifier only has a single channel, the maximum voltage differential is 0 V to 1000 V, so MFC’s entire range cannot be used.

Figure 4.2: Strain Develops in the same direction as the Voltage Differential in a d33 MFC

4.1.3 Test samples

To increase the likelihood that the beams could be actuated by MFCs, beams dimensions were chosen to limit the buckling load. The material for the samples was AISI 1095 steel with a thickness of 0.02 inch. This material was available, but it also had these advantages high strength and fatigue toughness. The theoretical buckling load was calculated using the Euler buckling equation.

$$P_{cr} = \frac{\pi^2 EI}{l^2} \qquad \text{Eq. 4.1}$$

By calculating the buckling for a range of length values, a range of acceptable lengths could be narrowed down. First, the beam had to be long enough to allow for mounting two MFCs along the length of the beam. Secondly, the buckling load should be greater than ten pounds. Because loads measured are already near the bottom edge of the loadcells range, this problem should not be compounded by having beams with very small buckling loads. From these requirements, beams with a theoretical length of six, seven, and eight inches were chosen. In experimentation, the seven-inch beam was used for the first two MFC experiments while the six-inch beam was used in the third experiment which did not use any MFCs.

The manufacturing process involved cutting and minor preparation steps. The beams and their holes were cut out using a CNC waterjet. The edges of the holes were beveled to remove any sharp edges. Finally, the top edge of the beam was painted white for high contrast images. For some beams, these processes were sufficient, but others required MFCs.

The beam for first two experiments needed MFCs mounted to the beam. Mounting an MFC involves a process like a strain gauge, but there are differences. The MFCs were mounted on each side 0.4 inches from the edge of the inside hole of the beam. Digital calipers were used to make the measurement, and a mark was scratched into the beam. A second mark is scratched for the other side of the MFC. The beam should be taped down onto a moveable support. Because the bonding process involves glue, released film is placed below and eventually above the beam. Tape the release film to the beam at both ends and tape the beam sample down in the middle. A piece of Scotch Brite is used to rough up the surface, which gives the adhesive a better surface to hang onto. After roughening the surface, align the MFC. The goal is to center the MFC along the length of the beam. The terminals were located towards the center of the beam. Apply a large length of tape (enough that it easily doubles back on itself) and tape the MFC down to the beam. Carefully, pull the tape back, leaving the inside edge as a hinge. Ammonia and Chem Wipes are used to clean the bonding area on the beam, and on the MFC. Clean in one direction to prevent the accidental addition of impurities. The next steps involve mixing the DP460 epoxy and applying it to the beam; these steps need to be completed in a timely manner before the epoxy begins to set (it has a work life of one hour). To ensure the proper ratios for mixing the epoxy, let out enough epoxy until both reactants flow evenly. Let out a small amount of epoxy into a mixing cup. The epoxy is mixed with a metal spatula. Since the amount is small, it is important to mix well. Any unmixed epoxy will lead to a weaker bond. Apply the mixed epoxy evenly to the bonding areas. Unstick the loose end of the tape, and let the MFC lay flat, while not applying tension to the tape. Repeat process for the second MFC. Now, place the second layer of release film, followed by the breather cloth, and the evening beam. Place the assembly in the oven and add the weight on top. The beam should be allowed to remain in the oven for two hours at 71 °C for maximum strength. After curing in the oven, the beam was left until the next day, but this is not expressly necessary. Some clean up with a razor blade was required where epoxy leaked beyond the intended area. The process should be repeated for the second side. Finally, high voltage wires were wired onto the terminals of the MFCs; the leads should be at 60 to 100 centimeters long. Additionally, polarity matters for MFCs, so there should be a method for marking positive and negative leads. With the MFCs bonded and their leads soldered, the samples are ready for testing.

4.2 Test process

The testing process is essentially uniform for each test and is useful and describing the components of the experiment. The first step is to load the test sample into the buckling device which involves placing the sample between the clamp faces and securing the clamp faces with the two bolts. Once secured, the test piece will be ready for multiple tests. Next, it is a good idea to test if there is an obvious propensity for buckling in a certain direction. Ideally, it would buckle with equal likelihood in either direction. Turn the hand crank to apply compression until the sample buckles. The goal should be to apply load until it buckles and then stop to avoid plastic deformation in the beam amplifying the problem. Test the beam several times in this way to see if there is obvious propensity to buckle in a certain direction. If so, the tail of the buckling device will need to be realigned.

To align the device, there are two main methods. The first method involves using a more rigid beam to align the loosened tail end. The second method involves leaving the slender test beam in (perhaps loosening) and making small rotations until there is an equal propensity to buckling. For the first method, unload the slender beam and load a beam of quarter inch hard steel leave the clamps loose for now. Loosen the screws that hold the tail end in place. Tighten the clamps which will cause the tail end to align itself. Now tighten the tail end. For more extreme misalignments, it might be necessary to rotate the tail end manually. In this case, one should be careful if the test sample is left in the clamps. If the tail end is rotated too far, the beam will experience large deflections. Rotate the end slowly until the beam will buckle in both directions. Mark the initial location, for reference before moving the tail end. In this experimentation, there was a large propensity to buckling in one direction, so there was not a point where there was an equal likelihood in buckling in either direction, so the tail stock was manually rotated until, holes for the tail stock brackets prevented further rotation.

Now that the beam is loaded and the device has been straightened, it is time to apply a preload to the beam. The preload is applied because the device is designed to be part of the object's structure, so it is not reasonable for it to be entirely relaxed. The preload is set to about ten percent of the theoretical buckling load. To apply the preload, turn the hand crank while watching the load on the VI's front panel. The preload is small, so it can be difficult to get to the right load. The threaded rod was too loose, so it allowed a relaxation in the load. For better actuator rigidity, a second support for the threaded rod should be considered.

If MFCs are required like in the first two experiments, the next step is to turn on the MFCs. This is simply done, but by pressing a button on the second VI, but it requires a set up set prior to applying the preload. To set up the MFCs, the connections for the MFCs that will be actuated should be made. Originally, the goal was to use relays and a LabVIEW virtual instrument (VI) to make these connections. This has the advantage of being able to distribute high voltage away from the voltage amplifier. In experimentation, there was an undiagnosed hardware problem that caused a significant retention of voltage on one of the MFCs. Because of this problem, the method of choosing MFCs was changed, and Wago lever nuts were used to make connections. To make a connection, the wire from the MFC's negative terminal is connected to the amplifier ground while the positive terminal is connected to the amplifier high voltage. All the MFCs that should not be connected, have their negative terminals connected with a lever nut, and their positive terminals connected with a separate lever nut. This was helpful for preventing a stray exposed wire from causing trouble, and it was useful in preventing incorrect connections. Additionally, the MFCs are marked to make clear which wire is attached to which MFC. All connections were checked with a digital multimeter (DMM) first at a lower voltage so that normal leads could be used and later at high voltage with 100x attenuated leads. Once the high voltage is turned on, the beam will experience deformation causing a change in preload. This is expected and was considered the fairest way to compare MFC scenarios. It should be the fairest method of comparison because no MFCs is the control case, and different MFC scenarios are the variable cases. Additionally, there is no intention to actively apply preload; the expectation is that it would be applied during assembly. The MFCs are an important part of the experimentation; connections should be made carefully and the difference in preload caused by the beams subsequent deformation should be recorded in data collection.

After turning on the MFCs, the next step is turning on the data collection. In this experimentation, three data collection devices are used to get force data, axial displacement data, and transverse displacement data. A National Instruments NI 9237 Data Acquisition Device (DAQ) is used to digitize signals. National Instruments LabView Virtual Instrument (VI) controls data collection, displays data, and saves it. The force data is collected by an MTS load cell in line with the applied load. It is a strain gauge type of load cell. Using National Instruments LabView, measurements are obtained, calibrated, and saved. The load cell's default calibration was used for the force measurements. For axial displacement, the 3540-025T-ST deflectometer by Epsilon Technology is used. The deflectometer is mounted on an Noga articulated arm, and measures displacement directly off the sliding support. The deflectometer is another strain gauge device. In this case, a custom calibration was completed. The final device to collect data is a Pixellink PL-D7620CU-T camera with 75 mm lens. Pixellink's software is used to control the camera and save images to a laptop. There is a basic burst mode that is used take images every half second for 60 seconds. This method allowed for images to be taken while the operator was turning the hand crank.

Hand crank operation is certainly not ideal, but there are some important factors that can help its accuracy. The first important factor is to apply load slowly. The loading scenario is quasistatic, so time is not a factor. Additionally, applying the load slowly allows for more data collection points which is beneficial for averaging out error. The second important factor is axial displacement is the important parameter. To slowly turn the crank, the operator's hand was applied to the outside radius of the crank, but this requires transitions to be made. Since displacement is the important variable, a transition can occur if the crank's position is maintained. The displacement versus time plot for the data should be continuous, but there is no need for it to be smooth. Force data should also be continuous, so it is important to prevent the crank from relaxing. By understanding the quasistatic loading condition, the operator can focus on maintaining smooth force and displacement data while maintaining a slow actuation speed.

When applying the compression, the operator should review force deflection curve. After the plot has shown the expected contacts, the operator should continue application until the plot bends over. At that point, the experiment can be ended after it has reached a reasonable threshold. Tests were continued until it was clear that data collected was outside the area of interest for two reasons. First, too much data is less of a problem than too little. Secondly, the displacement reading by the extensometer was not zeroed for every test, so data could start from -0.001 inches. It could be useful to have these different starting points, but it was eventually zeroed in post processing. It is recommended that there be a software method to arbitrarily zero data for each test as it would simplify ending a test. Once the test threshold has been reached, stop saving data from the VI. The camera data will continue automatically, but at this point relax the beam using the hand crank. The MFCs are turned off with the beam relaxed.

4.3 Post Processing

4.3.1 Measurements

To post process the experimental data, several steps are taken including zeroing the displacement, applying a Savitzky-Golay filter, and finally calculating the loads from the force displacement diagram. To zero the data, the average of the first 1000 points of displacement data is taken and subtracted from all the displacement data. Next, a first-order, Savitzky-Golay filter is implemented using a function from SciPy's signal module.

The important parameter to adjust for the filter is window size; a larger window size makes cleaner data, but also results in a loss of information. To optimize window size, the experimental force and displacement data is plotted against time. The filtered data is plotted on top of the raw data. The window size should be increased until the filtered data begins to lose features of the raw data. If the raw data has a sharp change, the filtered data should not round it off. As a rule, if the filtered data stays inside the raw data, the filtered data will be a good, if conservative, representation of the raw data. Figure 4.3 shows the plots of force and displacement data versus time. In filtering the data, the objective was to increase the window size until the filtered data went outside of the raw data. A window size of 300 was chosen; this may be too conservative. From the filtered data, the force-displacement curve is created, and further numeric analysis can be completed.

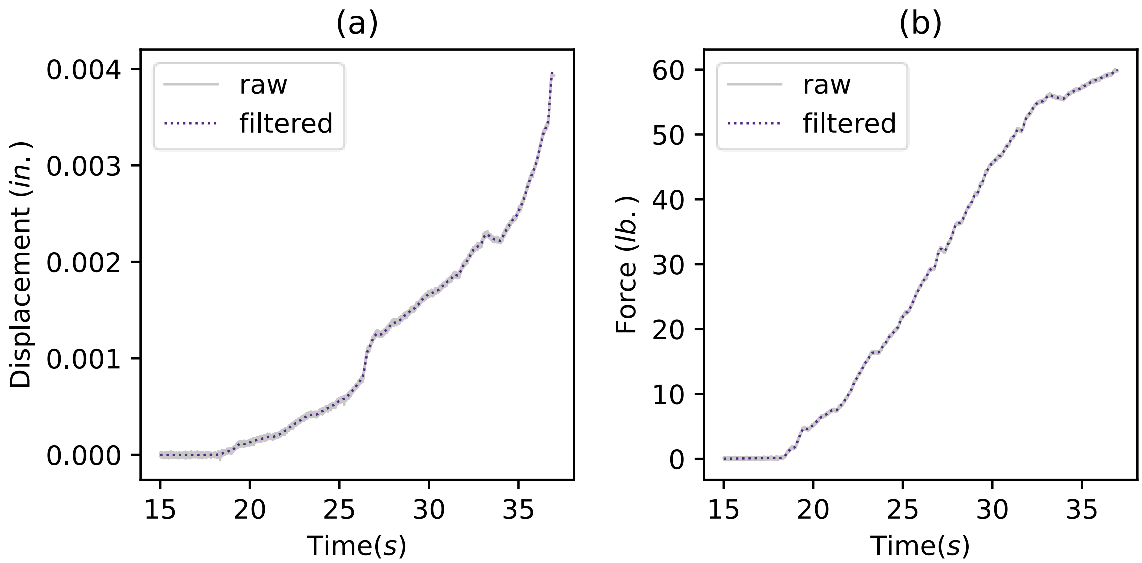

Figure 4.3: (a) Compares raw and filtered displacement data. (b) Compares raw and filtered force data. The filtered data should remain on top of the raw data.

4.3.2 Images

To use the digital images to measure distance, a calibration picture is taken of a ruler. The goal is to take the measurement picture at the same level as the top of the beam. This is important because the distance from the lens will change the distortion. The calibration for this experimentation is carried out in the transverse direction because that is the direction of greatest interest. From the calibration image, the pixel locations of two lines eleven inches away are recorded. This data is used to create the conversion factor shown in the following equation.

$$1\ inch = 486\ pixels \qquad \text{Eq. 4.2}$$

From this conversion factor, the data resolution can be calculated.

$$1\ pixel = 0.00206\ inch \qquad \text{Eq. 4.3}$$

Another consideration for the images is exposure. The goal is to produce a high contrast image, so the edge of the beam and stopper is clearly defined. To better define the edge of the beam, the image was intentionally overexposed.

For post processing, the image undergoes basic image manipulation, edge detection, and conversion to positional data. First, the image is converted from color to a black and white image. Next, the image is cropped more tightly to the beam to minimize extraneous distractions. Because of the nature of the image, it is possible to decrease the noise in the image by removing the midtones of the image. If the goal is to find the white beam, any pixel with a value below a threshold (200 for example) can be assigned a black value. When finding the black support, the opposite manipulation can be done. Using this type of filtering removes a lot of noise from edge detection results.

For edge detection, the Canny Edge Detection Algorithm is used. The function used is from the OpenCV Python Module. This algorithm convolves a matrix over the image to calculate partial derivatives. After the edges have been detected, a function was created to take the data found in the Canny Edge Detection output image and convert it to an array of coordinates. Since the output image from the Canny Edge detection contains only white (edges) and black values (not edges), the first part of this simply collects all the coordinates of the white pixels. In the second part of the program, loops through the data by columns, and takes data points within a range from a predicted row. The average of the new data becomes the predicted row for the next column. With this program, the data tracing the beam edges is collected. Similarly, one can find the edge of the black support.

Once there is positional data for the beam edge, the data should be straightened. First the offset of the right edge of the beam is removed since it is common to all steps. Next, a slope is calculated from the first and last points of data on the beam. The points used in the slope calculation were identified manually, to counteract the effect of shadows at the end of the beam. With the slope calculated, a transformation matrix is created to remove the slope from the data.

This transformation matrix is created for the first image, one before substantial deformation. The other images use the same transformation matrix to avoid applying multiple scales to the data.

Geometry is used to find the distance between the beam and the support. The point on the support with the largest y-value in the transformed domain is assumed to be the point of contact for the beam. From the polynomial fit data of the beam, a localized linear fit is made for the data. From this localized linear fit, the perpendicular line going through the point of contact can be found. The intersection point between these two perpendicular lines is used along with the point of contact to calculate the distance between the beam and the support. This distance is primarily important for creating simulations with similarly positioned supports.

It would be expected that the line found by least squares regression would come from the raw data, but in this work, it came from a polynomial fit of that data. A linear regression on the raw data for a localized area created a line that was clearly not tangent to the polynomial fit. When viewed over the entire data range, the polynomial fit creates an accurate representation of the data, so it is reasonable that it also should be able to provide a reasonable approximation for local slope of the beam.

4.4 Experiments

The experimental work is divided into three experiments designed to demonstrate a different attribute of the device.

4.4.1 Experiment 1

The first experiment has two main purposes. The first is to demonstrate that MFCs can actuate the direction the beam buckles. Secondly, the goal is to understand the influence that the MFCs have on buckling. No supports are used in this experiment because the pure buckling behavior is of interest.

To communicate the MFC scenario, a diagram is created. Figure # compares a more detailed drawing of the experiment with a much simpler icon form. If the beam is divided into quadrants, like Figure 4.4 (a), the MFC in that quadrant is graphically represented by the corresponding square in Figure 4.4 (b). If an MFC is on, it is represented by a red square. If it is off, it is represented by a white square. These diagrams will be used to communicate the scenario in the results section.

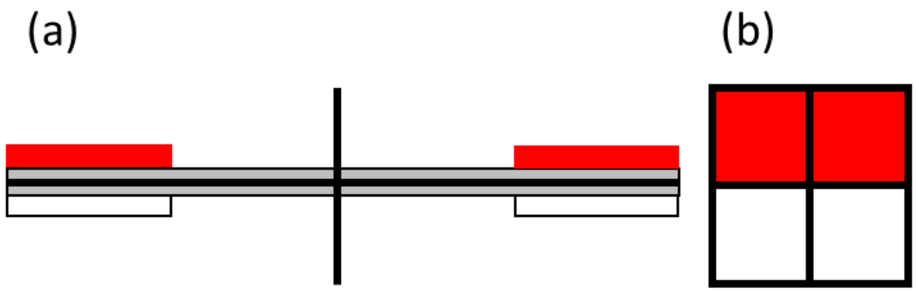

Figure 4.4: The two figures show two representations of how to communicate which MFCs are on. The red rectangles indicate that the MFC in that position is on. The white rectangle indicates the MFC is off.

With a standard form to communicate the scenario, it is possible to group scenarios based on expected common features. The first control group contains the scenario where no MFCs are turned on. The bending group contains scenarios where MFCs on one side (top or bottom) are turned on, but not the other side. The third important group is the axial group. The top MFCs mirror the bottom MFCs in these scenarios. The bending group is expected to have a lower buckling load because of the increased destabilization it experiences, as compared to the axial group. By grouping similar cases, it is possible to make predictions about their behavior.

4.4.2 Experiment 2

The second experiment is designed to demonstrate the ability to modify device stiffness with different support locations. Two extreme support scenarios were created based on the simulation work in Chapter 3. The first support scenario (Case 1) has a single support positioned near the end of beam positioned so it relatively far from the beam. The second support scenario (Case 2) has a single support located near the middle of the beams length and positioned laterally near to the beam. Finally, to demonstrate the ability to maintain a higher percentage of the device's stiffness, a third scenario (Case 3) was created where the support was located even closer to the beam near the same location. The three different support scenarios were actuated with one or two MFCs. Finally, they are compared to the case with no supports.

4.4.3 Experiment 3

The final experiment is designed for comparison with simulation. Since no simulation was done for a beam with MFCs, this experiment used a beam without MFCs. The experiment was completed with the six-inch beam. A single support was placed at 2.48 inches from the fixed edge 0.032 inches from the beam. These values come from digital image analysis, so they may not be accurate. A simulation was created with the same beam material, dimensions, and support locations for comparison.

4.5 Results

4.5.1 Experiment 1

The MFCs were able to affect the direction of buckling. Initially, this did not seem to be the case. Because of some physical problems with the buckling rig, the initial alignment using the thick steel sample caused a significant propensity to buckle in one direction. By adjusting alignment, so there was an initial angle in the opposite direction the situation was improved. With the alignment adjustment, the MFCs were easily able to actuate the direction of buckling, although there was still a favored direction for the case with no MFCs.

The force-displacement and stiffness displacement curves for the axial cases (Figure 4.5) and the bending cases (Figure 4.5) are shown. Notice that the case without any MFCs is the stiffest over all cases with MFCs turned on.

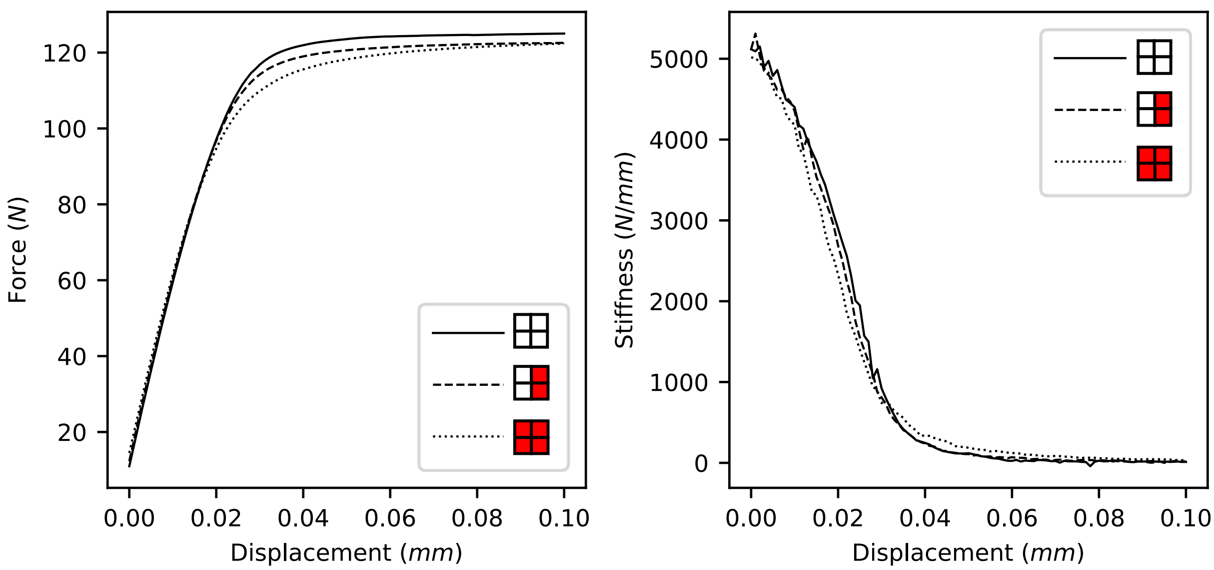

Figure 4.5: The force-displacement and stiffness-displacement plots for the axial cases.

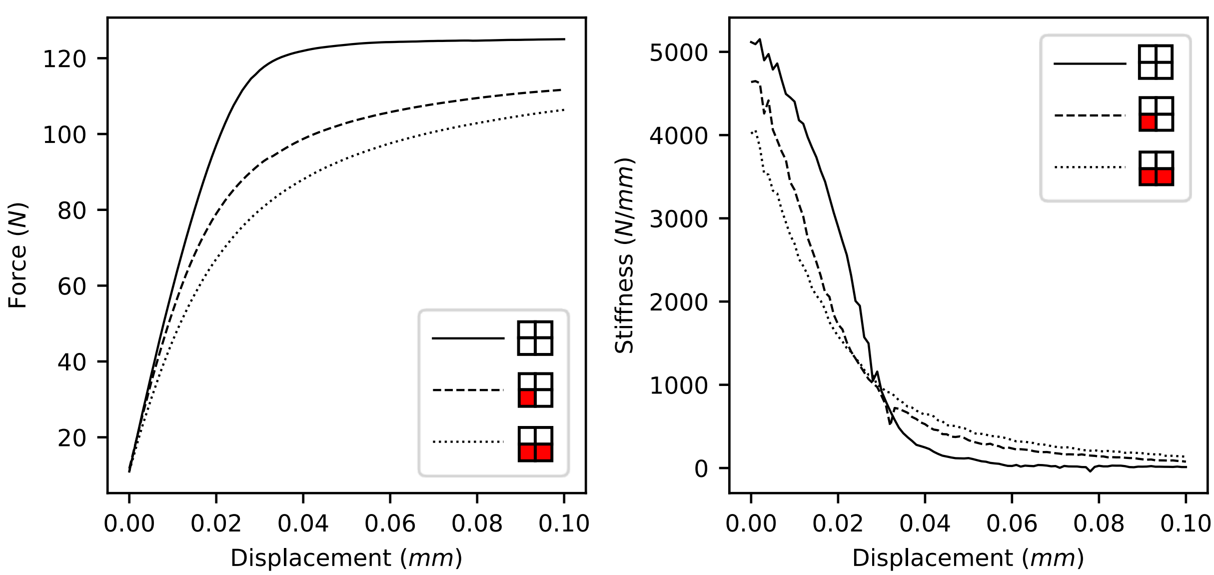

Figure 4.6: The force-displacement and stiffness-displacement plots for the bending cases.

4.5.2 Experiment 2

For the second experiment, the results for the three different support positions are compared in Figure 4.7.

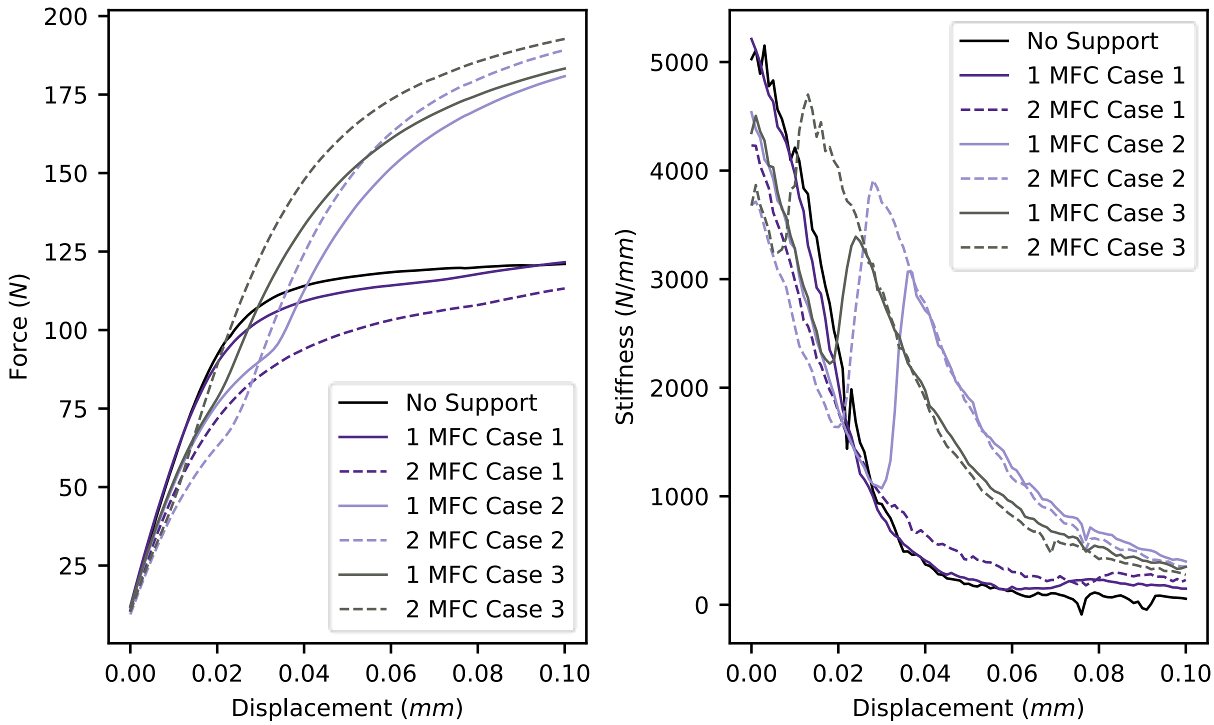

Figure 4.7: A comparison of the force-displacement and stiffness-displacement plots for three support cases.

4.5.3 Experiment 3

The goal of the final experiment is to provide experimental support for the simulation results from the last chapter. To provide a good comparison, a simulation is created with the same support scenario as the one experimentally tested. To improve these results, the picture data taken from the experimental test is used to solve for the location of the support used in simulation.

4.5.3.1 Measurement Results

The force deflection results for the experimental test are compared with simulated results. The plot showing the comparison is shown in Figure 4.8. For this comparison, the support of the simulation was located, so the contact point is located at the position described earlier. Additionally, the simulation continues to have the 2g gravity applied.

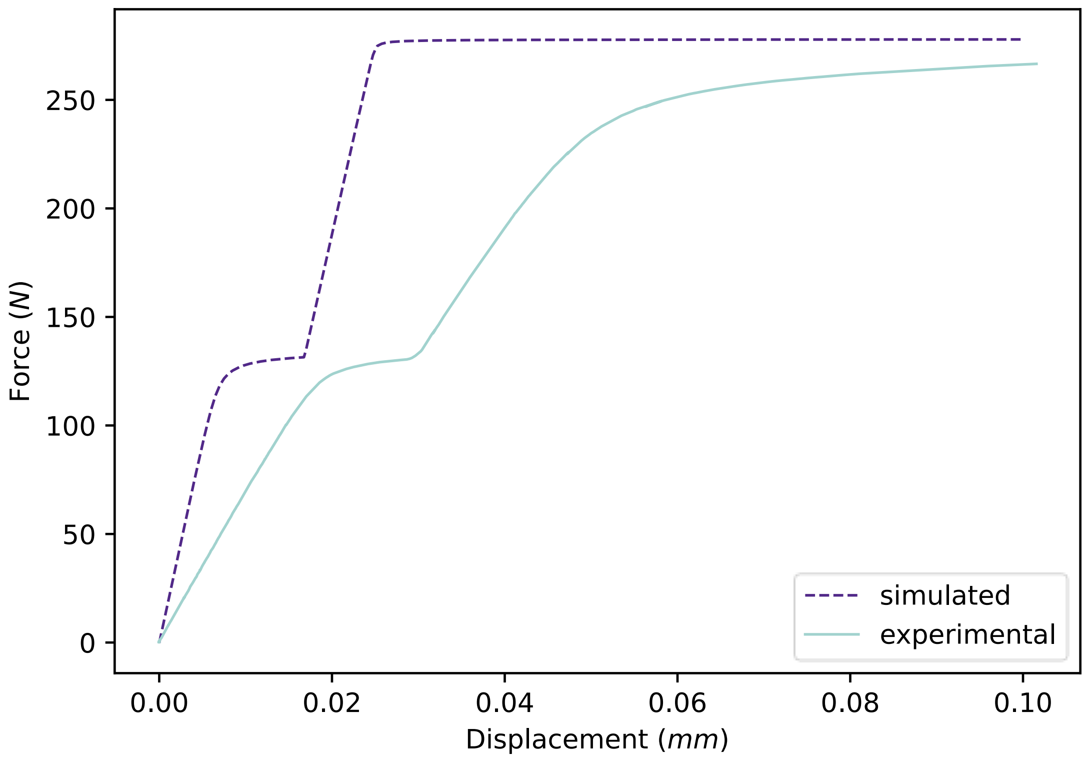

Figure 4.8: Comparison of simulated and experimental results for Experiment 3. There is an obvious difference in the slopes of the linear portions which is partially because of deflections experienced by the real, fixed supports.

The most obvious observation is that the curves do not line up with each other at all. The linear portions of the experimental curve have a much lower slope than those of the simulated curve. Initially surprising, there is an explanation for this. When viewing this problem from an analytical perspective the fixed support is perfectly fixed, and this is what is modeled in the simulation. In the experiment, the fixed end moves. The measurement for the extensometer measures the movement by the sliding end, but it does not account for the movement by the fixed end. By ignoring this displacement, the experimental measurements over predict the displacement the beam experiences which is consistent with what is seen in the plot.

Because of significant disparity between the displacements, the forces of key parts become important results for comparison, but they also need some clarification. As mentioned previously a 2g gravitational field, is applied to the beam. This gravitational field was used in simulations from Chapter 3 and ensures the simulation has a singular solution. The gravity is essentially affecting the eccentricity of the beams loading, so by changing its strength the buckling load changes.

4.5.3.2 Prelude to Picture Results

For comparison of picture data, simulated data and experimental data must be matched, so the data corresponds to the same portion of the experiment. From previous results, it was shown that the displacements between the experimental data and simulation data do not match, so for the points of comparison are found by comparing the critical points in the force displacement diagram. With the simulation data, it is a simple process of identifying displacement when the beam reaches the critical point, and then converting to time by the proportionality constant between displacement and time. Data for the displacements in the x and y direction is extracted for each node on the bottom edge of the beam. These displacements along with the original locations of these nodes will give the position for each point at the critical time step. For the experimental data, the process is more involved. First, the important displacement is found using Figure 4.8. From this point, that displacement is matched on the displacement versus time chart to get a relative time from the beginning of the data collection. From this relative time, it is possible to get an absolute time which can be compared with images timestamps, to select an image for comparison.

4.5.3.3 Picture Results

Using the previously discussed method to obtain similar results, the shape of the simulated beam can be compared to the shape of the experimental beam can be compared in Figure 4.9. One of the problems is that the support found by the method intersects the experimental beam.

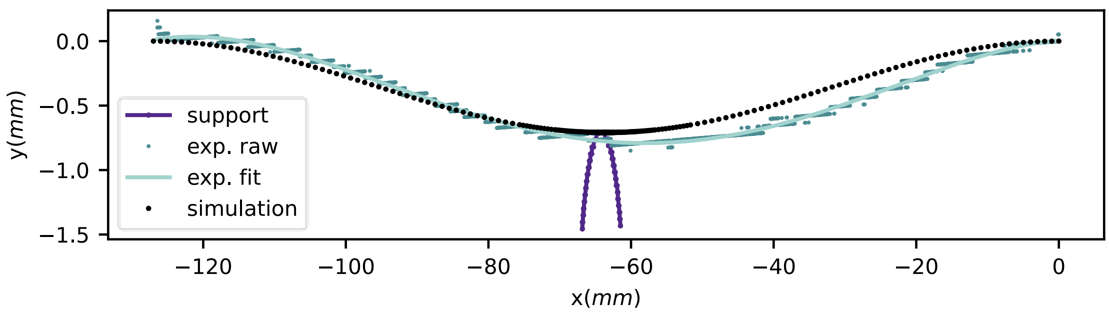

Figure 4.9: Comparison of experimental beam shape and simulated beam shape at contact.

The results appear to be quite different. One problem is that as the beam and the support get close together is that it becomes harder for the camera to make out the edge. The most extended edge of the support may not even be seen by the camera. A difference in between the position of two rows of pixels was shown in Eq. 4.3 to be 0.002 inches. Given the pixel size, the photograph may have too large of mesh size to effectively identify the contact point. With more resolution, the position of the support would be clearer which would likely increase the agreement. Besides image resolution, the angled supports of the buckling test rig affect the shape of the experimental data. These initial angles could explain the difference between the sloped portions of the graph. Finally, a polynomial was used to fit the picture data because of the simplicity of implementation. Better results may be obtained by applying the analytical complex trigonometric function, but this becomes more difficult with the contact.

Chapter 5 - Machine Learning Modeling

In this chapter, a neural network model of the physical system is developed. The purpose, methods, and results of this work will be discussed.

5.1 Objective of Machine Learning Work

In previous chapters, several factors were shown to change the characteristics of the system. Of these factors, stopper positions are modifiable and change the system's stiffness-displacement plot. Stopper positions should be optimized to allow for maximum stiffness control. In this phase of research, a neural network is trained to predict the stiffness paths based on stopper positions. A neural network can model the nonlinear behaviors of the system as well as provide for quick design iteration.

5.2 Machine Learning Methods

In this section, the process of creating a neural network, defining the data inputs and outputs for the neural network, and tuning the network will be discussed.

5.2.1 Neural Network: Creation

To create the neural network, TensorFlow-Keras code was used to create the network, so most of the creation process involves defining the parameters of the network. TensorFlow-Keras code preforms all the mathematical operations to create, train, and obtain results from the neural network. The creation, training, and testing of the neural network requires a few lines of TensorFlow-Keras code, so the main task involves understanding the parameters of the network and determining the data.

5.2.2 Neural Network: Data Definition

Defining the data given to the neural network is the most important part in improving the performance of the neural network. In this research, the inputs to the neural network include the x and y coordinates of each support and the displacement of that step. Initially, the network was designed to output force data, but outputting stiffness proved better.

The main purpose of the neural network is for the optimization of the support positions, so it should be able to be modified by the neural-network user. While it seems that the displacement data could be left out when considering a whole force-displacement or stiffness-displacement curve, the displacement input was included to increase the speed of training process. By including displacement data, it was thought the neural network would be able to categorize the force-deflection plot into distinct sections. It is expected that the displacement data could provide a future benefit by allowing the designer to specify a displacement range during the process of positioning the supports. For the outputs, since the authors more intuitively understood the force deflection diagram, force was initially chosen for the output. However, to improve neural network performance, it was thought that the stiffness data would provide greater differentiation between cases. The results for the network improved by adopting this change. The definition of the data affects the network purpose and performance.

Beyond defining what the data inputs and outputs are, an important step is the preprocessing of the data of the network. The first step of preprocessing was modifying the data set to exclude most of the initial buckling simulation data. Since this data is repeated for every support scenario, it gets weighted too heavily in the error calculations. Additionally, simulation cases with discontinuities in the force-deflection data were manually removed. The next step is to divide the data into three sets: training, testing and validation. The training set included 90% of the data, the testing set included 5%, and the validation set contained the remaining 5% of the data. In addition to cleaning data, preprocessing included Z-Score Normalization which creates a data set with a mean of zero and a standard deviation of one. The normalization gives the neural network a sense of scale for the data. For a fair analysis, the normalization transformation should only be created based on the training data, but it is applied to all data sets. All these preprocessing steps are important for improving training, but it is also critical for having a fair analysis.

5.2.3 Neural Networks- Hyperparameter Tuning

Hyperparameter tuning is the process of experimentally adjusting hyperparameters to increase the performance of the neural network. At first, the hyperparameters were adjusted manually, but to explore a greater portion of the hyperparameter space, Keras Tuner was used to conduct a random parameter search.

5.3 Hyperparameter Tuning Experiments

Hyperparameter tuning is essentially an experimental process. Previous experience is often a major factor in the decisions which are made. This section will develop the experiments complete, will describe results, and explain decisions made.

5.3.1 First Experiment

After some initial work to understand some parameters that seemed to work well, it was time to work on some hyper parameter tuning. There are two methods to approach hyper parameter tuning, grid search and random search. Given a selection of appropriate values, a grid search would train a model for every combination of values, while a random search would train models that would randomly select combinations from the values. According to Bergstra and Bengio, a random search is much faster [17]. A main idea they present is that often only a few hyperparameters significantly affect the results. In a random search, different values for a critical parameter will be observed much sooner than in a grid search. For a grid search, maybe the first 20 models have the same value for that parameter, while in a random search the first 5 models might have 5 different values for that parameter [17]. Based on this research it was decided a random search hyperparameter tuning approach was taken. To understand the effect of certain hyperparameters, sometimes this approach was modified to some extent.

After some ad hoc analysis to identify possible ranges that might work, the next step was to develop a more thorough, documentable approach. For this initial test, 99 models were created. The following hyperparameters were fixed: model type, batch size, activation, optimizer, and epochs. Early stopping was implemented to move on from bad models more quickly. The following parameters were tunable: neurons in each layer, number of layers, and the learning rate. A summary of the chosen values and ranges for these parameters is given in Table 5.1.

Table 5.1: Possible values for hyperparameters in random parameter search. *Chosen for each layer, but each layer in a network must be smaller or equal to the previous layer.

Hyperparameter	Possible Values
Model Type	Feed-Forward, Fully-Connected Neural Network
Number of Layers	3-10
Neurons in Each Layer*	1024, 512, 256, 128, 64, 32, 16, 8
Activation	RELU (hidden layers), None (output layer)
Optimizer	Adam
Learning Rate	0.0001, 0.001, 0.01
Epochs	1000 (Early Stopping Implemented)
Patience	100
Batch Size	8192

Many of the choices made by parameters was based on the ad hoc analysis. The goal for the epochs choice was to give enough time for the learning rate to make a difference, but to keep the training time reasonable. Previously, networks were trained on a range from 250 to 2000 epochs. Usually, a good network will improve rapidly, so it should stand out. However, the number of epochs does affect the best learning rate. With a hard cap on epochs, there is a bias against smaller learning rates, as they do not train as fast. Additionally, the number was chosen, so that most networks would complete the maximum number of epochs. The worst networks would be early stopped, but a majority of the networks would reach the full 1000 epochs. This also explains why a patience of 100 was chosen. If a network could not improve in a tenth of its training time, it probably is not going to be great, but it would allow most networks to keep training. For batch size, previous analysis seemed to indicate that the batch size was not a particularly important hyperparameter. With a larger batch size, training took less time. I chose 8192 because it increased training speed, but additionally, each epoch took long enough that the verbose training data was possible to read while it scrolled.

From the ad hoc experimentation, a learning rate of 0.001 proved to be a good choice, so it was bracketed by an order of magnitude on each side. Many layer neural networks seem to have their advantages, but there is also a problem with vanishing gradients. A broad range of options were chosen, although none would be considered super deep. Finally, there is the number of neurons in a layer. It was decided that the number of neurons should either stay the same or go down from layer to layer. By using a random number generator and some if else logic, there is a 40 percent likelihood that the next layer will have the same number of neurons, a 40 percent chance that the number of neurons will decrease by a power of 2, and a 20 percent chance that the number of neurons will decrease by 2 powers of 2. Additionally, the minimum allowed number of neurons in a hidden layer was 8.

5.3.2 Results

Now, that there is a basic understanding of the first experiment it is time to look at the results. The results are plotted below in Figure 5.1 (see Figure 5.2 for a zoomed in view of the best networks). There are four subplots that each analyze a certain parameter. Each plot has the validation set mean squared error plotted on the y-axis and the number of parameters plotted on a log scale on the x-axis. Lower error is better, and fewer parameters is better. In each subplot, the networks are grouped by a different hyperparameter.

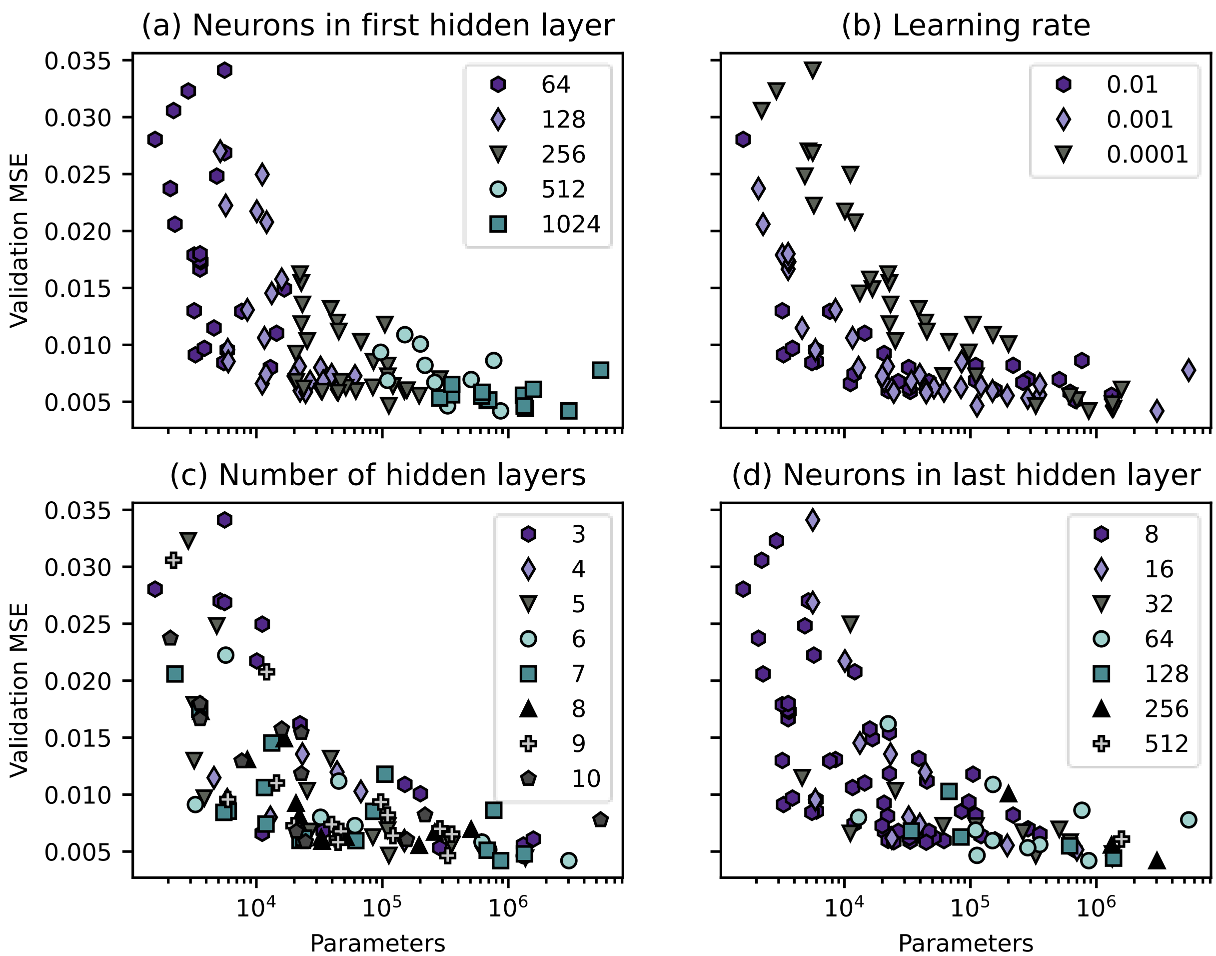

Figure 5.1: Comparison of validation set mean squared error versus number of parameters for 99 neural networks. The subplots group the networks according to different categories: (a) neurons in first layer, (b) learning rate, (c) number of hidden layers and (d) Neurons in the last hidden layer. The smaller the number of parameters the better (to minimize space requirements of the network) and lower validation error is better.

In Figure 5.1 (a), the number of neurons in the first hidden layer is compared. The trend in the graph indicates that the general rule is larger is better. However, the graph also shows the obvious result that more neurons in a layer means more parameters. As a rule of thumb, the number of parameters will grow proportionally to the square of the number of neurons. For a dense neural network, the equation for number of parameters is shown below.

$$p_i = n_{i-1} \times n_i + n_i \qquad \text{Eq. 5.1}$$

Where p_i is the number of parameters associated with layer i, n_{i-1} is the number of neurons in the previous layer, and n_i is the number of neurons in that layer. The networks with larger numbers of neurons in a layer performed better, but there is a significant cost in the number of parameters.

Figure 5.1b compares learning rates. The trend shows that as the neural network gets larger the best learning rate decreases. For networks with the fewest parameters, a learning rate of 0.01 gives the best results. In the range from 10^4 to 10^5 a learning rate of 0.001 is the best. Finally, for the really large networks a learning rate of 0.0001 is the best. It should be remembered that the best learning rate also has to do with the number of epochs, so it is possible that with more epochs a smaller learning rate may perform better.

Figure 5.1c compares number of layers. There does not seem to be any clear trends. If there is a trend, it seems that the moderately long networks perform well. Generally, the long and really short networks are worse, but there may be other factors involved.

Figure 1d compares number of neurons in the last layer. There may not be any clear trends, but there are a few areas of interest. In Figure 5.2d compare the data boxed in the 10^4 to 10^5 (dashed box) to the data boxed beyond that range (dotted box). The data in the dashed box essentially all have 8 neurons in their last layer, while none of them in the dotted region do. Other factors are likely more critical, but it could be a factor.

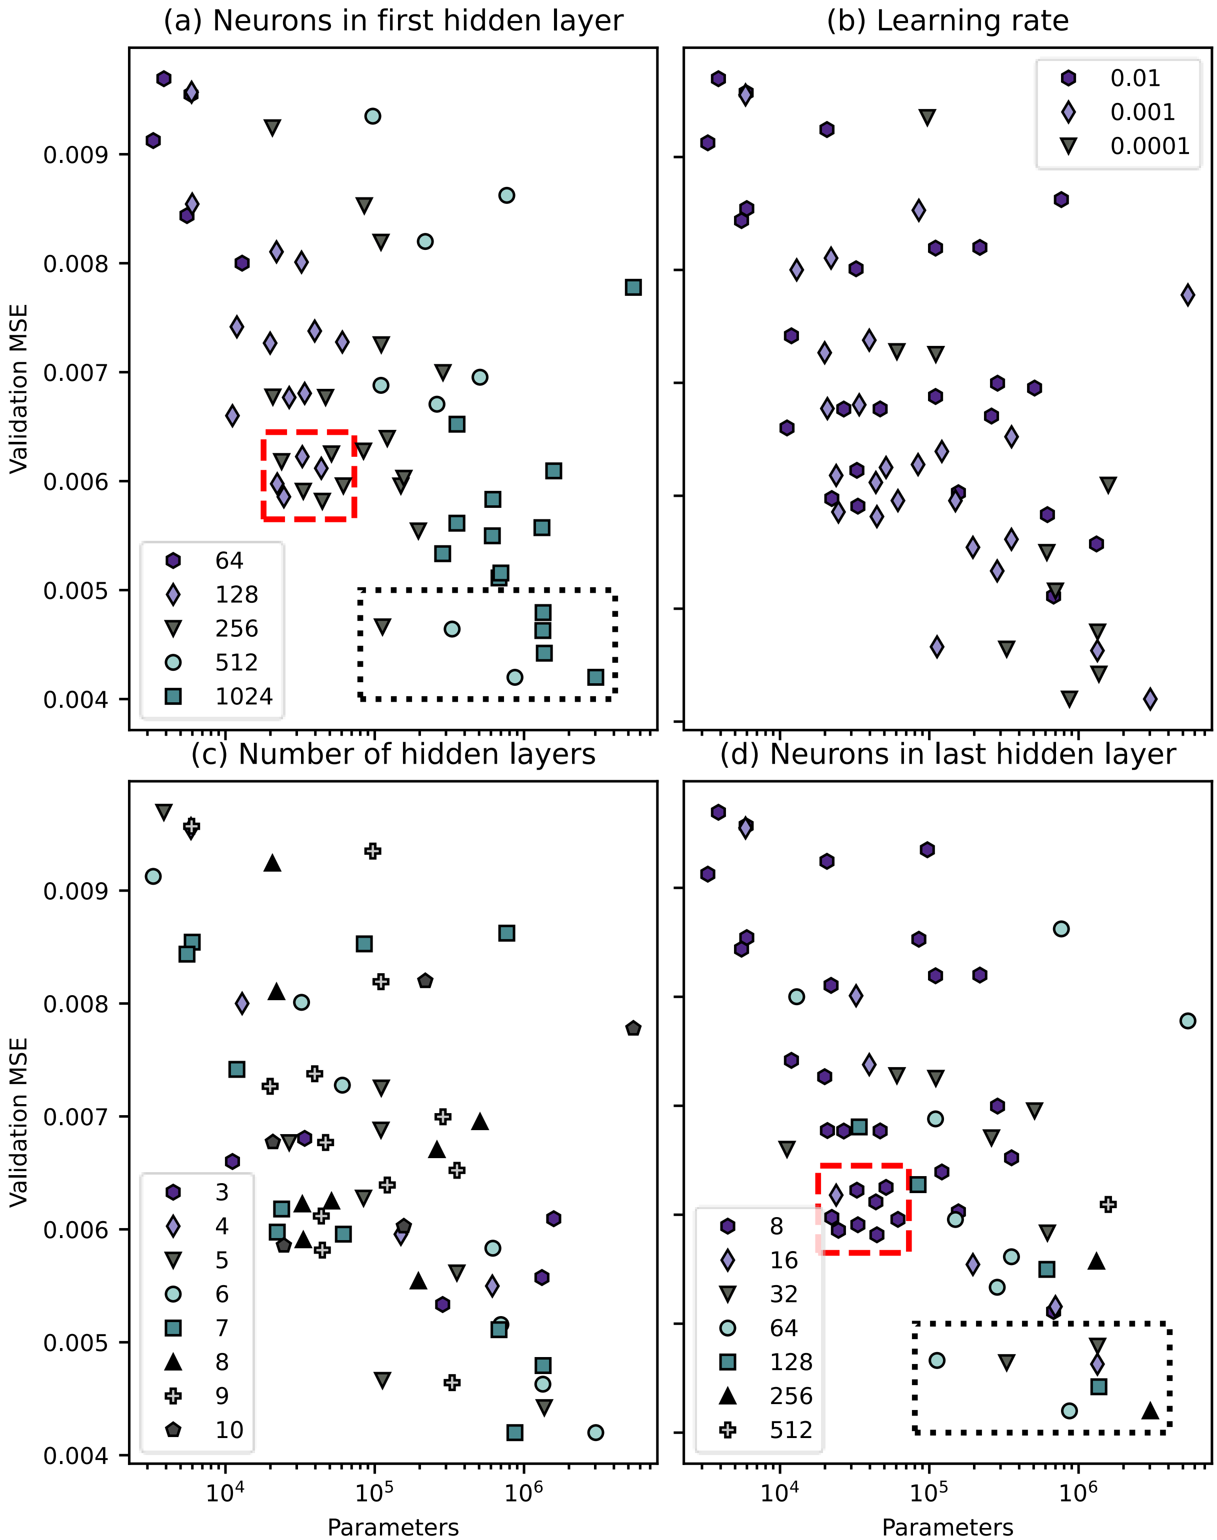

Figure 5.2: A zoomed in view of Figure 5.1 for the analysis of the best networks. The networks inside of the dashed region almost all end with 8 neurons, see (d). None of the better networks in the dotted region have 8 neurons.

5.3.3 Second Experiment

The next step for investigation is understanding the number of epochs that are required for each learning rate. To test this, nine neural networks with 6 dense hidden layers of 128 neurons were created, three for a learning rate of 0.01, 0.001, and 0.0001. Early stopping hides the number of epochs that each network gets trained on, so finding an appropriate number of epochs for each learning rate is the goal. The results from this test were surprising. The plot in Figure 5.3 shows how the mean squared error for the validation set and the mean square error for the training set varied with the number of epochs. The surprising part of the data is how few epochs, the model with a learning rate of 0.001 were able to complete. In a test of three models at each learning rate, all models with a learning rate of 0.0001 exceeded 1000 epochs, none of the other models made it to 1000. Surprisingly, the models with a learning rate of 0.01 made more epochs than the models with a learning rate of 0.001. These results indicate that there is a significant likelihood that models with a learning rate of 0.001 did not.

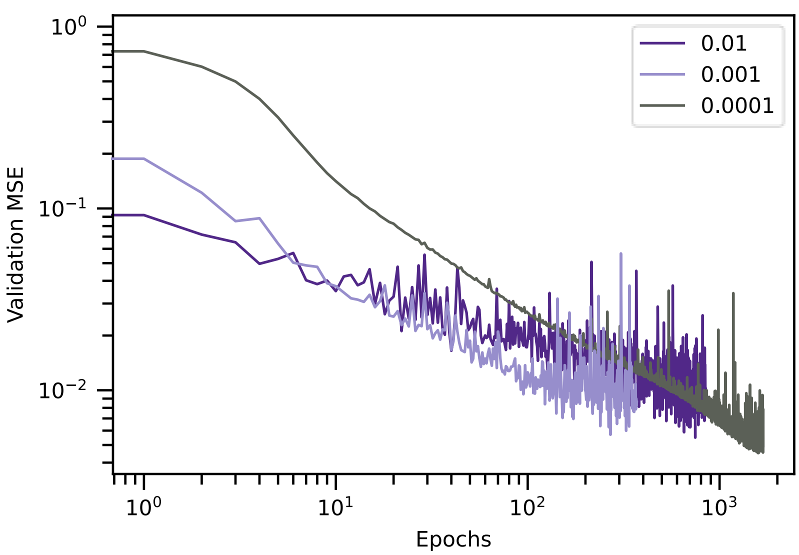

Figure 5.3: Figure comparing the validation MSE for the training of networks with three different learning rates. Notice that only the network with a learning rate of 0.0001 even made it to 1000 epochs.

From these results, it appears that there are two ways to approach the problem ahead, and maybe a requirement for more testing. The first approach would be to go long and slow with a learning rate of 0.0001. The other approach would be to train models with a stepped learning rate. Based on these results, it looks like a reasonable approach would be the following. Train model for 500 epochs with a learning rate of 10^{-3}. Train model for 100 epochs with a learning rate of 10^{-4}. Train model for 50 epochs with a learning rate of 10^{-5}. The stepped approach is definitely a faster approach, but there is another benefit.

An analogy for the loss function is some rugged terrain, and the goal is to get to the lowest point of that terrain. A large learning rate is like a helicopter, it allows you to survey many of the crevices of this terrain, but there are very few crevices that it can actually reach. A lower learning rate might be an ATV. It can go deeper in each crevice, but it cannot survey as many crevices. Maybe an even smaller learning rate is analogous to climbing down into a crevice on foot. The smaller vehicles allow greater access to the deep parts of the crevice, but it offers less ability to find a better crevice, especially if it gets stuck at the bottom. Going low and slow would be like trying to explore the whole terrain on foot. The stepped approach would be trying to use a helicopter, an ATV, and one's legs to get there instead.

A next step would be testing several models for each method to compare results. However, since the stepped method is much faster, it will be used primarily for future testing. A few promising models may be tested at a slow learning rate for many epochs.

5.3.4 Third Experiment

To try to develop some conclusion on the number of layers that is best, two sets of networks were created and tested. One set of networks was created with 64 neurons per layer while the other had 128 neurons per layer. All networks used the stepped learning rate approach. Figure

5.4 compares these results with those from Figure 5.2. Only two of the subplots are shown, because the others are not interesting. The results demonstrate that the stepped approach is effective as using it resulted in the best preforming methods with a fraction of the parameters. The stepped approach probably has too little training time for some of the larger networks, which may be the reason that the large 128 neuron per layer networks perform worse.

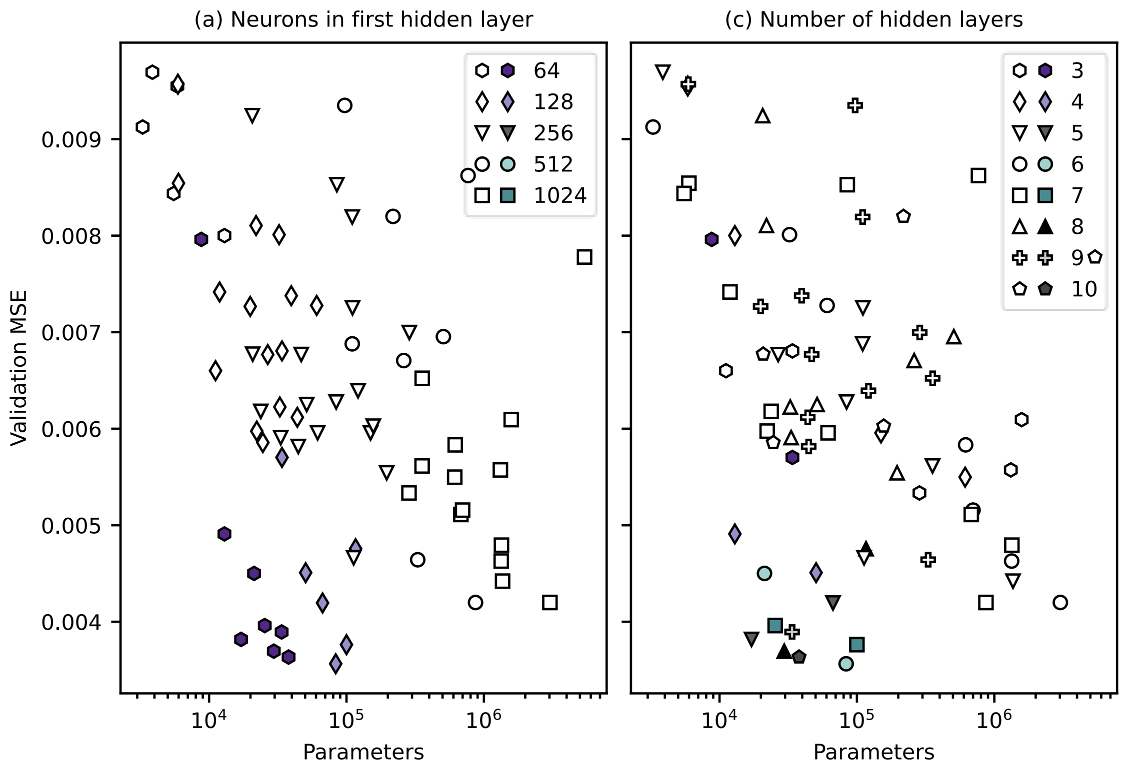

Figure 5.4: Subplots like Figure 5.2 for understanding the effect of layers for networks with the same number of neurons in each layer. All networks are trained used the step method described above.

5.3.5 Fourth Experiment

Based on previous analysis, another twenty networks were created with the goal of maximizing the performance of networks with less than 10^5 parameters. The initial layers of these networks had either 128 or 64 neurons. These networks once again had diminishing numbers of neurons per a network; however, 70% of the time the network will remain the same size. All networks were trained using the stepped learning rate approach. The fourth experiment's and third experiments are shown in color in Figure 5.5. The two best performing networks were chosen for final analysis.

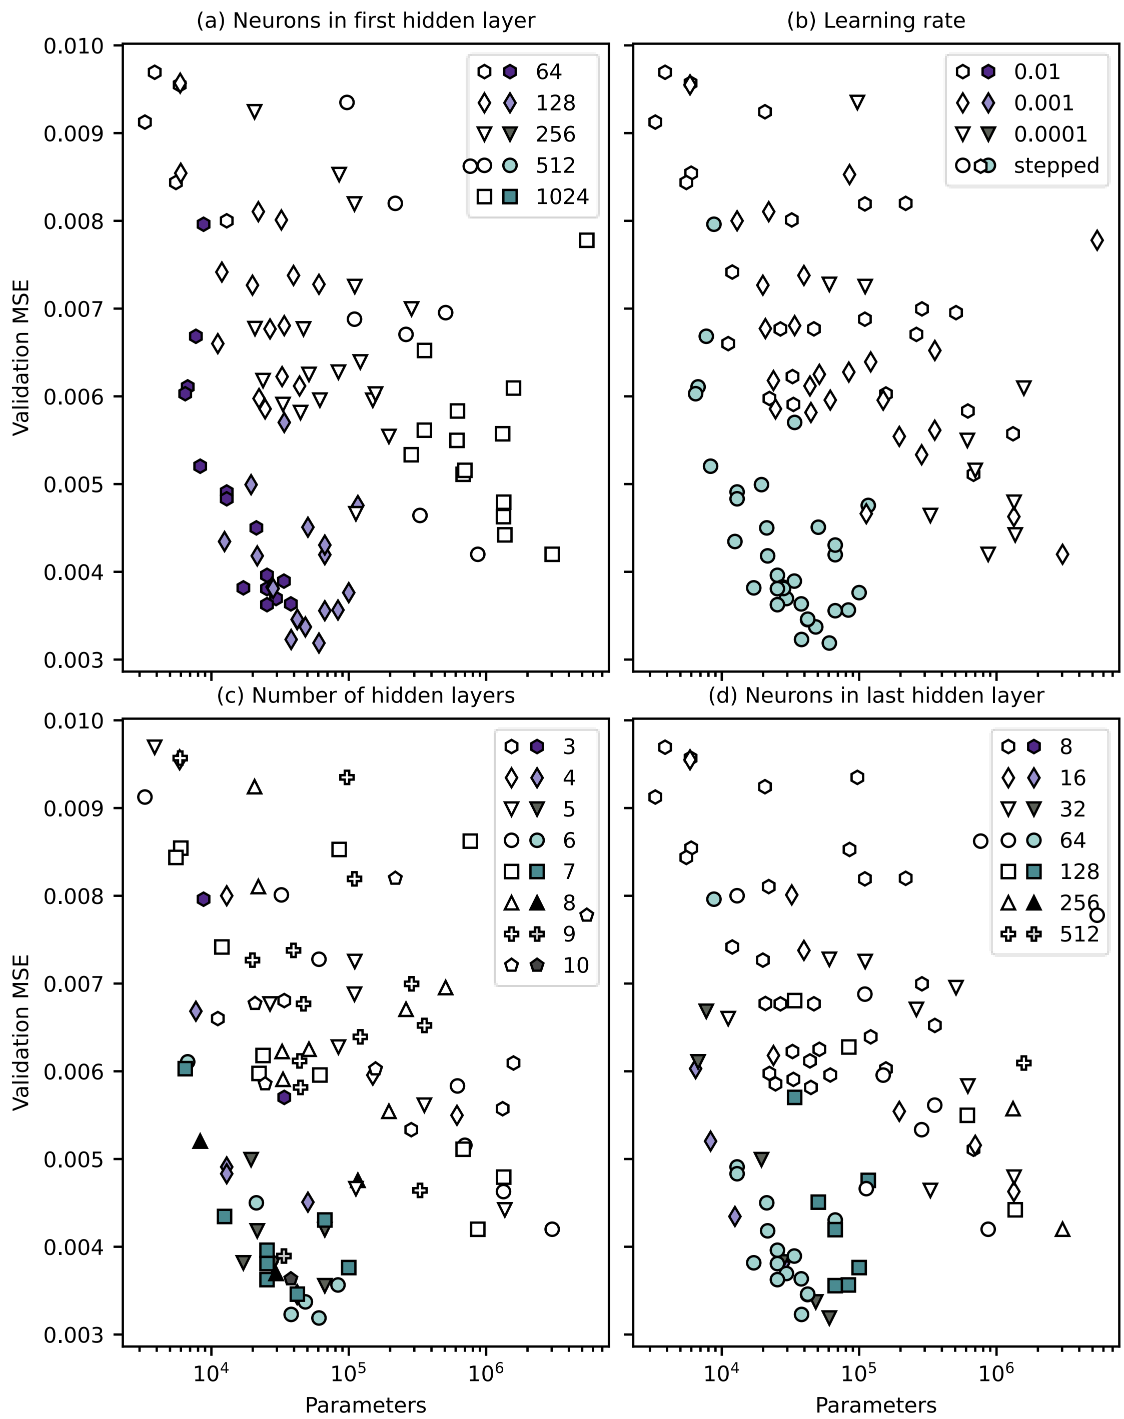

Figure 5.5: Comparing the results of networks from Experiment 3 and Experiment 4 to that of Experiment1.

5.3.6 Final Experiment

From previous experimentation two networks were chosen for more extensive training before obtaining final neural network performance results. Table 5.2 details the attributes of each of the networks. The more extensive included three repetitions of 1000 epochs at a learning rate of 10^{-3}, 500 epochs at a learning rate of 10^{-4}, and ended with 100 epochs at a learning rate of 10^{-5}. With the extensive training completed, the networks can be compared in a series of steps.

Table 5.2: Comparison of Architecture and total number of parameters for the two networks.

Model	Network 1	Network 2
Input Layer	5	5
Hidden Layer 1	128	128
Hidden Layer 2	128	128
Hidden Layer 3	64	128
Hidden Layer 4	64	128
Hidden Layer 5	64	64
Hidden Layer 6	64	32
Output Layer	1	1
Total Parameters	38 081	60 673

The final networks are compared with a series of tests. The first test is to get a prediction of the test set and the subsequent MSE value. These were actually collected for each network, but choices for which hyperparameters were made based of the validation set. Additionally, as a tool during training, all networks were used to predict the results for an entire stiffness-displacement curve that was not included in the training, validation, or testing sets. This is the second method of seeing the results. From these predictions, numerical integration is used to get the prediction for the force-displacement curve. These are plotted by taking the inverse transform of the predictions to return the data to real units. This was useful in understanding how bad, a MSE of 0.003 might be. This is easier than just plotting the predictions versus simulated values for the validation set because there is a familiar pattern to the data. Finally, the root mean squared error (RMSE) can be calculated for the final curve in terms of real units. The results for the stiffness standardized validation MSE, stiffness standardized test MSE, the stiffness RMSE in real units, and the force RMSE in real units is included in Table 5.3. Finally, the plot comparing the networks' predictions to the simulation force-displacement and stiffness displacement curves in Figure 5.6. Based on these results, Neural Network 1 is the best one concluding the hyperparameter study.

Table 5.3: Comparison of the two final networks using a variety of methods.

Model	Std. Stiff. Val. MSE	Std. Stiff. Test MSE	Real Stiff. RMSE (N/mm)	Real Force RMSE (N)
Network 1	0.003067	0.008687	230.5	0.5419
Network 2	0.002797	0.009375	600.9	1.0566

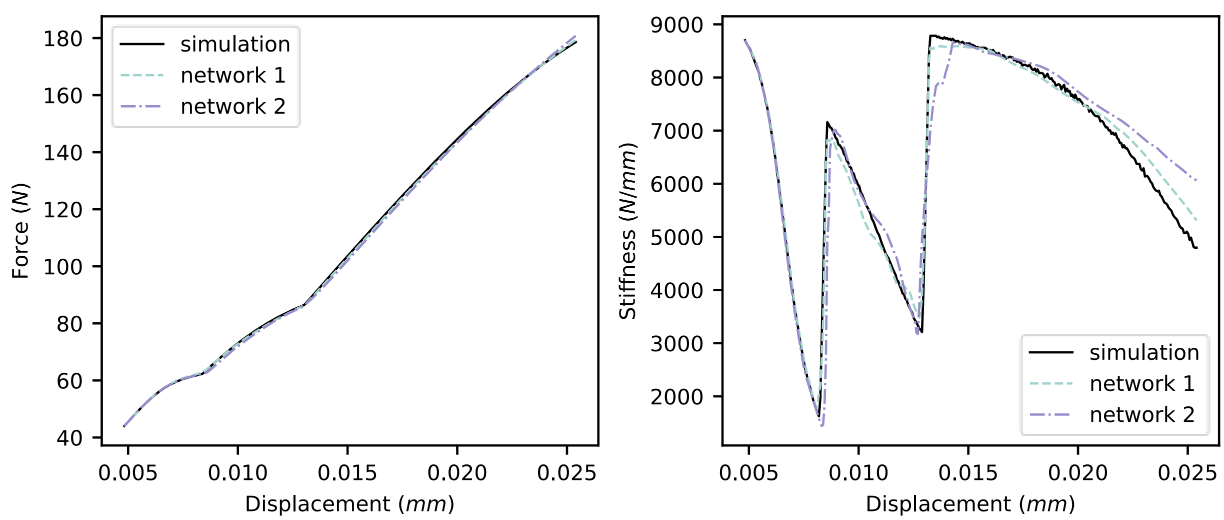

Figure 5.6: The two models' predicted force-displacement and stiffness-displacement curves are compared to the simulated case. This data was unique to all training, testing, and validation sets.

Chapter 6 - Conclusions and Recommendations for Future Work

6.1 Chapter Summaries

In Chapter 1, literature describing the problems caused by vibration and the methods of vibration control was introduced. The functions of piezoelectric materials in methods of active and semi-active vibration control were highlighted.

In Chapter 2, an idea was introduced to use buckling to magnify the ability of piezoelectric materials to influence the stiffness of slender beams. Different supporting scenarios, chosen by activation of piezoelectric materials, would allow for stiffness customizability of light weight metastructures. With customizable stiffness, there is an ability to control vibration.

In Chapter 3, finite element simulations were presented as a method to obtain information of the post buckling behavior of beams with different support scenarios. The nonlinear dynamic simulation demonstrated the ability to obtain different force deflection paths and ultimately and more significantly stiffness paths.

In Chapter 4, experimental methods were developed to provide corroboration for the simulated data. Additionally, these experiments needed to prove the feasibility of the conceptual idea. The first experiment primarily proved that MFCs were able to cause the beam to buckle in either direction as desired. The second experiment demonstrated significantly different force-displacement curves and stiffness-displacement curves depending on the location of the support. The third experiment was designed to support the FEA results. Primarily, it was able to demonstrate that load deflection curves produced by simulation were reasonable. Because of real deflections of the fixed support not experienced by the simulation, stiffness results were not comparable.

In Chapter 5, a neural network was created as a tool for positioning supports. Through the hyperparameter tuning process a neural network was developed that was able to predict results for a support scenario it had not seen before. For these predictions, it had a RMSE of 0.54 N for the force-displacement curve and a RMSE of 230 N/mm for the stiffness-displacement curve.

6.2 Conclusions

From this work, there are several important conclusions in the development of this device and of interest to the greater subjects of vibration suppression and mechanics of materials.

- The concept of positioning supports to influence the stiffness of a slender beam is demonstrated to be effective in simulation and experimentally. Multiple stiffness-displacement curves were created by solely varying the position of the intermediate supports.
- By influencing the direction of buckling with piezoelectric actuators, a support scenario can be chosen between two options, which allows for customization of stiffness.
- A nonlinear dynamic simulation was created to model the post-buckling behavior of a slender beam. Experiments supported the simulation, but due to experimental inadequacies, there was not close agreement.
- A neural network model was created for quicker design iteration of support positions.

6.3 Recommendations for Future Work

For continuing the work completed, the following recommendations are made.

- To continue this work, testing showing the device's influence on the dynamic behavior of a simple structure should be completed. For vibration suppression, the device must be able to demonstrate the ability to influence the natural frequency of a structure.

- To positively influence weight minimization, the beam structure's ability to carry load should be tested.
- Further experimental work will require a control system for deciding which active members to actuate and when they should be actuated.
- Another further development of this work will be through the use and further development of the simulation and neural network modeling tools used.

References

[1] D. J. Inman, *Engineering Vibration*, 4th ed. Upper Saddle River, NJ: Pearson, 2013.

[2] Q. Wang, X. Dong, L. Li, Q. Yang, and J. Ou, “Wind-induced vibration control of a constructing bridge tower with MRE variable stiffness tuned mass damper,” *Smart Materials and Structures*, vol. 29, no. 4, p. 045034, 2020.

[3] G. Caruso, “A critical analysis of electric shunt circuits employed in piezoelectric passive vibration damping,” *Smart materials and structures*, vol. 10, no. 5, p. 1059, 2001.

[4] N. Jalili and D. W. Knowles IV, “Structural vibration control using an active resonator absorber: modeling and control implementation,” *Smart Materials and structures*, vol. 13, no. 5, p. 998, 2004.

[5] Y. Bigdeli and D. Kim, “Damping effects of the passive control devices on structural vibration control: TMD, TLC and TLCD for varying total masses,” *KSCE Journal of Civil Engineering*, vol. 20, pp. 301–308, 2016.

[6] F. Jabbari and J. E. Bobrow, “Vibration suppression with resettable device,” *Journal of Engineering Mechanics*, vol. 128, no. 9, pp. 916–924, 2002.

[7] B. Yan, K. Wang, Z. Hu, C. Wu, and X. Zhang, “Shunt damping vibration control technology: A review,” *Applied Sciences*, vol. 7, no. 5, p. 494, 2017.

[8] V. Gupta, S. Adhikari, and B. Bhattacharya, “Exploring the dynamics of hourglass shaped lattice metastructures,” *Scientific reports*, vol. 10, no. 1, p. 20943, 2020.

[9] M. Xu *et al.*, “In-plane compression behavior of hybrid honeycomb metastructures: Theoretical and experimental studies,” *Aerospace Science and Technology*, vol. 106, p. 106081, 2020.

[10] K. H. Matlack, A. Bauhofer, S. Krödel, A. Palermo, and C. Daraio, “Composite 3D-printed metastructures for low-frequency and broadband vibration absorption,” *Proceedings of the National Academy of Sciences*, vol. 113, no. 30, pp. 8386–8390, 2016.

[11] A. Du Plessis *et al.*, “Properties and applications of additively manufactured metallic cellular materials: A review,” *Progress in Materials Science*, vol. 125, p. 100918, 2022.

[12] K. Saito, S. Pellegrino, and T. Nojima, “Manufacture of arbitrary cross-section composite honeycomb cores based on origami techniques,” *Journal of Mechanical Design*, vol. 136, no. 5, p. 051011, 2014.

[13] H. Fan, L. Yang, Y. Tian, and Z. Wang, “Design of metastructures with quasi-zero dynamic stiffness for vibration isolation,” *Composite Structures*, vol. 243, p. 112244, 2020.

[14] Y. Park and K. Kim, “Semi-active vibration control of space truss structures by friction damper for maximization of modal damping ratio,” *Journal of Sound and Vibration*, vol. 332, no. 20, pp. 4817–4828, 2013.

[15] A. Yousefi, S. Jolaiy, M. Lalegani Dezaki, A. Zolfagharian, A. Serjouei, and M. Bodaghi, “3D-Printed Soft and Hard Meta-Structures with Supreme Energy Absorption and Dissipation Capacities in Cyclic Loading Conditions,” *Advanced Engineering Materials*, vol. 25, no. 4, p. 2201189, 2023.

[16] H. Hassan, “The World is Nonlinear: Solving Buckling and Postbuckling Analysis of a Reinforced Panel,” 2019. https://blogs.solidworks.com/tech/2019/11/the-world-is-nonlinear-solving-buckling-and-postbuckling-analysis-of-a-reinforced-panel.html (accessed Jul. 08, 2023).

[17] J. Bergstra and Y. Bengio, “Random search for hyper-parameter optimization.,” *Journal of machine learning research*, vol. 13, no. 2, 2012.