

Deep Learning-Driven Modeling of Dynamic Acoustic Sensing in Biomimetic Soft Robotic Pinnæ

Sounak Chakrabarti

Dissertation submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Master of Science
in
Mechanical Engineering

Rolf Müller, Chair

Dylan P. Losey

Katie Rainey

Sep 1, 2024

Blacksburg, Virginia

Keywords: biosonar, deep learning, digital twin

Copyright 2024, Sounak Chakrabarti

Deep Learning-Driven Modeling of Dynamic Acoustic Sensing in Biomimetic Soft Robotic Pinnae

Sounak Chakrabarti

(ABSTRACT)

Bats possess remarkably sophisticated biosonar systems that seamlessly integrate the physical encoding of information through intricate ear motions with the neural extraction and processing of sensory information. While previous studies have endeavored to mimic the pinna (outer ear) dynamics of bats using fixed deformation patterns in biomimetic soft-robotic sonar heads, such physical approaches are inherently limited in their ability to comprehensively explore the vast actuation pattern space that may enable bats to adaptively sense across diverse environments and tasks. To overcome these limitations, this thesis presents the development of deep regression neural networks capable of predicting the beam pattern (acoustic radiation pattern) of a soft-robotic pinna as function of its actuator states. The pinna model geometry is derived from a tomographic scan of the right ear of the greater horseshoe bat (*Rhinolophus ferrumequinum*). Three virtual actuators are incorporated into this model to simulate a range of shape deformations. For each unique actuation pattern producing a distinct pinna shape conformation, the corresponding ultrasonic beam pattern is numerically estimated using a frequency-domain boundary element method (BEM) simulation, providing ground truth data. Two neural networks architectures, a multilayer perceptron (MLP) and a radial basis function network (RBFN) based on von Mises functions were evaluated for their ability to accurately reproduce these numerical beam pattern estimates as a function of spherical coordinates azimuth and elevation. Both networks demonstrate comparably low errors in replicating the beam pattern data. However, the MLP exhibits

significantly higher computational efficiency, reducing training time by 7.4 seconds and inference time by 0.7 seconds compared to the RBFN. The superior computational performance of deep neural network models in inferring biomimetic pinna beampatterns from actuator states enables an extensive exploration of the vast actuation pattern space to identify pinna actuation patterns optimally suited for specific biosonar sensing tasks. This simulation-based approach provides a powerful framework for elucidating the functional principles underlying the dynamic shape adaptations observed in bat biosonar systems.

Deep Learning-Driven Modeling of Dynamic Acoustic Sensing in Biomimetic Soft Robotic Pinnae

Sounak Chakrabarti

(GENERAL AUDIENCE ABSTRACT)

The aim is to understand how bats can dynamically change the shape of their outer ears (pinnae) to optimally detect sounds in different environments and for different tasks. Previous studies tried to mimic bat ear motions using fixed deformation patterns in robotic ear models, but this approach is limited. Instead this thesis uses deep learning neural networks to predict how changing the shape of a robotic bat pinna model affects its acoustic beam-pattern (how it radiates and receives sound). The pinna geometry is based on a 3D scan of a greater horseshoe bat ear, with three virtual "actuators" to deform the shape. For many different actuator patterns deforming the pinna, the resulting beam-pattern is calculated using computer simulations. Neural networks (multilayer perceptron and radial basis function network) are trained on this data to accurately predict the beam-pattern from the actuator states. The multilayer perceptron network is found to be significantly more computationally efficient for this task. This neural network based approach allows rapidly exploring the vast range of possible pinna actuations to identify optimal shapes for specific biosonar sensing tasks, shedding light on principles of dynamic ear shape control in bats.

Dedication

To my mom, dad, and my brother who sacrificed so much for me.

Acknowledgments

I extend my appreciation to Dr. Rolf Müller for fostering my interest in bioinspired robotics. Academic achievements wouldn't have been possible without his guidance and financial support. I would also like to thank Dr. Losey who introduced me to the academic world of robotics and for inspiring me to read further literature. I am grateful to Dr. Rainey for introducing me to NIWC PAC where I had the opportunity to intern and learn about swarm robots. I'm thankful to Michael Goldsworthy, Adam Hinson, Joseph Sutlive, and Ibrahim M. Eshera, lab colleagues who helped me along the way of completing the masters and made this all possible. Finally I'm deeply appreciative of the steadfast support from my parents, Bimal K. Chakrabarti and Mousumi Paulchakrabarti, and my brother Inesh Chakrabarti.

Contents

List of Figures	ix
List of Tables	xii
1 Introduction	1
1.1 Bat Biosonar	1
1.2 Biomimetic Soft Robot	3
1.3 Deep Learning Methods	5
2 Deep Learning-Driven Modeling	9
2.1 Executive Summary	9
2.2 Introduction	10
2.3 Methods	13
2.4 Results	20
2.5 Discussion	22
3 General Discussion and Conclusions	34
3.1 Research Accomplishments	35
3.2 Major Findings	36

3.3 Suggestions for future work	36
Bibliography	38

List of Figures

2.1	Workflow for the numerical beampattern predictions: (a) 3D tomographic scan of a greater horseshoe bat’s right ear used to create the pinna-shape model; (b) Digital pinna model rigged with actuators used to generate the different shape conformations; (c) Pinna surface mesh used for the BEM simulations; (d) Numerical beampattern prediction from the BEM model for a frequency of 40 kHz.	16
2.2	Orthographic view of a spherical mesh used to sample beampattern gains in the far field (at a radius of 3 m), sampling elements of partition generated by the HEALPix discretization [19].	25
2.3	Comparison of the deep neural network architectures used to approximate the beampattern functions: (a) multi-layer perceptron, (b) deep radial basis function neural network with von Mises basis functions in the first activation hidden layer that represent the three parameters of the respective basis function in each node. For both neural network architectures, the inputs were azimuth and elevation and the output was beam gain (on either a linear or on a dB scale).	26
2.4	Comprehensive MLP surrogate model of the pinna for used for beam-pattern prediction: The neural network’s input takes the actuation state of three soft actuators (T_1, T_2, T_3) alongside the spherical coordinates (θ, ϕ) to output the respective acoustic beampattern gains.	27

2.5	Impact of acoustic signal frequency and pinna shape conformation on the numerical (BEM) beampattern estimates: a) upright (dashed line) versus bent pinna (solid line) at 30 kHz; b) different frequencies, 30 kHz (dashed line) versus 40 kHz (solid line), for the upright pinna shape conformation. The visualizations of the beampatterns on the sphere are shown below with the cutting planes for the respective polar plot indicated.	28
2.6	Learning curve of the comprehensive MLP trained on the BEM simulation and actuator inputs. Average root mean square (RMSE, in dB) for training (solid black line) and validation (dashed gray line) over 15 trials conducted at frequencies of 30, 35, and 40 kHz. Standard deviations are denoted by the thinner dashed lines in the respective gray level.	29
2.7	Example beampattern predictions from the MLP as a function of pinna shape conformation and acoustic frequency: In this set of example shape configurations, the middle actuator transitions from moderately bent (2) to fully bent (10) in three intermediate steps. For each shape configuration, MLP beampattern predictions for 30, 35, and 40 kHz are shown.	30
2.8	Comparison of numerical and deep-learning-based beampattern predictions: Deep-learning-based estimates (left-hand column: acoustic MLP, right-hand column: acoustic VMBFNN) versus numerical predictions (BEM, middle column) for an example pinna deformation due to actuator states 0,6,0 and acoustic frequencies of 30, 35, and 40 kHz.	31

2.9	Comparison of example numerical (BEM) and deep-learning based comprehensive MLP predictions: cross-sectional plots of numerical (BEM, dotted line) and comprehensive MLP (dotted-dashed line) estimates. The examples shown are for (a) an upright shape conformation and (b) a bent shape conformation. In both cases, the acoustic frequency is 30kHz. The insets show the orientations of the cutting planes used to create the cross sections relative to the respective beampatterns.	32
2.10	Training and inference times for the acoustic neural network architectures evaluated: a) Distribution of training times for the acoustic MLP (dark gray) and the acoustic deep von Mises basis function neural network (VMBFNN, light gray); b) Distribution of inference times. Each box-and-whisker plot represents 15 data points from five trials for each of the three frequency (30, 35, and 40 kHz) that were based on 1,331 deformation models that were split into 60% for training, 20% for validation, and 20% for testing.	33

List of Tables

List of Abbreviations

σ The total mass of angels per unit area

NLP Natural Language Processing

σ is the eighteenth letter of the Greek alphabet, and carries the 's' sound. In the system of Greek numerals, it has a value of 200.

Chapter 1

Introduction

1.1 Bat Biosonar

Sonar, which stands for Sound Navigation and Ranging, is the use of acoustic waves to navigate and sense a specific environment. There are two main types of sonar: active and passive. Active sonar involves producing a sound with the intent of reaching or sensing a target, while passive sonar is used for listening to sounds produced by the target. Sonar is employed in many contexts, including biological systems, where it is referred to as biosonar. For example, gleaning bats use passive sonar to localize prey. [76, 91] Biosonar enables organisms to navigate their environment, allowing them to find, learn, and return to places of interest. [58].

Of the approximately 1,400 known bat species, around 1,200 [96] employ their biosonar systems for navigation, obstacle avoidance, and spatial orientation in diverse environments, enhancing their hunting and foraging abilities in dense vegetation [77]. The greater horseshoe bat (*Rhinolophus ferrumequinum*) exemplifies sophisticated biosonar usage, utilizing a combination of CF (constant frequency) segments and FM (frequency modulated) sweeps while searching for food. This species is exceptionally maneuverable when flying between branches, environmental clutter, and hanging in trees [26]. In addition, they are known to briefly stop in front of foliage and even do hovering flight while catching prey [26]. Similarly, insectivorous big brown bats (*Myotis fuscus*) exhibit adaptive echolocation strategies in complex

environments. When approaching prey or navigating cluttered spaces, they increase the rate of sound production while decreasing pulse length, effectively expanding their acoustic field of view. [68, 76] They can even adapt their sonar calls according to their own movements by timing emissions according to their wingbeat rates. [11] While their echolocative behaviors are well understood, the sensory basis behind these capabilities is not yet well understood.

Technical sonars are typically 100 to 1000 times larger than the wavelengths they operate on [13, 88]. In contrast, the bat biosonar system is very compact, comprising just three main components: an emitter (larynx and nostrils) and two ears functioning as receivers, collectively operating at frequencies up to 80 kHz. From a biological perspective this is remarkable, especially considering that the total brain mass of many bat species is less than 1 g [10, 39]. One postulate is that the geometry of the outer ear, known as the pinna, plays a critical role in dynamically sensing their environments, especially for (rhinolophids and hipposiderids) [18, 64]. During biosonar activity, the pinna assists with localization, using time delays and binaural cues for horizontal angle localization, and monaural cues for vertical localization [76]. The linearity in encoding signals with the pinna was proven when the removal of a flap in the ear of the Brown Long-Eared Bat (*Plecotus auritus*) drastically changed the sensitivity of the pinna, known as the gain, which is a function of both direction and frequency [51]. Moreover, immobilizing the pinna of the greater horseshoe bat caused a drastic decrease in performance of vertical localization [43]

Early studies demonstrate that the greater horseshoe bat commits to rapidly alternating ear movements coupled with short orientation pulses near close objects [18]. More recent findings have uncovered the potential involvement of nonlinear mechanisms in the encoding process. For instance, the surface velocities of the greater horseshoe bat's (*Rhinolophus ferrumequinum*) pinnae have been shown to create Doppler shifts that encode echoes incident on the ear canal [98]. The goal of our study is to describe the static changes in the pinna

using time invariant characteristics (beam patterns) in order to understand how the bat dynamically encodes signals upon reception.

1.2 Biomimetic Soft Robot

Pinna motions made by the greater horseshoe bat are of high interest due to the high variability in their dynamics including both rigid rotations as well as nonlinear deformations [100]. In particular, rigid rotations are known to have a large amount of variability [66], and are instrumental in reducing background clutter while increasing target localization abilities. Studies have shown that the pinna motions in horseshoe bats are non-rigid as well, resulting in noticeable geometrical changes to the pinna shape, and affecting beam patterns on reception. Importantly, the extent of these non-rigid shape deformations has been found to be comparable in magnitude to the ultrasonic wavelengths used by these bats [14]. This amount of variability is to be expected in bats based on qualitative observations. From a quantitative standpoint it's tantamount that the bat pinna contains a large number of muscles around 20 [73].

Given the number of degrees of freedom in the musculature of the original pinna, it is difficult to repeat all of the flexible behavioral patterns of these bats in real time [59]. In addition, there are physical limitations to how many pinna motions that can be extracted from this species of bats in a lab environment to understand how they echolocate. Moreover, obtaining substantial quantitative results from natural sensory behaviors in complex natural environments proves exceedingly difficult. Further complicating matters, bats might be generating additional types of deformations in their natural habitats. Studying such deformations in the wild presents its own set of challenges and experiments with live animals have to be limited in scope due to ethical considerations.

To simplify the search for pinna deformations, [52] identifies four major characterizations describing the impact of the pinna dynamics on encoding of sensory information: emission amplitude, direction angles, and frequency. Unfortunately, acquiring accurate data on these attributes from live bats is particularly challenging. Given these limitations, the need for more quantitative data necessitates a biomimimetic approach, the study of nature as a model for understanding special physical phenomena [9].

In this case, the bat’s biosonar system employs a cyclical perception and action model [82]. The process is initiated by transmitting sound pulses outward from the larynx or nostrils into the surrounding environment. These transmitted pulses then encounter objects, generating echoes that propagate back towards the organism. The outer ears capture these returning echoes, transducing them into neural signals within the inner ear structures. Subsequently, the neural system then extracts all pertinent information from the encoded signals. Based on this extracted information the bat dynamically modulates various components involved in sonar operation by issuing control signals to the appropriate motor pathways. Hence we have this closed-loop architecture that incorporates feedback at various stages allowing the opportunity for optimization of the bat’s sonar capabilities in response to echoes received and the data gleaned from them [48]. Within this loop, the bat pinna can be considered as a complex robotic sensor, where the function of the actuation is sensing the external environment. Given the intricate and flexible nature of this biological sensor, replicating its structure and function for research purposes presents a unique challenge.

To address this challenge, soft robotics offers a promising solution. This field, which focuses on the usage of flexible ‘soft’ materials for modelling real life objects [89], is necessary for replicating the bat’s pinna. Soft robotic models allow researchers to understand the dynamic effects without harming live bats, providing an ethical and practical approach to studying these complex structures [3, 82].

1.3 Deep Learning Methods

Modelling the controls of a redundant continuum soft robot with many degrees of freedom, and retaining the function of acoustical sensing is no easy task. Leveraging deep reinforcement learning has emerged as the cutting-edge methodology, enabling robots to iteratively explore and uncover optimal behavioral patterns through a process of trial and error interactions within their operational environments. [33] Obtaining realistic experience through a physical system is arduous, costly, and challenging to replicate in real time settings [33]. This issue is compounded by the large dimensionality of the state and action spaces inherent to such systems. For example, the greater horseshoe bat’s noseleaf involves around 14 muscles [92], while each of its pinnae (outer ears) requires approximately 20 muscles to function [73]. As a result, replicating the intricate complexity of this biological system, with its two pinnae, in a robotic platform could demand the integration of nearly 50 actuators. A model based approach with a reduced number of controllable parameters is necessary to address the overall control problem [33, 86, 93].

Otherwise, defining the search space containing all possible configurations for controlling a robot with many degrees of freedom is a difficult task. As an illustrative example, learning hand-eye coordination for grasping tasks with a robotic manipulator featuring seven degrees of freedom and a two-finger gripper necessitated executing over 800,000 grasping motions over a period of two months, employing six to 14 manipulators operating in parallel [37]. Similarly training a three fingered Jaco robot arm on a pixel-to-action learning task: reaching to a randomly placed target from a random start location would take 53 days on a real robot with continuous training for 24 hours a day [86]. Meanwhile, the MuJoCo simulator employing an Actor Critic method [42], has demonstrated higher efficiency in training the Jaco robot arm. This approach enables the arm to perform effectively just after 24 hours of training on a CPU compute cluster.

These advancements highlight the necessity of applying a deep-learning method to the problem of controlling the biomimetic bat ear [53]. The challenge lies in efficiently testing a vast number of dynamic shape conformation sequences. While physical simulation software could theoretically handle this task, the sheer number of possible conformations makes computational methods impractical due to time and resource constraints. To address this issue, a novel approach is required—one that can instantaneously evaluate the acoustic characteristics of each shape conformation.

In response to this challenge, this research investigates the application of deep learning techniques to model the full acoustic characteristics of a biomimetic pinna as a linear time-invariant system. Specifically, the aim is to characterize the receiver’s directional sensitivity, commonly referred to as its acoustic gain, as a function of both the far-field angle of incidence and the frequency of the incident sound waves. This characterization effectively captures the beam patterns associated with the biomimetic pinna’s reception capabilities.

Traditional modeling approaches that establish connections between the actuation inputs of a biomimetic pinna and the resulting acoustic characteristics require a two step process: (i) forecasting the pinna’s shape based on actuator inputs, and (ii) predicting the acoustic characteristics stemming from that specific shape [14]. As mentioned earlier, both steps can be computationally expensive to produce the desired results. While numerical methods such as finite elements [48] or boundary elements [40] can be optimized for efficiency, enabling higher angular resolution and the ability to evaluate the pinna’s acoustic response across a broad range of frequencies for a larger dataset of pinna shapes, the overall computational process can still demand up to several hours of computation time on a PC [38]. This computational burden is particularly pronounced when evaluating the Helmholtz equation at low to moderately large wavenumber ranges between m^{-1} to m^{-1} which correspond to CF component in the biosonar pulses of horseshoe bats [78, 87]. To address these computational

challenges, researchers have turned to alternative approaches. DL neural network-based regression methods provide a data-driven yet meshless avenue to capture the physics of traditional numerical methods [30, 95]. This fusion of numerical methods and machine learning offers unique advantages. Utilizing the numerical method as a data augmentation step for training the machine learning system, the actuator states can be linked to beam patterns directly addressing any observational biases [30].

Although training a deep neural network (DNN) can be computationally expensive and resource-intensive, once the training is complete, the process of inference or making predictions using the trained model is significantly more efficient and cost-effective. This characteristic makes DNNs well-suited for applications involving pinna motions, where the model needs to be trained only once, but the inference process must be repeated numerous times to generate predictions or classifications for new input data. Additionally, any inductive biases such as the use of boundary conditions or any other prior knowledge is also explored by designing a specialized neural network architecture using radial basis functions [2, 30].

In this study, two approaches have been tested, using a black box deep neural network to fit the data, and utilizing prior knowledge of the shape of the beampattern to train a specially curated network. The black box methodology uses multi-layer feed-forward networks, which are assumed to be well universal function approximators utilizing a combination of weights and activation layers which can easily find the relationship between the actuator states of the pinna and the associated beam patterns through back propagation [22, 69, 84]. The traditional RBFN assumes that the training data at hand is nonlinearly separable, and makes use of prior knowledge of the data in the form of weighted nonlinear basis functions such as the distances from center mean values of multiple gaussian distributions to approximate a smooth function [2]. Additionally, one convenient aspect is that when fitting data to the RBFN, the number of degrees of freedom is required, which is known. The maximum

number of basis functions is constrained to the number of data points which would lead to overfitting.

During the 1990s radial basis function neural networks (RBFNs) gained popularity as an alternative to deep neural networks (DNNs) due to their superior performance and lower network complexity compared to DNNs [81]. Multilayer perceptrons (MLPs), had higher sensitivity to initial conditions due to their numerous parameters, and often struggled to converge to global minimums, instead becoming trapped in undesirable local minima [20, 63]. However, powerful modern machine learning frameworks have made training black-box models more convenient compared to the effort of injecting prior knowledge through basis functions in RBFNs [60, 61].

To investigate whether a RBF network, that uses basis functions that are a match for the lobe structure of beampatterns, can improve the quality and efficiency of beampattern estimates, a deep RBFN using von Mises basis function [25] is implemented and tested. Unlike most RBF networks, which are symmetric around a 2D point, the von Mises basis function is symmetrical about a point in 3D spherical coordinates. The expression is given below. The input parameters are azimuth and elevation (θ, ϕ) while the centroids in azimuth and elevation are (α, β) . There's a concentration parameter κ which is also known as a shape parameter, and the larger the value is the narrower the function width is after transformation by the expansive function e .

$$VM(\theta, \phi, \alpha, \beta, \kappa) = e^{\kappa[\sin \theta \sin \beta \cos(\theta - \alpha) + \cos(\theta) \cos(\beta)]}$$

To assess whether the injection of this prior knowledge into a deep neural network could enhance inference and training performance, a comparative study is conducted. This study contrasts the performance of the deep von Mises basis function network against a standard 'vanilla' black-box multi-layer network.

Chapter 2

Deep Learning-Driven Modeling

2.1 Executive Summary

Biological function often depends on complex mechanisms of a dynamic, time-variant nature. An example are certain bat species (horseshoe bats - Rhinolophidae) that use intricate pinna musculatures to execute a variety of pinna deformations. While prior work has indicated a potential significance of these motions for sensory information encoding, it remains unclear how the complex time-variant pinna geometries could be controlled to enhance sensory performance. To address this issue, the present work has investigated deep neural network models as digital twins for biomimetic pinnae. The networks were trained to predict the acoustic impacts of the deformed pinna geometries. A total of three network architectures have been evaluated for this purpose using physical numerical simulations (boundary element method) as ground truth. The networks predicted the acoustic beampattern function from pinna shape or even directly from the states of actuators that were used to deform the pinna shapes in simulation. Inserting prior knowledge in the form of beam-shaped basis functions did not improve network performance. The ability of the networks to produce beampattern predictions with low computational effort (in about three milliseconds each) should lend itself readily to supporting learning methods such as deep reinforcement learning that require many such functional evaluations.

2.2 Introduction

Autonomous navigation in complex natural environments remains an unresolved technical challenge at present [23, 41, 56]. The current technological shortcomings apply to mobility as well as to the capability of sensing important features of the environment. However, there are a number of potential biological model systems provided by highly mobile animals that thrive in densely vegetated habitats and could hence offer new approaches to this problem. For example, there are a few groups of bat species that stand out for their ability to navigate and hunt their prey in dense vegetation [11, 26, 68, 76, 77]. While the sensory basis for many of the capabilities of these bats is not yet understood, it is known that bats with sophisticated biosonar systems can accomplish such tasks based on biosonar echoes as their sole source of information on their environments [12, 17]. In at least two bat groups with particularly sophisticated biosonar systems – horseshoe bats (Rhinolophidae) and Old World leaf-nosed bats roundleaf (Hipposideridae), a number of unique features that have been interpreted as adaptations for detecting prey in dense vegetation [54, 68, 75]. In addition, a conspicuous pinna dynamics appears to be an integral part of biosonar function. Animals from these bat species frequently carry out rigid pinna motions as well as pinna deformations while the biosonar echoes are received [14, 18, 64, 65, 100]. The frequency of occurrence of these motions along with the presence of an unusually differentiated pinna musculature [73, 74] suggests that the hypothesis that the pinna dynamics play a critical role for the biosonar performance. To support this hypothesis, the pinna motions have been shown to encode additional information through linear [51] as well as nonlinear mechanisms [97].

However, a detailed analysis of the function of the pinna motions *in vivo* is complicated by the complexity of the involved geometries and the associated acoustic diffraction processes [48, 50]. In addition, a large amount of variability has been demonstrated in rigid rotations [66] and qualitative observations suggest a high variability for the deformations as well [99, 100].

This raises the question whether and how this variability could be used to support different sensing tasks and conditions under which they have to be performed.

To reduce the complexity of studying this problem, a biomimetic approach based on robotic reproductions of the periphery of bat biosonar has been pursued [3, 82]. However, this line of inquiry has also proven to pose demanding challenges: Each bat pinna can be considered a complex “robot” where the function of the actuation is to enhance sensory function. Replicating this function requires designing and controlling a redundant continuum soft robot with many degrees of freedom which poses a major challenge to the state of the art in (soft) robotics and controls [36, 93]. The most prominent current approach to address the control part of this problem is deep reinforcement learning (RL) [28, 33, 83]

However, the high dimensionality of the state and action spaces of such a biomimetic soft robot presents a significant challenge: A greater horseshoe bat has about 20 muscles on each pinna [73]. Hence, a robot emulating the full complexity of such a biological model system (with two pinnae) could require about 40 soft-robotic actuators in total. Unless a model can be formulated that reduces the controls problem to a much smaller number of parameters [33, 86, 93], the enormous size of the search space spanned by all potential solutions to controlling a robot with that many degrees of freedom poses a serious problem. For example, learning hand-eye coordination for grasping with a robotic manipulator with seven degrees of freedom (including a two-finger gripper) required more than 800,000 grasping motions that were executed over a period of two months by running six to 14 manipulators in parallel [37]. Based on this and similar examples, applying a deep-learning method to the problem of controlling a biomimetic bat ear [53] could be expected to require testing a prohibitively large number dynamic shape conformation sequences. Hence, the most promising option for successfully completing a search for successful actuation patterns are simulation approaches. Nevertheless, even in simulation, the extremely large numbers of possible conformation se-

quences to be evaluated would still require a highly efficient computational to evaluate the functional (i.e., acoustic) characteristics of each shape conformation.

To address this critical issue on the road towards controlling a biomimetic pinna robot for an enhanced sensing performance, we have evaluated a deep-learning approach aimed at a complete linear acoustic characterization of a biomimetic pinna. Specifically, we have aimed at predicting the acoustic gain as a function of direction and frequency, i.e., the beampattern of a given pinna shape under actuation. Conventional modelling techniques that could be used to link the actuation inputs into a biomimetic pinna to the resulting acoustic characteristics are based on numerical approximations of the respective physics. Hence, such approaches would require two steps for (i) predicting the pinna shape from the actuator inputs and (ii) predicting the acoustic characteristics from the pinna shape [14]. Both of these steps can be computationally expensive:

Predicting the acoustic characteristics of shape via traditional numerical methods such as finite element [48] or boundary elements [40] could take many minutes to a few hours on a PC [38] for the low to moderately large wavenumber ranges (from 1.400 m^{-1} to 1.500 m^{-1}) that correspond to the frequencies of the CF component in the biosonar pulses of horseshoe bats [78, 87].

To radically cut down on this computational overhead, the current work has explored deep-learning-based methods that directly capture effects of the physical mechanisms [30, 95] that link actuator states to beampatterns. While training a powerful deep neural networks (DNN) is also computationally expensive, using the network to make predictions (inference stage) is much cheaper. The disparity between the cost of learning and inference is a good match for the purpose of evaluating a large number of pinna motions since the DNN has been trained only once for this task, whereas the inference needs to be repeated many times.

In order to assess which DNN approach is the most suitable for the task at hand, two different networks have been tested: a black box DNN and a DNN that included prior knowledge about shape of the beampatterns to be predicted. The black box approach used a multi-layer feed-forward network (i.e., a multi-layer perceptron) [84] as a universal function approximator [7, 22]. The network that was used for incorporating prior knowledge on the beampattern shapes was a radial basis function network (RBFN, [2]). In the 1990s, RBFNs were popular alternatives to black-box DNNs because they offered better performance and required less network complexity than DNNs [81]. In addition, MLPs with their many parameters and sensitivity to initial conditions tended to get trapped in undesirable local minima, hence failing to converge and reach the global minimum of the problem [20, 63]. However, since then powerful modern machine learning frameworks have developed that made training black-box models more convenient and reliable [57] and hence have made the effort of injecting prior knowledge through basis functions in an RBFN less worthwhile. To investigate whether an RBFN that uses basis function that are a match for the lobe structure of beampatterns can improve the quality and efficiency of beampattern estimates, a deep RBFN using von Mises basis function [25] was implemented and tested.

By finding a good, workable solution to the problem of estimating the beampattern properties of a simulated pinna as a function of actuator inputs, the present work could lay the foundation for training complex behaviors in biomimetic soft-robots designed to mimic the superior sensing skills of bats.

2.3 Methods

The current research was based on a digital shape representation of a horseshoe bat pinna that had been previously generated by virtue of a micro-computer tomographic scan of a

male greater horseshoe bat (*Rhinolophus ferrumequinum*, right ear). The pinna model was represented by a polygonal surface mesh that contained approximately 20,000 triangular elements. To match the dimensions of a biomimetic bat robot that has been based on the same pinna model [82], the mesh was resized to about twice the original size of the bat’s pinna resulting in a total model height of 5.8 cm [97] (from the base to the tip). The frequencies simulated were scaled accordingly. The mesh was modified to add an artificial ear canal (i.e., cylindrical pipe) that interfaced with the opening of the ear canal of the pinna model. The diameter of the canal (4.5 mm) was smaller than the shortest studied wavelength (8.6 mm at 40 kHz) and its length (length 16 mm) of its diameter was about 1.5 times the largest studied wavelength (11.4 mm at 30 kHz). Hence, it can be expected that the wavefront exiting ear canal will be planar, regardless of the type of sound source used [47].

The pinna model was deformed digitally using a 3-D animation technique (Blender, Blender Foundation, Amsterdam, Netherlands): To create the deformations, the mesh model was rigged with an animation skeleton that mimicked a soft-robotic biomimetic ear with three soft bending actuators spreading out from the base on the backside of the pinna (Figure 2.1, [82])

A bending trajectory was designed for each of the actuators to deform the pinna from an upright state to a fully bent state in a way that was qualitatively similar to the open-close deformation patterns that were visible in slowed-down high-speed video footage of real-life pinna deformations in a greater horseshoe bat [98]. For each actuator, 11 conformation states were selected by virtue of an even spacing along each actuator’s bending trajectory. By combining different bending states across the three individual actuators, a total of 1,331 different pinna conformation states (i.e., all 11^3 possible combinations) were generated.

The meshes representing the different pinna shape conformations were used to define an acoustic boundary element model (COMSOL Multiphysics v5.4 Acoustics Module, COMSOL AB, Stockholm, Sweden) to solve the Helmholtz equation [6, 44, 62]. To ensure a

sufficient mesh resolution for geometric features of potential acoustic relevance and increase computational efficiency at the same time, a physics-controlled remeshing strategy that adapts to the geometry and the physical process to be simulated ('physics-controlled free triangular remeshing', COMSOL Multiphysics v5.4 Acoustics Module, COMSOL AB, Stockholm Sweden, [5], maximum element growth rate 1.5 mm^{-1} , curvature factor 0.5, resolution parameter 0.6) was applied to deformed meshes.

To simulate free-field sound propagation away from the pinna, the pinna surface mesh was surrounded by an unbounded air space ('air material', COMSOL Multiphysics v5.4 Acoustics Module, COMSOL AB, Stockholm Sweden, [5], density 1.204 kg m^{-3} , bulk modulus 141.814 kPa , temperature 293.15 K , resulting in a sound speed of 343.2 m s^{-1}). The surface of the pinna was modeled as a sound-hard boundary, i.e., all incident sound energy was reflected; the initial pressure values on the pinna's surface were set to the ambient atmospheric pressure (1.013 hPa). For computationally efficient modeling of reflection-free sound propagation, a mapping-based approach to infinite elements was utilized where a mapping function transformed the infinite domain into a finite computational one, resulting in an outgoing-wave boundary condition [101]

Acoustic reciprocity [62], a fundamental property of wave propagation, allows for the interchangeability of source and receiver locations without altering the observed waveform. Hence, it is possible to obtain a densely-sampled acoustic characteristics of a pinna, i.e., a receiver, by operating the same structure as an emitter. This is advantageous, because it allows to determine the gain of the pinna for a large set of directions from a single simulation, whereas achieving this by simulating reception of an incoming waveform would require one simulation for each direction [72]. To realize this, a monopole source, embodied by a sphere with a radius of 0.1 mm was positioned 0.5 mm away from the terminal end of the artificial ear canal. Harmonic oscillations were induced in the sphere through a maximum

inward normal displacement of 0.05 mm on the boundary of the sphere generating waves of consistent amplitude at a single frequency.

To align the simulations with prior physical replicas of greater horseshoe bat pinna, the model was scaled $2\times$ over the biological paragon and the emitted ultrasonic frequencies were halved in order to preserve the ratio of the pinna dimensions to the acoustic wavelengths. The strongest (second) harmonics in the biosonar pulses of greater horseshoe bats have narrowband portions around 80 kHz with frequency modulated pulse portions sweeping down over a bandwidth of about 20 kHz [55]. In the current numerical simulation, the scaled model was analyzed for frequencies of 30, 35, and 40 kHz, i.e., frequencies that would cover most of the second harmonic of a greater horseshoe bat if scaled to the pinna size of these animals.




Figure 2.1: Workflow for the numerical beampattern predictions: (a) 3D tomographic scan of a greater horseshoe bat’s right ear used to create the pinna-shape model; (b) Digital pinna model rigged with actuators used to generate the different shape conformations; (c) Pinna surface mesh used for the BEM simulations; (d) Numerical beampattern prediction from the BEM model for a frequency of 40 kHz.

To solve the large linear systems associated with the BEM formulation of the Helmholtz equation, the left-preconditioned generalized minimal residual (GMRES, [15, 70]) algorithm was utilized. This solver was run until the residual norm, a dimensionless measure of solution error, fell below a value of 0.1 or until a maximum of 10,000 iterations was reached.

To enhance the solver’s performance and reduce the likelihood of convergence to a local

minimum, the Sparse Approximate Inverse (SAI) preconditioning method [4, 71] was applied every two iterations to the left side of the linear system matrix (COMSOL Multiphysics v5.4 Acoustics Module, COMSOL AB, Stockholm, Sweden, maximum column size factor 5, preconditioner adapted based on the detected symmetry properties of the system matrix, relaxation factor one).

The directivity pattern of the pinna was estimated from the simulation results by collecting sound pressure level values on a sphere with a radius of 3 m, i.e., well beyond the Fraunhofer distance associated with the highest studied frequency and the largest pinna dimension (0.6 m at 40 kHz [34]). The sphere's surface was discretized into 3,072 uniformly distributed points by virtue of a hierarchical equal area iso-latitude pixelization (HealPix with grid resolution parameter of 16 (Figure 2.2, [19])). The sound pressure amplitudes were then obtained by virtue of piecewise linear interpolation at the points on the sphere from a cubic grid with dimensions of 6 m and a minimum resolution of 2 mm centered on the pinna model.

To facilitate the visualization of the 3D beam patterns, a surface representations of the beam gain (logarithmic sound-pressure-level values normalized to a maximum of 0 dB and clipped at minimum of -40 dB) as a function of direction was created using Delaunay triangulations [8] (Matlab, [24]).

The widths of the main-lobes of the 3D beam patterns were measured using threshold value that was half-power down from the maximum (i.e., -3 dB). The sidelobes were counted based on prominence, i.e., the number of peaks beyond the -6 dB threshold down to -20 dB.

For the deep learning experiments, the entire dataset of 1,331 beampatterns was partitioned into 60% training data, 20% validation data, and 20% testing data [27, 35].

Three deep neural networks (DNN) approaches to beampattern prediction were tested: Two multilayer perceptrons (MLP, [69]) and a deep radial basis function neural network that in-

incorporated Von Mises functions (“acoustic VMBFNN”, [25]). Of the two MLPs, the “acoustic MLP”, predicted the beampatterns based on the respective shape of the model pinna, whereas the other, the “comprehensive MLP”, predicted the beampattern based on the actuator states. Hence, whereas the acoustic MLP covered only the outcomes of the acoustic diffraction process, the predictions of the comprehensive MLP spanned the mechanics (as represented by the skeletal animation technique used) as well as the acoustic diffraction process. The performance of the two MLPs was compared to gain insight into whether prediction of the entire causal chain from actuator input to beampattern is more difficult than predicting only beampattern from the pinna shape. Like the acoustic MLP, the acoustic VMBFNN was limited to predicting beampatterns from the pinna shape conformation. Its performance was compared to the acoustic MLP to assess the potential advantages of using basis functions in terms of accuracy or efficiency of the predictions.

Inspired by the Inception network, a deep convolutional neural network that extracts features at multiple scales using varying kernel sizes [85, 90], the network architectures in this study were fully connected variants, employing linear layers instead of convolutional filters for both upsampling and downsampling. This design choice is aimed at enhancing feature learning and overall performance [85, 90]. Hyperbolic tangent (tanh) activation functions were inserted between linear layers, given their proven efficacy with physics-informed datasets in general [29, 67] and to solve problems governed by partial differential equations (PDEs), such as the Helmholtz equation in this study, in particular. To address the vanishing/exploding gradient challenge through weight initialization, the Kaiming uniform method was adopted, a standard technique that is well suited for tanh activation functions (among others, [21]).

All DNNs tested in this study produced an output of predicted sound pressure levels on a normalized decibel scale as a function of spherical coordinates that ranged over 360 °in azimuth and from 0 to 90°in elevation. Training was based on an L2 mean squared error

(MSE) as the loss function. All models were implemented in Pytorch (v1.9, [60, 61]).

The acoustic networks (MLP and VMBFNN) were relatively shallow with three hidden layers whereas the comprehensive MLP had eight fully connected layers. In either case, each layer was succeeded by a hyperbolic tangent ($\tanh$) activation function (Figure 2.4). The acoustic DNNs (VMBFNN and MLP) started with an input size of two (azimuth and elevation values) and immediately expanded to a maximum node count of 128, followed by two downsampling layers, ending in a single output layer (Figure 2.3). In contrast to this, the comprehensive MLP network had an input of length five (three actuator states, azimuth, and elevation) and employed a hierarchical upsampling scheme, expanding the feature space through the layers to sized of 16, 32, 64, and 128 neurons respectively, followed by a symmetrical downsampling process ultimately leading to a single output neuron [85, 90].

The von Mises basis functions were incorporated into the initial activation layer of the acoustic VMBFNN (Figure 2.3) as has been previously done for a single-layer neural network with von Mises basis functions that was designed to model head-related transfer functions in humans [25]. Each von Mises basis function in the network was characterized by three tunable parameters: two location parameters (α, β) , corresponding to the azimuth and elevation of the function’s centroid on the sphere, and a concentration parameter (κ) that controls the function’s width [25]. The two location parameters of the von Mises basis functions in the acoustic VMBFNN were initialized using a normal distribution

while the concentration parameter was initialized with a constant value of zero signifying a function with a constant amplitude across the entire sphere.

Due to the differences in the number of hidden or initial activation layers, each DNN model type had a unique number of trainable parameters for which varying optimization strategies were used: The shallower acoustic VMBFNN and acoustic MLP architectures contained

10,881 and 10,753 parameters respectively, that were trained over 10,000 epochs across five independent trials, leveraging the Adam optimization algorithm (learning rate 10^{-3} , weight decay 10^{-8}) for L2 regularization [32]. The cumulative sequential training duration for the 39,930 acoustic VMBFNN and acoustic MLP networks ($1,331$ shape conformations $\times 3$ frequencies $\times 2$ network types $\times 5$ repetitions) of approximately 300 hours. Using parallel computing distributed over three NVIDIA A100-80G GPUs, the training could be completed in less than a day. Meanwhile, the comprehensive MLP architecture contained 88,017 trainable parameters which were trained using an identical Adam optimizer [32] with a new initial learning rate of 5×10^{-4} , weight decay of $1e-8$ and a batch size of 1,024 for faster convergence within 200 epochs and 5 trials. To overcome the comprehensive MLP's tendency to get trapped in undesirable local minima during the early training phase, a linear learning rate scheduler with a start factor of $1e-3$ for a total of 200 iterations was tested for estimating reasonable bounds. This approach proved sufficient, obviating the need for further employing a cyclical learning rate [80].

2.4 Results

For both the upright and the bent pinna shape conformations, the width of the mainlobes of the numerical (BEM) beampattern estimates appeared to decrease with increasing frequency (Figure 2.5. At 30 kHz, the -3 dB widths of the mainlobes were 19 and 14 degrees for the upright and bent shape conformations, whereas the respective values were 12 and 6 degrees for 40 kHz. Across all three studied frequencies (30, 35, 40 kHz), the numerical (BEM) beampattern estimates associated with the bent shape conformation had noticeable stronger sidelobes than those associated with the upright shape conformation (Figure 2.5b). For the bent shape conformation, the average gain of the sidelobe maxima was -12.4 dB, whereas

for the upright shape conformation, it was -21.4 dB. However, due to the small sample size ($N=9$), this difference was not statistically significant (paired t -test, $p=0.17$). Another effect that occurred consistently across all three frequencies was that the orientation of the beampattern’s mainlobe shifted from the upright to bent shape conformation by 11° in azimuth and 9° in elevation on average ($N=3$).

All three deep neural networks were able to learn the function approximation task of predicting the beampatterns for the different frequencies and pinna shape conformations without any evidence of overfitting (Figure 2.6). The beampattern predictions obtained by virtue of each of the three deep neural networks (acoustic MLP, acoustic VMBFNN, and comprehensive MLP) overall matched the numerical (BEM) predictions well for all pinna shape conformations and studied frequencies (Figure 2.8).

Furthermore, the deep-neural-network predictions showed the same basic behaviors of the mainlobe and sidelobes across the different shape conformations and frequencies that had been observed in the numerical (BEM) results (Figure 2.7). The differences seen when comparing the deep-learning and numerical (BEM) predictions were typically in the shape detail such as the shape of the peaks or the depth of the notches (Figure 2.8a). Major deviations, e.g., in the number of lobes were found occasionally, but were typically limited to lower gain amplitudes (Figure 2.8a). The two types deep neural networks that were used to predict acoustic properties from shape differed substantially in the amount of time that was needed for training and making inferences (Figure 2.10): The median training time of the acoustic MLP was found to be 25% faster than that of the acoustic VMBFNN (27.6 versus 34.6 seconds, Figure 2.10a). Similarly, the median inference time for the acoustic MLP was 24% faster than the median inference time of the acoustic VMBFNN (3.6 versus 2.9 seconds, Figure 2.10b). The training time for the comprehensive MLP was on average 1.76 hours while the inference time took merely an average of 0.0231 seconds for each model.

2.5 Discussion

The numerical (BEM) and deep-learning beampattern predictions obtained in the current work were in line with expectations from the acoustic fundamentals of diffraction of sound by a finite aperture [91]: Qualitatively, the beampatterns were composed of a mainlobe surrounded by sidelobes that behaved as expected from basic acoustics. Specifically, the widths of the mainlobes were reduced as the simulated acoustic frequencies were increased and the sidelobes became more numerous. Hence, it can be assumed that all numerical predictions made in this work (by virtue of BEM as well as deep neural networks) are representative of physical systems – even if the specific modelled systems may not have an exact physical counterpart. Furthermore, the beampattern predictions shared a feature that had been previously predicted with different numerical methods for bat pinna shapes [14] or observed in biomimetic pinnae [59] in that the strength of the sidelobes increased with bending of the pinna. This demonstrates that the numerical model studied here is close to not only the fundamental physics of an aperture similar to a bat pinna but shares at least one non-trivial property associated with the dynamics of bat and biomimetic pinnae. Taken together, these findings can be regarded as strong indications that the system simulated here is representative of its physical peers, both biological and biomimetic.

The first major finding of the current work is that all tested deep neural networks could reproduce the numerical (BEM) beampattern predictions fairly well. While there were deviations between the BEM predictions and their deep-learning approximations, qualitative inspection of the results indicates that the major features of the beampatterns, e.g., the peaks and notches at high gains, were pretty much always reproduced correctly. The differences between the estimates were in details such as the exact location of the notches and their depth as well as in the shapes of the peaks. To understand the potential uses of a beampattern, its major qualitative features are probably much more important than these

details. For example, an estimation approach that relies on the presence of a notch could function all the same if the notch was shifted by a small angle, it would just have to be recalibrated for the new notch location. Hence, any robust functional concepts developed from experimenting with a deep-learning representations of a biomimetic pinna could in all likelihood be also realized with a physical system.

The comparison of the acoustic MLP and the acoustic VMBFNN indicates that adding the basis functions into the estimator network did not result in any improvements for the current work – neither in terms of accuracy, nor in terms of training and inference times. Hence, the current attempt to improve estimator performance by inserting prior knowledge of the beampattern functions in the form of the von Mises basis functions can be regarded as a failure. However, von Mises basis functions have been previously used successfully in conjunction with a shallow three layer neural network to model beampatterns associated with human hearing [25]. This difference in the outcome could have several, non mutually exclusive reasons: It could be that the employed basis functions were too far from the actual peak shapes in the bat-like beampatterns to be useful. Alternatively, it could be that the deep neural network could not take advantage of the prior information because of its architecture or the way in which it had been initialized and trained. Finally, it may be the case that the much greater adaptability of current deep neural networks is in itself sufficient to approximate the beampattern shapes quickly and accurately.

The second major finding of the current research is that the comprehensive MLP was able to establish a direct link between actuation parameters and acoustic properties, thereby skipping the physical complexities of the mechanics of the soft robotic pinna as well as those of the acoustic diffraction process. This finding holds considerable potential for creating a digital twin [16] of a deforming biological or biomimetic pinna that could be used to discover strategies for employing the time-variant acoustical properties of these structures to improve

sensing performance. The present results indicate that a DNN model is capable of capturing the outcomes of all physical processes (i.e., mechanical deformation of the geometry and acoustic diffraction). The computational effort associated with these predictions is readily compatible with the needs of deep-learning methods such as deep reinforcement learning that require a vast number of model evaluations.

To further improve performance, especially with respect to accuracy, it could be of interest to test other types of networks against the comprehensive MLP on the current dataset such as physics informed neural networks. These network could use the nature of the wave solutions from the Helmholtz differential equation as part of their loss function to improve the physical fidelity of the results produced.




Figure 2.2: Orthographic view of a spherical mesh used to sample beam pattern gains in the far field (at a radius of 3 m), sampling elements of partition generated by the HEALPix discretization [19].




Figure 2.3: Comparison of the deep neural network architectures used to approximate the beam pattern functions: (a) multi-layer perceptron, (b) deep radial basis function neural network with von Mises basis functions in the first activation hidden layer that represent the three parameters of the respective basis function in each node. For both neural network architectures, the inputs were azimuth and elevation and the output was beam gain (on either a linear or on a dB scale).




Figure 2.4: Comprehensive MLP surrogate model of the pinna for used for beam-pattern prediction: The neural network's input takes the actuation state of three soft actuators (T_1, T_2, T_3) alongside the spherical coordinates (θ, ϕ) to output the respective acoustic beam-pattern gains.




Figure 2.5: Impact of acoustic signal frequency and pinna shape conformation on the numerical (BEM) beampattern estimates: a) upright (dashed line) versus bent pinna (solid line) at 30 kHz; b) different frequencies, 30 kHz (dashed line) versus 40 kHz (solid line), for the upright pinna shape conformation. The visualizations of the beampatterns on the sphere are shown below with the cutting planes for the respective polar plot indicated.




Figure 2.6: Learning curve of the comprehensive MLP trained on the BEM simulation and actuator inputs. Average root mean square (RMSE, in dB) for training (solid black line) and validation (dashed gray line) over 15 trials conducted at frequencies of 30, 35, and 40 kHz. Standard deviations are denoted by the thinner dashed lines in the respective gray level.




Figure 2.7: Example beampattern predictions from the MLP as a function of pinna shape conformation and acoustic frequency: In this set of example shape configurations, the middle actuator transitions from moderately bent (2) to fully bent (10) in three intermediate steps. For each shape configuration, MLP beampattern predictions for 30, 35, and 40 kHz are shown.




Figure 2.8: Comparison of numerical and deep-learning-based beam pattern predictions: Deep-learning-based estimates (left-hand column: acoustic MLP, right-hand column: acoustic VMBFNN) versus numerical predictions (BEM, middle column) for an example pinna deformation due to actuator states 0,6,0 and acoustic frequencies of 30, 35, and 40 kHz.




Figure 2.9: Comparison of example numerical (BEM) and deep-learning based comprehensive MLP predictions: cross-sectional plots of numerical (BEM, dotted line) and comprehensive MLP (dotted-dashed line) estimates. The examples shown are for (a) an upright shape conformation and (b) a bent shape conformation. In both cases, the acoustic frequency is 30kHz. The insets show the orientations of the cutting planes used to create the cross sections relative to the respective beampatterns.




Figure 2.10: Training and inference times for the acoustic neural network architectures evaluated: a) Distribution of training times for the acoustic MLP (dark gray) and the acoustic deep von Mises basis function neural network (VMBFNN, light gray); b) Distribution of inference times. Each box-and-whisker plot represents 15 data points from five trials for each of the three frequency (30, 35, and 40 kHz) that were based on 1,331 deformation models that were split into 60% for training, 20% for validation, and 20% for testing.

Chapter 3

General Discussion and Conclusions

This research makes significant contributions to several interconnected fields, demonstrating the power of interdisciplinary approaches in advancing robotics and artificial intelligence. In the realm of biomimetic sonar, the study provides valuable insights into the biosonar system of bats by verifying the variable dynamics and acoustics of the bat’s pinna [14, 66]. The process involves creating a biomimetic continuum soft robotic replica in Blender software for exploring the complex deformable pinna structures, thus bridging the gap between biomimetic sonar and soft robotics [1, 46]. Soft robots, characterized by their adaptability and shape-changing abilities, face challenges in modeling and controlling complex, non-linear behaviors. This study’s approach to linking actuation parameters with acoustic properties addresses these challenges directly, contributing to the growing integration of AI and machine learning in soft robotics for improved sensory capabilities and closed loop control [31].

This work is particularly relevant as it aligns with ongoing efforts to integrate biomimetic sonar with AI for enhanced sensing capabilities [94, 97]. By focusing on pinna dynamics and their impact on beampatterns, the research lays groundwork for more sophisticated, nature-inspired sonar systems [45].

The application of deep neural networks (DNNs) to predict beampatterns and establish direct links between actuation parameters and acoustic properties places this work at the forefront of machine learning integration in robotics. By comparing different neural network architectures, the study contributes to ongoing research in optimizing AI models for specific

intelligent soft robotic design applications [79]. This approach aligns with the broader trend of using AI to simplify complex physical models [30, 95].

The concept of creating a digital twin [16] for deforming pinnae, with potential applications in reinforcement learning, enables real-time control and adaptation in the biomimetic bat robot. The computational efficiency achieved with the DNN models is crucial for real-time applications in robotics, especially in soft robotics where rapid adaptation to changing environments is often required. This efficiency highlights the practical applicability of the research findings.

Overall, this research showcases the effectiveness of interdisciplinary integration, combining acoustics, soft robotics, and machine learning. Such a biomimetic design approach of advanced robots is necessary for moving autonomous systems from constrained to natural environments in applications such as agriculture, environmental surveillance, and defense [49].

3.1 Research Accomplishments

This research significantly advances our understanding of bat biosonar systems, particularly the complex role of the pinna. By characterizing pinna dynamics through beampatterns, we've developed a time-invariant system that provides crucial insights into bat echolocation. The significance of this work lies in its potential applications for biomimetic sonar systems and our deeper comprehension of biological adaptations. Key research accomplishments include:

- Validation of Blender simulations against physical replicas of greater horseshoe bat pinna, enhancing the accuracy and reliability of future computational models

- Development of a novel numerical boundary element method for evaluating pinna directivity, providing a powerful tool for future acoustic studies
- Creation of a comprehensive pipeline to analyze pinna deformation variability, utilizing 1332 beampatterns across multiple frequencies, which offers unprecedented insights into pinna functionality
- Pioneering investigation of deep neural networks as digital twins for biomimetic pinnae, paving the way for advanced reinforcement learning algorithms in pinna control

3.2 Major Findings

- All tested deep neural networks were capable of reproducing the numerical (BEM) predictions well. For instance, the peaks and notches of the numerical beampatterns were almost always reproduced correctly
- The comprehensive MLP was established as a linear time invariant model between actuation parameters and acoustic characteristics of the pinna. This way, a surrogate model was produced that skipped the physical complexity of the soft robot's mechanics as well as the acoustic diffraction upon the pinna surface
- a simple actuation model of the pinna with three degrees of freedom was created and utilized to extract 1,332 meshes that correlated with the actuator states.

3.3 Suggestions for future work

- A linear time invariant model using the Helmholtz equation was made in this work. However, studies show that there are time variant properties in non rigid pinna defor-

mations. Hence, a linear time variant model using the nonlinear Westervelt equation could be devised

- In this study, the shallow three layer von mises basis function was tested as a spherical function approximator of beam patterns. A deeper multi-layered radial basis function network using von mises basis functions as the intermediate layers could be tested.
- A physics informed neural network utilizing the helmholtz equation as a part of the deep neural network's loss function could be implemented to create a physical surrogate model of the bioimimetic pinna.

Bibliography

- [1] Blender Foundation. Blender, 2022. URL <https://www.blender.org/>.
- [2] David S. Broomhead and David Lowe. Multivariable functional interpolation and adaptive networks. *Complex Syst.*, 2, 1988.
- [3] Philip Caspers and Rolf Müller. A design for a dynamic biomimetic sonarhead inspired by horseshoe bats. *Bioinspiration & Biomimetics*, 2018. doi: 10.1088/1748-3190/aac788.
- [4] Edmond Chow and Yousef Saad. Parallel approximate inverse preconditioners. In *SIAM Conference on Parallel Processing for Scientific Computing*, 1997.
- [5] COMSOL AB. *COMSOL Multiphysics Reference Manual*, 2018. URL https://doc.comsol.com/5.4/doc/com.comsol.help.comsol/COMSOL_ReferenceManual.pdf.
- [6] Martin Costabel. Principles of boundary element methods. *Computer Physics Reports*, 6(1):243–274, 1987. ISSN 0167-7977. doi: 10.1016/0167-7977(87)90014-1. URL <https://www.sciencedirect.com/science/article/pii/0167797787900141>.
- [7] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, Dec 1989. ISSN 1435-568X. doi: 10.1007/BF02551274. URL <https://doi.org/10.1007/BF02551274>.
- [8] Boris Delaunay et al. Sur la sphere vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, 7(793-800):1–2, 1934.
- [9] Henry Dicks. The philosophy of biomimicry. *Philosophy & Technology*, 29(3):

- 223–243, 2016. doi: 10.1007/s13347-015-0210-2. URL <https://doi.org/10.1007/s13347-015-0210-2>.
- [10] John F. Eisenberg and Don E. Wilson. Relative brain size and feeding strategies in the Chiroptera. *Evolution*, 32(4):740–751, 1978. Received July 22, 1977. Revised December 1, 1977.
- [11] Benjamin Falk, Joseph Kasnadi, and Cynthia F. Moss. Tight coordination of aerial flight maneuvers and sonar call production in insectivorous bats. *The Journal of Experimental Biology*, 2015. doi: 10.1242/jeb.122283.
- [12] M. Brock Fenton, Alan D. Grinnell, Arthur N. Popper, and Richard R. Fay, editors. *Bat Bioacoustics*. Springer Handbook of Auditory Research. Springer, New York, NY, 1 edition, 2016. ISBN 978-1-4939-3525-3. doi: 10.1007/978-1-4939-3527-7. 49 b/w illustrations, 21 illustrations in colour.
- [13] Yanqing Fu, Philip Caspers, and Rolf M’uller. A dynamic ultrasonic emitter inspired by horseshoe bat noseleaves. *Bioinspiration Biomimetics*, 11(3):036007, 2016. doi: 10.1088/1748-3190/11/3/036007.
- [14] Li Gao, Sreenath Balakrishnan, Weikai He, Zhen Yan, and Rolf Müller. Ear deformations give bats a physical mechanism for fast adaptation of ultrasonic beam patterns. *Phys. Rev. Lett.*, 107:214301, Nov 2011. doi: 10.1103/PhysRevLett.107.214301. URL <https://link.aps.org/doi/10.1103/PhysRevLett.107.214301>.
- [15] Anne Greenbaum. Iterative methods for solving linear systems. *Frontiers in applied mathematics*, 1987. doi: 10.1137/1.9781611970937.
- [16] Michael Grieves. Digital twin: manufacturing excellence through virtual factory replication. *White paper*, 1(2014):1–7, 3 2014. Cited by 2694 as of citation creation.

- [17] Donald R. Griffin. *Listening in the dark: the acoustic orientation of bats and men*. Yale University Press, 1958.
- [18] Donald R. Griffin, D. C. Dunning, D. A. Cahlander, and F. A. Webster. Correlated orientation sounds and ear movements of horseshoe bats. *Nature*, 1962. doi: 10.1038/1961185a0.
- [19] K.M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. Hansen, M. Reinecke, and M. Bartelmann. Healpix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *The Astrophysical Journal*, 622(2):759–771, 2005.
- [20] Mohamad H. Hassoun. Fundamentals of artificial neural networks. *Proceedings of the IEEE*, 1995. doi: 10.1063/1.4822376.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *IEEE International Conference on Computer Vision*, 2015. doi: 10.1109/iccv.2015.123.
- [22] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989. ISSN 0893-6080. doi: [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8). URL <https://www.sciencedirect.com/science/article/pii/0893608089900208>.
- [23] Thomas Howard, Colin Green, David Ferguson, and Alonzo Kelly. State space sampling of feasible motions for high-performance mobile robot navigation in complex environments. *Journal of Field Robotics*, 25(7):325 – 345, June 2008.
- [24] The MathWorks Inc. Matlab version: 9.14.02239454 (r2022a), 2023. URL <https://www.mathworks.com>.

- [25] Rick L. Jenison. A spherical basis function neural network for approximating acoustic scatter. *The Journal of the Acoustical Society of America*, 99(5):3242–3245, 05 1996. ISSN 0001-4966. doi: 10.1121/1.414869. URL <https://doi.org/10.1121/1.414869>.
- [26] Gareth Jones and Jeremy M. V. Rayner. Foraging behavior and echolocation of wild horseshoe bats *rhinolophus ferrumequinum* and *r. hipposideros* (chiroptera, rhinolophidae). *Behavioral Ecology & Sociobiology*, 1989. doi: 10.1007/bf00302917.
- [27] V. Roshan Joseph and V. Roshan Joseph. Optimal ratio for data splitting. *Statistical Analysis and Data Mining*, 2022. doi: 10.1002/sam.11583.
- [28] Leslie Pack Kaelbling, Michael L. Littman, and Andrew W. Moore. Reinforcement learning: A survey. *arXiv: Artificial Intelligence*, 1996. doi: 10.1613/jair.301.
- [29] B.L. Kalman and S.C. Kwasny. Why tanh: choosing a sigmoidal function. *[Proceedings 1992] IJCNN International Joint Conference on Neural Networks*, 1992. doi: 10.1109/ijcnn.1992.227257.
- [30] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6): 422–440, January 2021. doi: 10.1038/s42254-021-00314-5.
- [31] Daekyum Kim, Sang-Hun Kim, Taekyoung Kim, Brian Byunghyun Kang, Minhyuk Lee, Wookeun Park, Subyeong Ku, DongWook Kim, Junghan Kwon, Hochang Lee, Joonbum Bae, Yong-Lae Park, Kyu-Jin Cho, and Sungho Jo. Review of machine learning methods in soft robotics. *PLOS ONE*, 16(2):e0246102, 2 2021. doi: 10.1371/journal.pone.0246102.
- [32] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

- [33] Jens Kober and Jan Peters. Reinforcement learning in robotics: A survey. *Int. J. Robotics Res.*, 2012. doi: 10.1177/0278364913495721.
- [34] John Kraus. *Antennas for All Applications*. McGraw-Hill, New York, 3 edition, 2002.
- [35] Anders Krogh and Jesper Vedelsby. Neural network ensembles, cross validation, and active learning. In G. Tesauro, D. Touretzky, and T. Leen, editors, *Advances in Neural Information Processing Systems*, volume 7. MIT Press, 1994. URL https://proceedings.neurips.cc/paper_files/paper/1994/file/b8c37e33defde51cf91e1e03e51657da-Paper.pdf.
- [36] Sanmeel Vijay Lagad, Ibrahim M. Eshera, Sounak Chakrabarti, and Rolf Müller. Development of a tension-controlled soft-robotic actuation system for a biomimetic bat robot. *Journal of the Acoustical Society of America*, 2021. doi: 10.1121/10.0008446.
- [37] Sergey Levine, Peter Pastor, Alex Krizhevsky, Julian Ibarz, and Deirdre Quillen. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *The International Journal of Robotics Research*, 2018. doi: 10.1177/0278364917710318.
- [38] Yijun Liu. On the bem for acoustic wave problems. *Engineering Analysis With Boundary Elements*, 2019. doi: 10.1016/j.enganabound.2019.07.002.
- [39] Brian K McNab and Meike Köhler. The difficulty with correlations: Energy expenditure and brain mass in bats. *Comparative Biochemistry and Physiology Part A: Molecular & Integrative Physiology*, 212:9–14, Oct 2017. doi: 10.1016/j.cbpa.2017.06.017. URL <https://doi.org/10.1016/j.cbpa.2017.06.017>.
- [40] F. De Mey, F. De Mey, Jonas Reijniers, Herbert Peremans, Makoto Otani, M. Otani, and Uwe Firzlaß. Simulated head related transfer function of the phyllostomid bat

- phyllostomus discolor. *Journal of the Acoustical Society of America*, 2008. doi: 10.1121/1.2968703.
- [41] Piotr Mirowski, Razvan Pascanu, Fabio Viola, Hubert Soyer, Andrew J. Ballard, Andrea Banino, Misha Denil, Ross Goroshin, Laurent Sifre, Koray Kavukcuoglu, Dharmashan Kumaran, and Raia Hadsell. Learning to navigate in complex environments, 2017.
- [42] Volodymyr Mnih, Adrià Puigdomènech Badia, Mehdi Mirza, Alex Graves, Timothy P. Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning, 2016.
- [43] Joachim Mogdans, Joachim Ostwald, and Hans-Ulrich Schnitzler. The role of pinna movement for the localization of vertical and horizontal wire obstacles in the greater horseshoe bat, *rhinolopus ferrumequinum*. *Journal of the Acoustical Society of America*, 1988. doi: 10.1121/1.397183.
- [44] P.M.C. Morse, K.U. Ingard, and P.M. Morse. *Theoretical Acoustics: By Philip M. Morse and K. Uno Ingard*. McGraw-Hill, 1968. URL <https://books.google.com/books?id=NmQAYAEACAAJ>.
- [45] Rolf Müller. Dynamics of biosonar systems in horseshoe bats. *The European Physical Journal Special Topics*, 224:3393–3406, 2015. doi: 10.1140/epjst/e2015-50089-7.
- [46] Rolf Müller and Roman Kuc. The evolution of bat robots. *Journal of the Acoustical Society of America*, 16(4):30–38, 2020.
- [47] Rolf Müller. A numerical study of the role of the tragus in the big brown bat. *The Journal of the Acoustical Society of America*, 116(6):3701–3712, 12 2004. ISSN 0001-4966. doi: 10.1121/1.1815133. URL <https://doi.org/10.1121/1.1815133>.

- [48] Rolf Müller. Numerical analysis of biosonar beamforming mechanisms and strategies in bats). *The Journal of the Acoustical Society of America*, 128(3):1414–1425, 09 2010. ISSN 0001-4966. doi: 10.1121/1.3365246. URL <https://doi.org/10.1121/1.3365246>.
- [49] Rolf Müller. Bioinspiration from bats and new paradigms for autonomy in natural environments. *Bioinspiration & Biomimetics*, 19(3):033001, 4 2024. doi: 10.1088/1748-3190/ad311e. Special Issue on Bioacoustics.
- [50] Rolf Müller, Hongwang Lu, Shuyi Zhang, and Herbert Peremans. A helical biosonar scanning pattern in the Chinese Noctule, *Nyctalus plancyi*. *The Journal of the Acoustical Society of America*, 119(6):4083–4092, 06 2006. ISSN 0001-4966. doi: 10.1121/1.2200151. URL <https://doi.org/10.1121/1.2200151>.
- [51] Rolf Müller, Hongwang Lu, and John R. Buck. Sound-diffracting flap in the ear of a bat generates spatial information. *Physical Review Letters*, 2008. doi: 10.1103/physrevlett.100.108701.
- [52] Rolf Müller, Anupam K Gupta, Hongxiao Zhu, Mittu Pannala, Uzair S Gillani, Yanqing Fu, Philip Caspers, and John R Buck. Dynamic substrate for the physical encoding of sensory information in bat biosonar. *Physical Review Letters*, 118(15):158102, 4 2017. doi: 10.1103/PhysRevLett.118.158102.
- [53] Rolf Müller, Xiaoyan Yin, Ruihao Wang, LiuJun Zhang, and Michael Goldsworthy. A brain for a batbot: Combining deep learning and biomimetic robots to understand and replicate bat biosonar. *Journal of the Acoustical Society of America*, 2020. doi: 10.1121/1.5147577.
- [54] G. Neuweiler, W. Metzner, U. Heilmann, et al. Foraging behaviour and echolocation

- in the rufous horseshoe bat (*Rhinolophus rouxi*) of Sri Lanka. *Behavioral Ecology and Sociobiology*, 20:53–67, 1987. doi: 10.1007/BF00292166.
- [55] Gerhard Neuweiler, Gerd Schuller, and Hans-Ulrich Schnitzler. On- and off-responses in the inferior colliculus of the greater horseshoe bat to pure tones. *Journal of Comparative Physiology A-neuroethology Sensory Neural and Behavioral Physiology*, 1971. doi: 10.1007/bf00297790.
- [56] Anh Nguyen, Ngoc Nguyen, Kim Tran, Erman Tjiputra, and Quang D. Tran. Autonomous navigation in complex environments with deep multimodal fusion network. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5824–5830, 2020. doi: 10.1109/IROS45743.2020.9341494.
- [57] G. Nguyen, S. Dlugolinsky, M. Bobák, et al. Machine learning and deep learning frameworks and libraries for large-scale data mining: a survey. *Artificial Intelligence Review*, 52:77–124, 2019. doi: 10.1007/s10462-018-09679-z.
- [58] olivier Truellier, Sidney I. Wiener, Alain Berthoz, and Jean-Arcady Meyer. Biologically based artificial navigation systems: Review and prospects. *Progress in Neurobiology*, 51(5):483–544, 1997. ISSN 0301-0082. doi: [https://doi.org/10.1016/S0301-0082\(96\)00060-3](https://doi.org/10.1016/S0301-0082(96)00060-3). URL <https://www.sciencedirect.com/science/article/pii/S0301008296000603>.
- [59] Mittu Pannala, Sajjad Zeinoddini Meymand, and Rolf Müller. Interplay of static and dynamic features in biomimetic smart ears. *Bioinspiration & Biomimetics*, 8(2): 026008, may 2013. doi: 10.1088/1748-3182/8/2/026008. URL <https://dx.doi.org/10.1088/1748-3182/8/2/026008>.
- [60] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic

- differentiation in pytorch. In *NIPS 2017 Workshop on Autodiff*, 2017. URL <https://openreview.net/forum?id=BJJsrnfCZ>.
- [61] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [62] A.D. Pierce. *Acoustics: An Introduction to Its Physical Principles and Applications*. McGraw-Hill series in mechanical engineering. McGraw-Hill Book Company, 1981. ISBN 9780070499614. URL <https://books.google.com/books?id=8IgpAQAAMAAJ>.
- [63] Jose C. Principe, Neil R. Euliano, N.R. Euliano, Neil R. Euliano, W. Curt Lefebvre, and W. Curt Lefebvre. Neural and adaptive systems : fundamentals through simulations. *null*, 2000. doi: null.
- [64] J. D. Pye and L. H. Roberts. Ear movements in a hipposiderid bat. *Nature*, 1970. doi: 10.1038/225285a0.
- [65] J. D. Pye, M. Flinn, and A. Pye. Correlated orientation sounds and ear movements of horseshoe bats. *Nature*, 1962. doi: 10.1038/1961186a0.
- [66] Peiwen Qiu and Rolf Müller. Variability in the rigid pinna motions of hipposiderid bats and their impact on sensory information encoding. *Journal of the Acoustical Society of America*, 2020. doi: 10.1121/10.0000582.

- [67] M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2018.10.045>. URL <https://www.sciencedirect.com/science/article/pii/S0021999118307125>.
- [68] John M. Ratcliffe, Coen P. H. Elemans, Lasse Jakobsen, Lasse Hjort Jakobsen, and Annemarie Surlykke. How the bat got its buzz. *Biology Letters*, 2013. doi: 10.1098/rsbl.2012.1031.
- [69] Frank Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408, 1958. doi: 10.1037/h0042519. URL <https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.588.3775&rep=rep1&type=pdf>.
- [70] Youcef Saad and Martin H. Schultz. Gmres: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869, 1986. doi: 10.1137/0907058. URL <https://doi.org/10.1137/0907058>.
- [71] Yousef Saad. *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, 2003. doi: 10.1137/1.9780898718003.
- [72] P.N. Samarasinghe, T. D. Abhayapala, and W. Kellermann. Acoustic reciprocity: An extension to spherical harmonics domain. *Journal of the Acoustical Society of America*, 142(4):EL337–EL343, 2017.
- [73] H Schneider. Die ohrmuskulatur von asellia tridens geoffr.(hipposideridae) und myotis myotis borkh.(vespertilionidae)(chiroptera). *Zoologische Jahrbuch (Anat.)*, 79:93–122, 1961.

- [74] H. Schneider and FP Möhres. Die ohrbewegungen der hufeisennasenfledermäuse (chiroptera, rhinolophidae) und der mechanismus des bildhörens. *Z Vergl Physiol.*, 44(1): 1–40, January 1960. doi: 10.1007/BF00339091. URL <https://doi.org/10.1007/BF00339091>.
- [75] H. U. Schnitzler and E. Flieger. Detection of oscillating target movements by echolocation in the Greater Horseshoe bat. *Journal of Comparative Physiology*, 153:385–391, 9 1983. doi: 10.1007/BF00612592. Accepted: 20 June 1983.
- [76] Hans-Ulrich Schnitzler and Elisabeth K. V. Kalko. Echolocation by insect-eating bats. *BioScience*, 2001. doi: 10.1641/0006-3568(2001)051[0557:ebieb]2.0.co;2.
- [77] Hans-Ulrich Schnitzler, Cynthia F. Moss, and Annette Denzinger. From spatial orientation to food acquisition in echolocating bats. *Trends in Ecology & Evolution*, 18(8):386–394, 2003. ISSN 0169-5347. doi: [https://doi.org/10.1016/S0169-5347\(03\)00185-X](https://doi.org/10.1016/S0169-5347(03)00185-X). URL <https://www.sciencedirect.com/science/article/pii/S016953470300185X>.
- [78] Gerd Schuller and George D. Pollak. Disproportionate frequency representation in the inferior colliculus of doppler-compensating greater horseshoe bats: Evidence for an acoustic fovea. *Journal of Comparative Physiology A-neuroethology Sensory Neural and Behavioral Physiology*, 1979. doi: 10.1007/bf00617731.
- [79] Benjamin Shih et al. Electronic skins and machine learning for intelligent soft robots. *Science Robotics*, 5(41):eaaz9239, 2020. doi: 10.1126/scirobotics.aaz9239.
- [80] Leslie N. Smith. Cyclical learning rates for training neural networks. *IEEE Workshop/Winter Conference on Applications of Computer Vision*, 2017. doi: 10.1109/wacv.2017.58.

- [81] N Sundararajan, Narasimhan Sundararajan, P Saratchandran, Paramasivan Saratchandran, and Lu Ying Wei. A review of radial basis function (rbf) neural networks. *null*, 1999. doi: 10.1142/9789812812506_0001.
- [82] Joseph Sutlive, Agoshpreet Singh, Shuxin Zhang, and Rolf Müller. A biomimetic soft robotic pinna for emulating dynamic reception behavior of horseshoe bats. *Bioinspiration & Biomimetics*, 16(1):016016, dec 2020. doi: 10.1088/1748-3190/abbc73. URL <https://dx.doi.org/10.1088/1748-3190/abbc73>.
- [83] Richard S. Sutton, Richard S. Sutton, and Andrew G. Barto. Reinforcement learning: An introduction. *null*, 1988. doi: null.
- [84] Daniel Svozil, Vladimír Kvasnička, and Jiri Pospichal. Introduction to multi-layer feed-forward neural networks. *Chemometrics and Intelligent Laboratory Systems*, 1997. doi: 10.1016/s0169-7439(97)00061-0.
- [85] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–9, 2015. doi: 10.1109/CVPR.2015.7298594.
- [86] Niko Sünderhauf, Oliver Brock, Walter Scheirer, Raia Hadsell, Dieter Fox, Jürgen Leitner, Ben Upcroft, Pieter Abbeel, Wolfram Burgard, Michael Milford, and Peter Corke. The limits and potentials of deep learning for robotics. *Cornell University - arXiv*, 2018. doi: 10.48550/arxiv.1804.06557.
- [87] Biao Tian and Hans-Ulrich Schnitzler. Echolocation signals of the greater horseshoe bat (*rhinolophus ferrumequinum*) in transfer flight and during landing. *Journal of the Acoustical Society of America*, 1997. doi: 10.1121/1.418272.

- [88] Bryan D Todd and Rolf Müller. A comparison of the role of beamwidth in biological and engineered sonar. *Bioinspiration & Biomimetics*, 13(1):016014, 2018. doi: 10.1088/1748-3190/aa9a0f.
- [89] Darshan Trivedi, Christopher D Rahn, William M Kier, and Ian D Walker. Soft robotics: Biological inspiration, state of the art, and future research. *Applied Bionics and Biomechanics*, 5(3):99–117, 2008.
- [90] Benjamin Tsui, William A. P. Smith, and Gavin Kearney. Low-order spherical harmonic hrtf restoration using a neural network approach. *Applied Sciences*, 10(17), 2020. ISSN 2076-3417. doi: 10.3390/app10175764. URL <https://www.mdpi.com/2076-3417/10/17/5764>.
- [91] Robert Urick J. *Principles of Underwater Sound*. McGraw-Hill, 1975.
- [92] Kaoru Usui, Eraqi R. Khannoon, and Masayoshi Tokita. Facial muscle modification associated with chiropteran noseleaf development: insights into the developmental basis of a movable rostral appendage in mammals. *Developmental Dynamics*, 2022. doi: 10.1002/dvdy.472.
- [93] Xiaomei Wang, Xiaomei Wang, Yingqi Li, Ka-Wai Kwok, and Ka-Wai Kwok. A survey for machine learning-based control of continuum robots. *Frontiers in Robotics and AI*, 2021. doi: 10.3389/frobt.2021.730330.
- [94] Ben Westcott, Ibrahim M. Eshera, and Rolf Müller. A robotic platform for integrating biomimetic sonar and AI. *The Journal of the Acoustical Society of America*, 155 (3_{supplement}) : A244 – –A244, 032024. ISSN0001 – 4966. doi : . URL <https://doi.org/10.1121/10.0027374>.

Karen E. Willcox, Omar Ghattas, and Patrick Heimbach. The imperative of physics-based