

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Mixed-Variable Multi-Objective Bayesian Optimization, Design-by-Morphing and their Applications

Permalink

<https://escholarship.org/uc/item/66m2b3m9>

Author

Sheikh, Haris Moazam

Publication Date

2022

Peer reviewed|Thesis/dissertation

Mixed-Variable Multi-Objective Bayesian Optimization, Design-by-Morphing and their
Applications

by

Haris Moazam Sheikh

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Mechanical Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Philip S. Marcus, Chair

Professor Ömer Savaş

Professor Per-Olof Persson

Fall 2022

Mixed-Variable Multi-Objective Bayesian Optimization, Design-by-Morphing and their
Applications

Copyright 2022
by
Haris Moazam Sheikh

Abstract

Mixed-Variable Multi-Objective Bayesian Optimization, Design-by-Morphing and their Applications

by

Haris Moazam Sheikh

Doctor of Philosophy in Engineering - Mechanical Engineering

University of California, Berkeley

Professor Philip S. Marcus, Chair

Fluid flows are non-intuitive. Even with years of experience, non-intuitive behavior of fluids can mean the optimal geometry of fluid machinery is surprising or even extreme (consider, for instance, the bulbous bow of a ship). Finding the optimal design of a hydrodynamic or aerodynamic surfaces is often impossible due to the expense of evaluating the cost functions (say, with computational fluid dynamics) needed to determine the performances of the flows that the surface controls. In addition, inherent limitations of the design space itself due to imposed geometric constraints, conventional parameterization methods, and user bias can restrict all of the designs within a chosen design space regardless of whether traditional optimization methods or newer, data-driven design algorithms with machine learning are used to search the design space.

This dissertation presents two methodologies to address these difficulties: (1) Design-by-Morphing (DbM), a novel strategy for creating a design search space by morphing homeomorphic shapes to create a continuous and constraint-free design search space that can produce radical extrapolated shapes, something which is unique from existing design strategies; and (2) an optimization algorithm to search that space that uses a novel Mixed-variable, Multi-Objective Bayesian Optimization that we call MixMOBO, that can optimize such expensive, black-box problems with minimum number of functions calls. We apply these methodologies for optimization of several problems and present shape optimization of airfoils, draft-tubes for hydrokinetic turbines, and architected meta-materials. In all cases, we show significantly improved and radical designs.

Chapter One of this thesis focuses on the details of the MixMOBO algorithm, the first mixed-variable, multi-objective Bayesian optimization algorithm. MixMOBO outperforms existing algorithms for mixed-variable problems. It details HedgeMO strategy for hedging acquisition function portfolios for multi-objective problems. MixMOBO is then applied for optimization of strain energy density of an architected meta-material structure with

categorical variables. From a design space of 8.5 billion possible candidates, our algorithm is able to optimize the design space with only 250 function evaluation and achieve 10^4 times improvement in strain energy density over existing structures [1, 2]. Chapter Two focuses on applying MixMOBO for design of Cauchy-Symmetric architected meta-material structures. With only 69 function calls, MixMOBO is able to find such a structure from a design space of 10^7 possible structures. Chapter Three demonstrates the use of Design-by-Morphing for optimization of airfoils. We show that with just 25 baseline shapes, we are able to reproduce the UIUC airfoil database with high fidelity and optimize this space to create aerodynamically superior and safer airfoils [3]. Chapter Four focuses on application of design of a draft tube for a hydrokinetic turbine to maximize pressure recovery at the exit of the turbine [4].

Dedicated to

my grandparents who didn't get to see me graduate, but who always believed in me.

my family for giving me unconditional support and love, now and always, even from afar.

I have the high ground.
-*O.W. Kenobi*

Contents

Contents	iii
List of Figures	v
List of Tables	x
1 MixMOBO	1
1.1 Abstract	1
1.2 Introduction	1
1.3 MixMOBO	4
1.4 MixMOBO Problem Statement	5
1.5 Methodology	6
1.6 Validation Tests	10
1.7 Application to Architected Materials	13
1.8 Conclusions	15
1.9 Additional Details for MixMOBO	17
2 Cauchy Symmetric Structures	21
2.1 Abstract	21
2.2 Introduction	22
2.3 Problem Setup	24
2.4 Materials and Methods	28
2.5 Results	31
2.6 Discussion	34
2.7 Conclusions	37
3 Airfoils	38
3.1 Abstract	38
3.2 Introduction	38
3.3 Design-by-Morphing	40
3.4 Optimization Methodology	45
3.5 Results	48

3.6	Discussion	52
3.7	Conclusion	53
3.8	Additional Details for Airfoil Optimization and Design-by-Morphing	54
4	Draft Tubes	61
4.1	Abstract	61
4.2	Introduction	62
4.3	Preliminaries	65
4.4	Methodology	67
4.5	Results	76
4.6	Conclusion	78
4.7	Additional Details for Draft-Tube and Hub Optimization	79
	Bibliography	82

List of Figures

1.1	Performance comparison of MixMOBO against other mixed-variable algorithms	10
1.2	Performance comparison of HedgeMO against other acquisition functions	12
1.3	<i>Top Left:</i> The 4 unit cells, labelled A – D . <i>Top Right:</i> The 2 orientations in which they can be joined. <i>Bottom Left:</i> Optimization results using MixMOBO. <i>Bottom Right:</i> SEM images of Unblemished and Optimum structures.	14
1.4	HIM images of the loaded and unloaded unblemished and optimum structures. (a) Unloaded Unblemished structure (b) Unblemished structure after loading, showing severe fracture and collapse of many beam members. (c) Focused image revealing several fractured beams and the internal collapse of the upper layer that subsequently instigated the accumulation of damage in the underlying layers. (d) Unloaded MixMOBO Optimum structure. (e) MixMOBO Optimum structure after the structure was subjected to the same maximum compressive load as the structure shown in (b). Unloading of the optimum structure showed only excessive plastic deformation without catastrophic collapse and the manifestation of the buckling mode. (f) Focused revealing the effect of buckling that led to deformation but no fracture due to the occurrence of densification. (g) Side view of the unloaded optimum structure. (h) Side view of the unloaded optimum structure revealing that fracture was inhibited throughout the structure due to the densification precipitated by the low critical buckling load. Scale: Each scale bar is equal to $10^{-5}m$	16
2.1	The stiffness matrix of a cubic symmetric unit cell can be described by three independent stiffness components which corresponded to the stress components orthogonal to the direction of strain (left) the stress components in the direction of loading (middle) and the shear modulus (right).	25
2.2	(a) Five types of cubic unit cells, labeled A-E, used to construct a 5 x 5 x 5 lattice. Unit cells are all face-centered cubic geometry to ensure compatible connectivity within the lattice (b) Unit cells are placed together into three layers that are cubic symmetric. Unit cells of the same type are placed into one of 10 positions, labeled 1-10, which the k-vector is composed of. Layer 1, located on the outside of the structure is composed of 5 x 5 unit cells. Layer 2 is the middle layer and is composed of 3 x 3 unit cells, and Layer 3 is a singular unit cell located in the center of the microlattice.	25

2.3	Process diagram beginning with the selection of 50 initial random microlattices. Afterwards the process iteratively loops between (i) application of the MixMOBO algorithm to choose new parameters and (ii) the modeling, simulation, and evaluation of the Cauchy-symmetry cost function. Once a microlattice with sufficiently small cost function value is found, the iterative loop is exited and the optimal design is fabricated and tested to validate its properties.	27
2.4	Benchmarks for MixMOBO.	29
2.5	Optimization progress for Cauchy-Symmetric structures	32
2.6	(a) Orthogonal view of the Cauchy Symmetric structure taken in a helium ion microscope. (b) Top view of structure highlighting with a large depth of focus to resolve internal beam members. Scale bars are both $20 \mu m$	33
2.7	(a) Mechanical loading response curves of the optimal Cauchy Symmetric lattice structure measured using a PI-87 picoindenter. Points are added to both curves to identify slopes that can be employed for the validation of the cost function. (b) Representative frame of the compression measurement. The indenter applies a compressive load to the structure in the lateral direction. (c) Representative frame of the shear measurement. The indenter is attached to the bar that is fabricated on the edge of the structure, applying a shear load to the geometry. Scale bars are both $20 \mu m$	34
2.8	In situ analysis of the force displacement curves for compression and shear. (a) The response of the structure during compression shows a smooth linear response without discernible boundary effects (stage A-B), until the critical failure load is reached and the first layer collapses (stage C), leading to permanent deformation (stage D). (b) The response of the structure during shows the transition from linear behavior to instability and buckling of the lattice members (stage A-B), which result the initiation of microbuckling phenomena and densification, creating a serrated profile (stage B-C) without extensive permanent deformation (stage C-D). The black scale bar is $50 \mu m$	35
3.1	An example of DbM. The coordinates of the baseline shapes are weighted, summed, and normalized to form the coordinates of a morphed shape.	41
3.2	An Example of DbM. Column 1 shows the baseline shapes. Column 2 depicts the elements of the collocation vectors of the baseline shapes plotted as a function of the index i of the collocation vector. Column 3 shows the weighted elements of the collocation vector plotted as a function of the index i of the collocation vector. Column 4 shows the resultant collocation vector of the morphed shape and the morphed shape itself.	41
3.3	Conditioning for intersection removal. (a) Intersections are detected; (b) Blown up image of one intersection. Shape coordinates direction is depicted by arrows; (c) Intersection removed by flipping vector between intersection; (d) Zero area removed by linear interpolation to remove the intersecting area and then smoothed over, shown by hat coordinates	43

3.4	Twenty-five baseline shapes picked from the UIUC airfoil coordinates database[174]. See Appendix 3.8 for more details.	45
3.5	Reconstruction of randomly chosen 100 pre-existing airfoil shapes via DbM using twenty-five baseline shapes in Figure 3.4.	46
3.6	General flowchart of airfoil optimization via DbM	46
3.7	The Pareto front consisting of the optimal airfoil shapes as a result of the 3,000 generation runs of NSGA-II. Twenty-five red hollow circles with indices indicate twenty-five baseline airfoil shape evaluations. See Appendix 3.8 to understand how the clustering is performed.	49
3.8	Nine representative Pareto-optimal airfoil shapes. a-c are in cluster 1, d-f are in cluster 2 and g-i are in cluster 3.	50
3.9	Mean weight distributions of the Pareto-optimal airfoil shapes with respect to twenty-five baseline airfoil shapes.	51
3.10	Morphed airfoil shapes generated by the optimal weight vectors, representing a the total mean of all optimal airfoils' weights, b the coefficients of the PCA axis having the most variance and c the coefficients of the PCA axis having the second-most variance. The black and red surfaces correspond to the first and second half of the collocation points, respectively.	53
3.11	Airfoil performance curves	55
3.12	Multi-objective optimization of benchmark test functions using GA	56
3.13	Projection of the 25-dimensional optimal weight vectors to the 2-dimensional subspace spanned by 2 PCA axes of the dominant variance. <i>k</i> -means clustering with the cluster size of 3 is used to identify the clusters.	58
4.1	Simplified schematic of a hydroelectric power plant. The enumerated cross-sections in red are: (0) the entrance to the turbine; (1) the inlet of the hub/draft tube assembly; and (2) the outlet of the draft tube. The relative size of the draft tube (depicted in red) and hub (depicted in green) has been enlarged for clarity. The turbine blades are between points (0) and (1) . The reservoir surface is open to atmosphere. The pressure at the draft tube exit (2) is P_2 and is a given, fixed reference or gauge pressure.	66
4.2	Optimization flowchart. The baselines for our draft tube and hub shapes are shown. The baselines are morphed together to create new shapes that are evaluated for their performance using CFD. This search space is sequentially optimized using parallel batches (represented by diamonds) of 5 shapes using MixMOBO, until the evaluation budget is reached. Note that the hub and the draft tube shapes are optimized simultaneously.	67

- 4.3 The origin curve is an arbitrarily chosen curve, shown in the figure with a dot-dash, that starts at the inlet and ends at the outlet. Its only constraint is that it must lie within all of the baseline draft tubes. Because all of the baseline draft tubes used in this study are reflection symmetric about the same x - z plane, we chose to embed the origin curve within this symmetry plane. A unit vector \hat{s} lies along the origin curve pointing from the inlet to the outlet. The arc-length, or coordinate, along the origin curve is denoted as s , with a value of $s = 0$ at the entrance and $s = S$ at the exit. We choose N_s equally-spaced collocation points $\{s_i\}$, $i = 0, 1, 2, \dots, N_s - 1$ along the origin curve, with s_0 in the inlet plane, and s_{N_s-1} in the outlet plane. Planes are defined at each value of s_i that are locally perpendicular to the origin curve, with the planes at s_0 and s_{N_s-1} , at the inlet and outlet, respectively. The figure shows three of these planes as dashed lines at s_0 , s_j , and s_{N_j-1} , where $1 \leq j \leq N_s - 2$. The solid lines in each of these planes show the local polar coordinates (r, φ) within each plane. In each of the N_s planes, we chose the origin of φ to lie in the x - z symmetry plane, and the angle $\varphi = 0$ is indicated in each of the three illustrated planes as broken lines. 69
- 4.4 An example of design-by-morphing (DbM). Column 1 shows the baseline shapes. Column 2 shows the radius $R_{k=0,j}^p \equiv r^p(\varphi_0 \equiv 0, s_j)$ of each baseline shape as a function of j . Note that for ease of visualization, we are only plotting the radius at the “top” of the baseline shape, rather than for all φ . Column 3 shows the weighted radius of each baseline at its top (i.e., the product of the baseline’s weight with its radius at the top). Column 4 shows the morphed shape produced from the weights as well as the radius of the top of the morphed draft tube as a function of s 71
- 4.5 Numerical evaluation for all five baseline draft tube shapes using the first baseline hub shape. In our study, each simulation was run for 128,571 time steps. We chose this number of time steps to insure that the computed flow has come to a statistically steady state. Note that it requires approximately 50,000 time steps for the C_{prm} to come to a statistically steady state. (We assume that the flow itself comes to a statistically steady state some time between 50,000 and 128,571 time steps). The sharp heel draft tube (baseline shape **1**), has the highest C_{prm} among the tube baseline shapes. 73
- 4.6 Pressure Recovery at the “Top” or “Bottom” of the Sharp-Heel draft tube and Hub 1 (in the x - z symmetry plane). According to our collocation defined in Section 4.4, the “Top” corresponds to $R_{k=0,j}^1$ and “Bottom” corresponds to $R_{k=N_\varphi/2,j}^1$ as a function of j , where $j = 0$ is at the inlet. We define $C_p(j) \equiv \frac{P_{wall}(j) - P_{wall}(0)}{\frac{1}{2}\rho\left(\frac{Q}{A_1}\right)^2}$, where $P_{wall}(j)$ is the pressure at the draft tube wall. Solid circles are the experimental data; continuous curves are our numerically computations. The time step is $0.00134 D/u_{avg}$ 74

- 4.7 Optimization progress showing the Normalized C_{prm} of the current epoch's optimum as a function of the epoch number. The Normalized C_{prm} is defined in the text. By definition, the Normalized C_{prm} must increase monotonically with the epoch number. The morphed draft tube/hub with the second highest, and third highest C_{prm} 's are shown within the main figure. The designs that are illustrated in the panels to the right of the main figure are the shape of the morphed draft tube/hub design with the highest C_{prm} (framed with a blue box) and also the sharp heel draft tube with the first baseline hub (framed with a red box). The main figure begins at epoch zero and the Normalized C_{prm} at epoch 0 is that of the design of 50 random designs used to initialize MixMOBO with the highest C_{prm} . 77
- 4.8 Benchmarks for determining the number of epochs of MixMOBO needed to find the likely global maximum. The number of random evaluations of the function to be maximize that are needed to initialize MixMOBO, was set at 50. Our MixMOBO algorithm was then run with 5 parallel batch evaluations per epoch. The evaluation budget was 425 black-box function evaluations including the 50 initial random evaluations. Thus, we ran the code for $75 = (425 - 50)/5$ epochs. Because each optimization begins with random evaluations, the algorithm was run 5 times for each test function to determine the average and standard deviation of the number epochs needed to find the global maximum. The mean Normalized Reward, defined as $(current\ optimum - random\ sampling\ optimum)/(global\ optimum - random\ sampling\ optimum)$ is plotted along with the 0.2*standard deviation plotted as colored bands. A Normalized Reward of unity means that the algorithm has successfully found the global maximum. 81

List of Tables

1.1	Experimental values of the critical buckling P_c (minimization objective for Mix-MOBO), strain energy density at buckling and fracture, u_b and u_f respectively, elastic stiffness S , and ratio of normalized strain energy density compared to the Unblemished structure.	13
2.1	Benchmark Test Functions	30
3.1	The model names, features, shape outlines, and XFOIL evaluation results of the 25 baseline shapes used by DbM in this paper. The coordinates of the baseline shapes are obtained from the UIUC airfoil coordinates database[174]. The airfoil evaluation results are obtained for an incompressible outer-flow of $Re = 1 \times 10^6$. The reference evaluation results are interpolated from the Airfoil Tools online database[212], where N/A indicates that there is no data available for this airfoil.	59
3.2	Benchmark Test Functions. All of the test functions are bi-objective with extended to n -dimensional constrained search space.	60
4.1	Draft Tube Mesh Statistics	72
4.2	C_{prm} and DbM weights of draft tube/hub designs. The designs listed are those with the best, second best, and third best C_{prm} 's found with MixMOBO. We also show the epoch number when the design was found by MixMOBO. (See Fig. 4.7.) In addition, the sharp heel draft tube with the baseline 1 hub is listed.	77
4.3	Benchmark test functions	80
4.4	Details about the test functions, other than the Amalgamated function, are given in Tušar, Brockhoff, and Hansen [52]. All of the test functions have known global maxima. We created the Amalgamated function, a piece-wise function formed from commonly used analytical test functions with different features. The Amalgamated function is non-convex and anisotropic (as is the design space of the draft tube/hub), unlike the other test functions listed here, which are isotropic. These other test functions are commonly used for testing optimization algorithms. Similar to the optimization of the draft tube/hub, each test function here has 6 dimensions.	80

Acknowledgments

First and foremost, I would like to thank God for giving me the strength through this long and arduous journey.

I owe a debt of gratitude to my advisor, Prof. Philip Marcus. Thank you sir for putting up with me, listening to my wacky ideas, being supportive of my life and academic pursuits, for seeing the potential in me when I arrived in Berkeley and teaching me to avoid run-on sentences. Thank you for allowing me to pursue my areas of interest freely and for always listening to me. Thank you for being patient with me and tirelessly proofing all of my papers.

I would like to thank my parents who always believed in me and always encouraged me to go further. Thank you for all your sacrifices so that I could pursue my dreams. I have only gotten here because of you. Thank you to my entire family for their unconditional love and support.

Thank you to Professor Grigoropolous, Zacharias Vangelatos and the entire Laser Thermal Lab, for collaborating with me and for some truly innovative ideas. Thank you to Jinge Wang and Sangjoon Lee for their collaboration on the airfoil project, and spending lot of time debugging XFOIL. It would have been a frustrating experience without you guys.

Thank you to all the faculty in the Mechanical Engineering department, particularly Professor Ömer Savaş and Professor Simo Mäkiharju for stimulating discussions on fluids and life. I owe a special thanks for the NERSC group at LBNL, where I interned for a summer during COVID-19.

To all my 6116 lab mates, thank you for providing an invigorating experience. Thanks for listening to my far-fetched ideas, ramblings, rants and putting up with me in general. Thank you for selflessly giving your time whenever I needed. And thanks for keeping the snack counter always full. Special shout out to Chiyu ‘Max’ Jiang for being a great mentor and friend during these years.

To all my friends and peers, thank you so much for keeping me sane during this journey, particularly when we were all stuck inside with COVID. Thank you for the countless hours on chatting on Zoom and COD. I wouldn’t have made it without you guys.

To Arsenal, thank you for giving me hope every season. We’ll get to the top soon. 2022/23 is our season. Go Gunners!

And finally thank you to the kind taxi-cab driver, who picked me up when I first arrived at SFO. I left my backpack that contained all my money, laptop and documents in the backseat that night. Thank you for coming back the same night, unprompted, and handing me the unopened bag. I don’t know your name, since I was too delirious with joy to thank you properly, but you saved me that day.

Chapter 1

Bayesian Optimization For Multi-Objective Mixed-Variable Problems

1.1 Abstract

Optimizing multiple, non-preferential objectives for mixed-variable, expensive black-box problems is important in many areas of engineering and science. The expensive, noisy, black-box nature of these problems makes them ideal candidates for Bayesian optimization (BO). Mixed-variable and multi-objective problems, however, are a challenge due to BO's underlying smooth Gaussian process surrogate model. Current multi-objective BO algorithms cannot deal with mixed-variable problems. We present MixMOBO, the first mixed-variable, multi-objective Bayesian optimization framework for such problems. Using MixMOBO, optimal Pareto-fronts for multi-objective, mixed-variable design spaces can be found efficiently while ensuring diverse solutions. The method is sufficiently flexible to incorporate different kernels and acquisition functions, including those that were developed for mixed-variable or multi-objective problems by other authors. We also present HedgeMO, a modified Hedge strategy that uses a portfolio of acquisition functions for multi-objective problems. We present a new acquisition function, SMC. Our results show that MixMOBO performs well against other mixed-variable algorithms on synthetic problems. We apply MixMOBO to the real-world design of an architected material and show that our optimal design, which was experimentally fabricated and validated, has a normalized strain energy density 10^4 times greater than existing structures.

1.2 Introduction

Optimization is an inherent part of design for complex physical systems. Often optimization problems are posed as noisy black-box problems subject to constraints, where each function

call requires an extremely expensive computation or a physical experiment. Many of these problems require optimizing a mixed-variable design space (combinatorial, discrete, and continuous) for multiple non-preferential objectives. Architected material design [5, 6, 7, 8, 9, 2], hyper-parameter tuning for machine learning algorithms [10, 11, 12], drug design [13, 14], fluid machinery [3, 4] and, controller sensor placement [15] pose such problems. Due to their cost of evaluation, Bayesian optimization is a natural candidate for their optimization.

Much research has gone into Bayesian optimization for continuous design spaces using Gaussian processes (GP) as a surrogate model and efficiently optimizing this design space with a minimum number of expensive function calls [16, 17, 18]. Despite the success of continuous variable Bayesian optimization strategies, multi-objective and mixed-variable problems remain an area of open research. The inherent continuous nature of GP makes dealing with mixed-variable problems challenging. Finding a Pareto-front for multi-objective problems, and parallelizing function calls for batch updates, Q-batch, also remain challenges in the sequential setting of the BO algorithm. Hedge strategies, which use a portfolio of acquisition functions to reduce the effect of choosing a particular acquisition function, have not been formulated for multi-objective problems.

Mixed-Variable BO Algorithms:

We provide a brief description of the current approaches in recent studies for dealing with mixed variables.

One Hot Encoding Approach: Most BO schemes use Gaussian processes as surrogate models. When dealing with categorical variables, a common method is ‘one-hot encoding’ [19]. Popular BO packages, such as GPyOpt and Spearmint [10], use this strategy. However, this can result in inefficiency when searching the parameter space because the surrogate model is continuous. For categorical variables, this approach also leads to a quick explosion in dimensional space [20].

Multi-Armed Bandit (MAB) Approach: Some studies use the MAB approach when dealing with categorical variables where a surrogate surface for continuous variables is defined for each bandit arm. These strategies can be expensive in terms of the number of samples required [21, 22], and they do not share information across categories. An interesting approach, where coupling is introduced between continuous and categorical variables, is presented in the CoCaBO algorithm [20], and it is one of the baselines that we test MixMOBO against.

Latent Space Approach: A latent variable approach has also been proposed to model categorical variables [23, 24, 25, 26]. This approach embeds each categorical variable in a \mathcal{Z} latent variable space. However, the embedding is dependent on the kernel chosen, and for small-data settings can be inefficient.

Modified Kernel Approach: There is a rich collection of studies in which the underlying kernel is modified to work with ordinal or categorical variables. For example, Ru, Alvi, Nguyen, Osborne, and Roberts [20] considers the sum + product kernel; Deshwal, Belakaria, and Doppa [27] proposes hybrid diffusion kernels, HyBO; and Oh, Gavves, and Welling [28] proposes frequency modulated kernels. The BOCS algorithm [29] for categorical variables

uses a scalable modified acquisition function. Pelamatti, Brevault, Balesdent, Talbi, and Guerin [30], Oh, Tomczak, Gavves, and Welling [31], Nguyen, Gupta, Rana, Shilton, and Venkatesh [22], and Garrido-Merchán and Hernández-Lobato [32] all use modified kernels to adapt the underlying surrogate surface. Our approach is unique in that any modified kernel can be incorporated into our framework. Currently we use the modified radial basis function (RBF) kernel for modelling the surrogate surface, with our future research focused on using different kernels in our framework.

Other Surrogate Models: Other surrogate models can be used in place of the GP to model mixed-variable problems such as random forests, an approach used by SMAC3 [33] or tree based estimators, used in the Tree-Parzen Estimator (TPE) [34]. Daxberger, Makarova, Turchetta, and Krause [35] considers a linear model with cross-product features. BORE [36] leverages the connection to density ratio estimation.

Multi-Objective BO Algorithms:

Multi-objective Bayesian optimization (MOBO) has been the subject of some recent studies. BoTorch [37], the popular BO framework, uses the EHVI and ParEGO based on the works of Fonseca, Paquete, and Lopez-Ibanez [38] and Daulton, Balandat, and Bakshy [39] and Daulton, Eriksson, Balandat, and Bakshy [40]. Hyper-volume improvement is the main mechanism used to ensure diversity in generations. ‘ Q -batch’ parallel settings of the above two acquisition functions use hyper-volume improvement and the previously selected point in the same batch to choose the next set of points. For most single-objective BO algorithms with parallel batch selection, the next batch of test points is selected by adding the ‘fantasy’ cost-function evaluation, usually the predicted mean, to the previously selected test point within that batch. However, this commonly used method often leads to overly confident test point selection, and the surrogate surface then needs to be optimized, and sometimes refitted Q times. Using a genetic algorithm (GA), we can select a ‘ Q -batch’ of points with a single optimization of the surrogate surface from the GA generation.

Suzuki, Takeno, Tamura, Shitara, and Karasuyama [41] provide an interesting Pareto-frontier entropy method as an acquisition function, and Shu, Jiang, Shao, and Wang [42] use Pareto-frontier heuristics to formulate new acquisition functions. Their approaches were not extended to mixed-variable problems.

Hedge Strategies

Hedge algorithms have proven to be efficient in dealing with a diverse set of problems by using a portfolio of acquisition functions. ‘GP-Hedge’, introduced by Brochu, Hoffman, and de Freitas [43] is a well-known and efficient algorithm. However, current Hedge algorithms have not been extended for multi-objective problems and, to the authors’ knowledge, there is no existing Hedge strategy implementation that solves such problems.

1.3 MixMOBO

In this paper, we present a Mixed-variable, Multi-Objective Bayesian Optimization (MixMOBO) algorithm, the first generalized framework that can deal with mixed-variable, multi-objective problems in small data setting and can optimize a noisy black-box function with a small number of function calls.

Genetic algorithms, such as NSGA-II [44], are well known for dealing with mixed-variable spaces and finding an optimal Pareto-frontier. However, these algorithms require a large number of black-box function calls and are not well suited to expensive small-data problems. Our approach is to use a GA to optimize the surrogate model itself and find a Pareto-frontier. Diversification is ensured by the distance metrics used while optimizing the surrogate model. This method allows cheap Q-batch samples from within the GA generation, and also allows the use of commonly used acquisition functions such as Expected Improvement (EI), Probability of Improvement (PI) and Upper Confidence Bound (UCB) [43], which work well for single objective problems. We note here that other metrics can easily be incorporated instead of a distance metric within the GA setting and is one of the areas of our future work. Using a GA on a mixed variable surrogate model in a multi-objective setting allows us to work with modified kernels that were developed for mixed-variable problems in literature. We also present a new acquisition function, ‘Stochastic Monte-Carlo’ (SMC), which performs well for multi-objective categorical variable problems [2].

Hedge strategies for Bayesian optimization are efficient for single objective algorithms. We present here our Hedge Multi-Objective (HedgeMO) algorithm, which uses a portfolio of acquisition functions for multi-objective problems and can work with Q-batch updates. It is an extension of GP-Hedge [43], which has regret bounds, and the same bounds hold for HedgeMO.

We note here that MixMOBO is designed for mixed-variable, multi-objective problems. Although there are algorithms in the literature that can solve problems with a *subset* of these attributes (e.g. mixed-variable single-objective or multi-objective continuous variable problems), no algorithm, to our knowledge, can deal with all of these attributes. In addition, MixMOBO outputs a batch of query points and uses HedgeMO, the first multi-objective hedging strategy. To the authors’ knowledge, no existing approaches can achieve all this within a single framework.

In summary, the main contributions of our work are as follows:

- We present Mixed-variable, Multi-Objective Bayesian Optimization (MixMOBO), the first algorithm that can deal with mixed-variable, multi-objective problems. The framework uses GA to optimize the acquisition function on a surrogate surface, so it can use modified kernels or surrogate surfaces developed to deal with mixed-variable problems in previous studies. This extends the capabilities of previous approaches in literature to handle mixed-variable and multi-objective problems as well if adopted within our framework, since our framework is agnostic to the underlying GP kernel over mixed-variables.

- GA is used to optimize surrogate models, which allows the optimization of multi-objective problems. ‘ Q -batch’ samples can be extracted in parallel from within the GA generation without sacrificing diversification.
- We present a Hedge Multi-Objective (HedgeMO) strategy for multiple objectives for which regret bounds hold. We also present an acquisition function, Stochastic Monte-Carlo (SMC), which performs well for combinatorial multi-objective problems, and use it as part of our HedgeMO portfolio.
- We benchmark our algorithm against other mixed-variable algorithms and prove that MixMOBO performs well on test functions. We applied MixMOBO to a practical engineering problem: the design of a new architected meta-material that was optimized to have the maximum possible strain-energy density within the constraints of a design space. The fabrication and testing of this new material showed that it has a normalized strain energy density that is 10^4 times greater than existing unblemished microlattice structures in literature.

The rest of the paper is organized in the following manner: Section 1.4 defines the optimization problem to be solved with MixMOBO. The detailed workings of MixMOBO and HedgeMO are presented in Section 1.5. Section 1.6 details the validation tests performed on our framework to test its efficiency and comparison to existing algorithms. Our application of MixMOBO for design of architected materials and its results are presented in Section 1.7.

1.4 MixMOBO Problem Statement

We pose the multi-objective and mixed-variable problem as:

$$\vec{w}_{opt} = \underset{\vec{w} \in \mathcal{W}}{\operatorname{argmax}}(\vec{f}(\vec{w})) \quad (1.1)$$

for maximizing the objective. Here $\vec{f}(\vec{w}) = [f_1(\vec{w}), f_2(\vec{w}), \dots, f_k(\vec{w})]$ are the K non-preferential objectives to be maximized, and \vec{w} is a mixed-variable vector, defined as $\{\vec{w} \in \mathcal{W}\} = \{\vec{x} \in \mathcal{X}, \vec{y} \in \mathcal{Y}, \vec{z} \in \mathcal{Z}\}$. \vec{x} is an m -dimensional vector defined over a bounded set $\mathcal{X} \subset \mathbb{R}^m$ representing m continuous variables. Ordinal and categorical variables are defined as $\vec{y} = [y_1, \dots, y_n]$ and $\vec{z} = [z_1, \dots, z_o]$, respectively. Each variable $y_j \in \{O_1, \dots, O_j\}$ takes one of O_j ordinal ‘levels’ (discrete numbers on the real-number line) and each categorical variable takes a value $z_j \in \{C_1, \dots, C_j\}$ from C_j unordered categories (that cannot, by definition, be ordered on the real-number line). \mathcal{Y} and \mathcal{Z} are the ordinal and combinatorial spaces respectively.

Generally, $\{\vec{w}_{opt}\}$ is a set of Pareto-optimal solution vectors i.e., vectors that are not Pareto-dominated by any other vector. A vector \vec{w} is Pareto-dominated by \vec{w}' , iff $f_k(\vec{w}) \leq f_k(\vec{w}') \forall k = 1, \dots, K$. This $\{\vec{w}_{opt}\}$ is the optimal set found by MixMOBO, details of which are presented in the following section.

1.5 Methodology

Preliminaries

Single-objective Bayesian optimization is a sequential optimization technique, aimed at finding the global optimum of a single objective noisy black-box function f with minimum number of evaluations of f . For every i^{th} iteration, a surrogate model, g , is fit over the existing data set $\mathcal{D} = \{(w_1, f(w_1)), \dots, (w_i, f(w_i))\}$. An acquisition function then determines the next point \vec{w}_{i+1} for evaluation with f , balancing exploration and exploitation. Data is appended for the next iteration, $\mathcal{D} = \mathcal{D} \cup (w_{i+1}, f(w_{i+1}))$, and the process is repeated until the evaluation budget for f or the global optimum is reached.

Gaussian processes are often used as surrogate models for BO [17, 45]. A GP is defined as a stochastic process such that a linear combination of a finite set of the random variables is a multivariate Gaussian distribution. A GP is uniquely specified by its mean $\mu(\vec{w})$ and covariance function $ker(\vec{w}, \vec{w}')$. The GP is a distribution over functions, and $g(\vec{w})$ is a function sampled from this GP:

$$\vec{g}(\vec{w}) \sim GP(\mu(\vec{w}), ker(\vec{w}, \vec{w}')). \quad (1.2)$$

Here, $ker(\vec{w}, \vec{w}')$ is the covariance between input variables \vec{w} and \vec{w}' . Once a GP has been defined, at any \vec{w} the GP returns the mean $\mu(\vec{w})$ and variance $\sigma(\vec{w})$. The acquisition function $\mathcal{A}(\vec{g})$, balances exploration and exploitation, and is optimized to find the next optimal point \vec{w}_{i+1} . The success of BO comes from the fact that evaluating \vec{g} is much cheaper than evaluating f .

MixMOBO Approach

Our Mixed-variable Multi-Objective Bayesian Optimization (MixMOBO) algorithm extends the single-objective, continuous variable BO approach presented in the preceding section, to more generalized optimization problems and is detailed in Algorithm 1.

A single noisy GP surrogate surface is fit for multiple objectives, $\vec{g}(\vec{w}) \sim GP(\vec{\mu}(\vec{w}), ker(\vec{w}, \vec{w}'))$. Note that this is different from Eq. 1.2, since the GP would predict mean for multiple objectives. For details on fitting a single GP to multi-objective data, we refer the reader to [17, Eq. 2.25-2.26]. For multiple objectives, the response vector, with n -data points, is of size $k \times n$. The predicted variance remains the same, but the predicted mean is a $k \times 1$ vector. This is equivalent to fitting K GP surfaces with the same kernel for all of the surfaces, where K is the total number of objectives. All K objectives are assumed to have equal noise levels. Only one set of hyper-parameters needs to be fit over this single surface, rather than fitting K sets of hyper-parameters for K different surfaces; thus, when K is large, the overall computational cost for the algorithm is reduced. Note that we could fit K different GP surfaces, particularly if different noise levels for different objectives is to be considered, with different hyper-parameters to the data to add further flexibility to the

fitted surfaces. This idea will be investigated in our future work. We use LOOCV [45] for estimating hyper-parameters since we are dealing with small-data problems.

Algorithm 1 Mixed-variable Multi-Objective Bayesian Optimization (MixMOBO) Algorithm

- 1: **Input:** Black-box function $\vec{f}(\vec{w}) : \vec{w} \in \mathcal{W}$, initial data set size $N_{.i}$, batch points per epoch Q , total epochs N , mutation rate $\beta \in [0, 1]$
 - 2: **Initialize:** Sample black-box function \vec{f} for $\mathcal{D} = \left\{ \left(\vec{w}_j, \vec{f}(\vec{w}_j) \right) \right\}_{j=1:N_{.i}}$
 - 3: **for** $n = 1$ **to** N **do**
 - 4: Fit a noisy Gaussian process surrogate function $\vec{g}(\vec{w}) \sim GP(\vec{\mu}(\vec{w}), \text{ker}(\vec{w}, \vec{w}'))$
 - 5: For L total acquisition functions, from each \mathcal{A}^l acquisition function, propose Q -batch test-points, $\{(\vec{u}_n^l)\}_{1:Q} = \{argmax_{\vec{u} \in \mathcal{W}} \mathcal{A}^l(\vec{g})\}_{1:Q}$ within the constrained search space \mathcal{W} using multi-objective GA
 - 6: Mutate point $\{(\vec{u}_n^l)_q\}$ within the search space \mathcal{W} with probability rate β if L_2 -norm of its difference with any other member in set $\{(\vec{u}_n^l)\}_{1:Q}$ is below tolerance
 - 7: Select batch of Q points using HedgeMO, $\{\vec{w}_n\}_{1:Q} = \text{HedgeMO}(\vec{g}, \{(\vec{u}_{1:n}^l)\}_{1:Q}, \mathcal{D})$
 - 8: Evaluate the selected points from the black-box function, $\{\vec{f}(\vec{w}_n)\}_{1:Q}$
 - 9: Update $\mathcal{D} = \mathcal{D} \cup \left\{ \left(\vec{w}_n, \vec{f}(\vec{w}_n) \right) \right\}_{1:Q}$
 - 10: **end for**
 - 11: **return** Pareto-optimal solution set $\left\{ \left(\vec{w}_{opt}, \vec{f}(\vec{w}_{opt}) \right) \right\}$
-

Gaussian processes are defined for continuous variables. For mixed variables, we need to adapt the kernel so that a GP can be fit over these variables. Cited works in Section 1.2 dealt with modified kernels that were designed to model mixed variables. Those kernels can be used in the MixMOBO algorithm. For the current study, we use a simple modified squared exponential kernel:

$$\text{ker}(\vec{w}, \vec{w}') \equiv \epsilon_f^2 \exp \left[-\frac{1}{2} |\vec{w}, \vec{w}'|_C^T \underline{\underline{M}} |\vec{w}, \vec{w}'|_C \right], \quad (1.3)$$

where $\vec{\theta} = (\{\underline{\underline{M}}\}, \epsilon_f)$ is a vector containing all the hyper-parameters, $\{\underline{\underline{M}}\} = \text{diag}(\vec{h})^{-2}$ is the covariance hyper-parameter matrix and \vec{h} is the vector of covariance lengths. The distance metric, $|\vec{w}, \vec{w}'|_C$, is an concatenated vector, with the distance between categorical variables defined to be the Hamming distance, and the distance between continuous variables and the distance between ordinal variables defined to be their Euclidean distances. Noise is added to the diagonal of the covariance matrix. We emphasize that *any* modified kernel discussed in the citations of Section 1.2 can be used within our framework and is a focus of our future work.

Once the GP is fit over multi-objective data, acquisition functions, \mathcal{A}^l , explore the surrogate model to maximize reward by balancing exploration and exploitation. Using a standard optimization scheme is problematic when dealing with mixed-variable and multi-objective problems due to non-smooth surrogate surface and conflicting objectives. We propose using a constrained, multi-objective GA to optimize the acquisition functions, which, although expensive to use on an actual black-box function, is an ideal candidate for optimizing the acquisition function working on the surrogate surface. For multi-objective problems, multi-objective GA algorithms, such as [44], are ideal candidates for obtaining a Pareto-front of optimal solutions.

Within a GA generation, for multi-objectives, diversification is ensured by a ‘distance crowding function’ which ranks the members of a non-dominated Pareto-front. The ‘distance crowding function’ can be computed in decision-variable space, in function space or a hybrid of the two, and ensures that the generations are distinct and diverse. This inherent feature of GA is exploited to ensure diversity in the ‘ Q -batch’ of points. The ranking takes place when choosing the test points from an acquisition function for a multi-objective problem because the choice must take into account the diversity of the solution and propagate the Pareto-front. Because the members of the population are ranked by the GA, we can easily extract a ‘ Q -batch’ of points from each of the acquisition functions without needing to add any ‘fantasy’ cost function evaluations or optimizing the acquisition functions again. This is a great advantage of using GA as our optimizer since we can output a ‘ Q -batch’ of diverse query points using the inherent GA features.

For dealing with mixed-variable problems, GA are again ideal candidates. Genetic algorithms (GA) can be easily be constrained to work in mixed variable spaces. These variables can be dealt with by using probabilistic mutation rates. The genes are allowed to mutate within their prescribed categories, thereby constraining the proposed test points to the \mathcal{W} space.

Common acquisition functions, such as EI, PI, and UCB, can be used within this framework and can be used to nominate a ‘ Q -batch’ of points. If a candidate in a Q -batch is within the tolerance limit of another candidate in the same batch or a previous data point (for convex functions), we mutate the proposed point within \mathcal{W} to avoid sampling the same data point again.

Test points are selected from \mathcal{W} to evaluate their \vec{f} using HedgeMO algorithm which is detailed in the next section. HedgeMO selects a ‘ Q -batch’ of test-points from the candidates proposed by each of the acquisition functions. These points are then, along with their function evaluations \vec{f} s, appended to the data set.

HedgeMO Algorithm

Hedge strategies use a portfolio of acquisition functions, rather than a single acquisition function. It is an extension to multi-objective problems of GP-Hedge algorithm proposed by [43]. HedgeMO is part of our MixMOBO algorithm that not only extends the Hedge

Algorithm 2 HedgeMO Algorithm

- 1: **Input:** Surrogate function $\bar{g}(\vec{w}) : \vec{w} \in \mathcal{W}$, proposed test points by AFs $\left(\{(\vec{u})_{1:n}^{1:L}\}_{1:Q}\right)$, batch points per epoch Q , current epoch n , total objective K , parameter $\eta \in \mathbb{R}^+$
 - 2: **for** $l = 1$ **to** L **do**
 - 3: For l^{th} acquisition function, find rewards for Q -batch points nominated by that AF from epochs $1:n-1$, by sampling from \vec{g} , $\left\{\bar{\theta}_{1:n-1}^l\right\}_{1:Q} = \bar{\mu}(\{(\vec{u})_{1:n-1}^l\}_{1:Q})$, where $\bar{\theta} = \{\theta\}^k$ for each objective k
 - 4: **end for**
 - 5: Normalize rewards for each l^{th} AF and k^{th} objective, $\phi_l^k = \sum_{j=1}^{n-1} \sum_{q=1}^Q \frac{\{\theta_j^l\}_q^k - \min(\Theta)}{\max(\Theta) - \min(\Theta)}$, where Θ is defined as $\Theta = \{\theta_{1:n-1}^{1:L}\}_{1:Q}^k$
 - 6: Calculate probability for selecting nominees from l^{th} acquisition function, $p^l = \frac{\exp(\eta \sum_{k=1}^K \phi_l^k)}{\sum_{i=1}^L \exp(\eta \sum_{k=1}^K \phi_i^k)}$
 - 7: **for** $q = 1$ **to** Q **do**
 - 8: Select q^{th} nominee as test-point \vec{w}_n^q from l^{th} AF with probability p^l
 - 9: **end for**
 - 10: **return** Batch of test points $\{\vec{w}_n\}_{1:Q}$
-

strategy to multi-objective problems, but also allows ‘ Q -batches’. Our algorithm is shown in Algorithm 2.

Extending the methodology presented by [43], HedgeMO chooses the next ‘ Q -batch’ of test points from the history of the candidates nominated by all of the acquisition functions. Rewards are calculated for each acquisition function from the surrogate surface for the entire history of the nominated points by the L acquisition functions. The rewards are then normalized to scale them to the same range for each objective. This step is fundamentally important because it prevents biasing the probability of any objective. This type of bias, of course, cannot occur in single objective problems. The rewards for different objectives k are then summed and the probability, p^l , of choosing a nominee from a specific acquisition function is calculated using step 6 in Algorithm 2. For a ‘ Q -batch’ of tests points, the test points are chosen Q times.

Regret Bounds: The regret bounds derived by Brochu, Hoffman, and de Freitas [43] hold for HedgeMO if and only if the Upper Confidence Bound (UCB) acquisition function is a part of the portfolio of acquisition functions. The regret bounds follow from the work of Srinivas, Krause, Kakade, and Seeger [46] who derived cumulative regret bounds for UCB. In essence, the cumulative regret in our case is bounded by two sublinear terms as for UCB and an additional term which depends on proximity of the chosen point with the test point

proposed by UCB. The interested reader is directed to Srinivas, Krause, Kakade, and Seeger [46] and Brochu, Hoffman, and de Freitas [43] for a description of the exact regret bounds and their derivation.

SMC Acquisition Function

We introduce a new acquisition function, Stochastic Monte-Carlo (SMC), which for the maximization of an objective, is defined as:

$$SMC \equiv \operatorname{argmax}_{\vec{w} \in \mathcal{W}} [\tilde{\mu}(\vec{w}) + r(\vec{w})], \tag{1.4}$$

where $r(\vec{w})$ is sampled from $U(0, 2\sigma(\vec{w}))$, and $\tilde{\mu}(\vec{w})$ and $\sigma(\vec{w})$ are the mean and standard deviation returned by the GP at \vec{w} , respectively. This is equivalent to taking Monte-Carlo samples from a truncated distribution. For categorical and ordinal variable problems, this acquisition function performs well across a range of benchmark tests [2]. We use this acquisition function as part of our portfolio of HedgeMO in the MixMOBO algorithm.

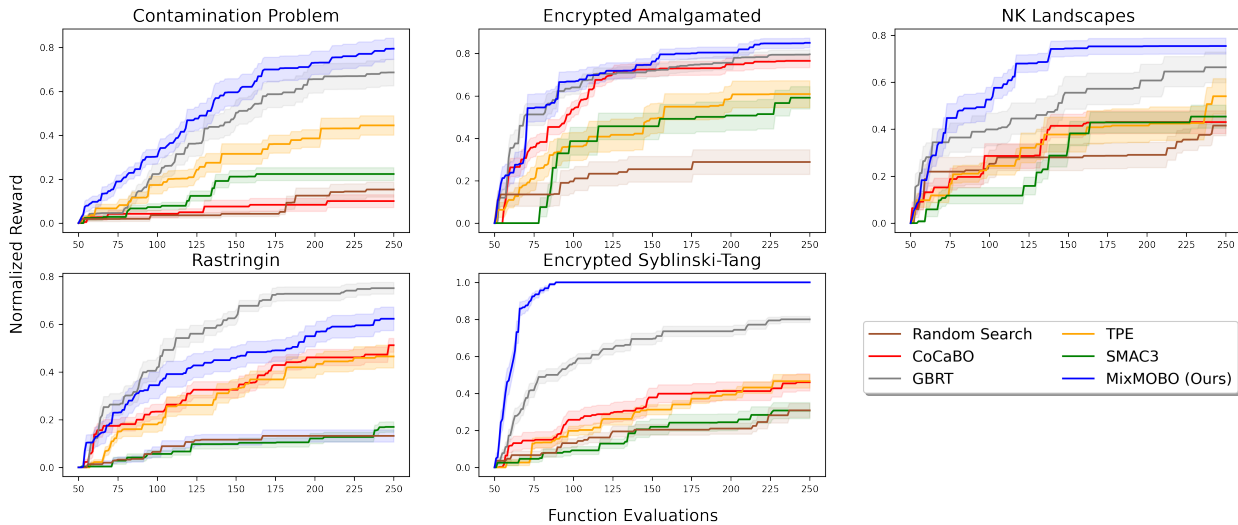


Figure 1.1: Performance comparison of MixMOBO against other mixed-variable algorithms

1.6 Validation Tests

MixMOBO is designed to deal with mixed-variable, multi-objective problems. However, no other small-data algorithm, to the authors’ knowledge, can similarly deal with all the attributes of such problems to provide an honest comparison. In the absence of such competition, we use the specific case of mixed-variable, single-objective problems to provide a comparison to state-of-the-art algorithms present for such problems and prove that even for this subset

case, MixMOBO is able to perform better than existing algorithms in the literature. We then perform further experiments in both single and multi-objective settings to show the efficacy of the HedgeMO algorithm compared to stand-alone acquisition functions and the performance of SMC in the multi-objective setting.

We benchmarked MixMOBO against a range of existing state-of-the-art optimization strategies that are commonly used for optimizing expensive black-box functions with mixed-variable design spaces. We chose the following single objective optimization algorithms for comparison: **CoCaBO** [20], which combines the multi-armed bandit (MAB) and Bayesian optimization approaches by using a mixing kernel. CoCaBO has been shown to be more efficient than GPyOpt (one-hot encoding approach [47]) and EXP3BO (multi-armed bandit (MAB) approach [21]). We used CoCaBO with a mixing parameter of 0.5. We also tested MixMOBO against **GBRT**, a sequential optimization technique using gradient boosted regression trees [48]. **TPE_Hyperopt** (Tree-structured Parzen Estimator) is a sequential method for optimizing expensive black-box functions, introduced by Bergstra, Bardenet, Bengio, and Kégl [49]. **SMAC3** is a popular Bayesian optimization algorithm in combination with an aggressive racing mechanism [50]. Both of these algorithms, in addition to **Random Sampling**, were used as baselines. Publicly available libraries for these algorithms were used.

Six different test functions for mixed variables were chosen as benchmarks. A brief description of these test functions and their properties is given below with further details in Appendix 1.9:

Contamination Problem: This problem, introduced by Hu, Hu, Xu, Wang, and Cao [51], considers a food supply chain with various stages in the chain where food may be contaminated with pathogens. The objective is to maximize the reward of prevention efforts while making sure the chain does not get contaminated. It is widely used as a benchmark for binary categorical variables. We use the problem as a benchmark with 21 binary categorical variables.

Encrypted Amalgamated: An anisotropic, mixed-variable function created using a combination of other commonly used test functions [52]. We modify the combined function so that it can be used with mixed variables. In particular, it is adapted for categorical variables by encrypting the input space with a random vector, which produces a random landscape mimicking categorical variables [2]. Our Encrypted Amalgamated function has 13 inputs: 8 categorical, 3 ordinal variables (with 5 categories or states each) and 2 continuous.

NK Landscapes: This is a popular benchmark for simulating categorical variable problems using randomly rugged, interconnected landscapes [53, 54]. The fitness landscape can be produced with random connectivity and number of optima. The problem is widely used in evolutionary biology and control optimization and is NP -complete. The probability of connectivity between NK is controlled by a ‘ruggedness’ parameter, which we set at 20%. We test the Li, Emmerich, Eggermont, Bovenkamp, Bäck, Dijkstra, and Reiber [54] variant with 8 categorical variables with 4 categories each.

Rastringin: This is an isotropic test function, commonly used for continuous design spaces [52]. We use a 9-D Rastringin function for testing a design space of 3 continuous and 6 ordinal variables with 5 discrete states.




Figure 1.2: Performance comparison of HedgeMO against other acquisition functions

Encrypted Syblinski-Tang: This function is isotropic [52], and we have modified it as we did with the Encrypted Amalgamated test function so that it can work with categorical variables and was used as a representative benchmark for N -categorical variable problems. The 10-D variant tested here consists only of categorical variables with 5 categories each.

Encrypted ZDT6: This is a multi-objective test function introduced by Zitzler, Deb, and Thiele [55] that we modified with encryption so that it can deal with mixed variables. The test function is non-convex and non-uniform in the parameter space. We test ZDT6 with 10 categorical variables with 5 states each. ZDT6 was only used for testing HedgeMO.

To the extent of our knowledge, no other optimization algorithm is capable of handling mixed-variable, multi-objective problems in small-data settings. Thus, we have no direct comparisons between MixMOBO and other published algorithms. Therefore, we tested MixMOBO against a variant of NSGA-II [44] with the ZDT4 and ZDT6 test functions with mixed variables. However, we found that using a GA required more than 10^2 more function calls to find the Pareto front to a similar tolerance. For visualization purposes, we do not plot the GA results.

All of the optimization algorithms were run as maximizers, with a 0.005 noise variance built into all the benchmarks. The budget for each benchmark test was fixed at 250 function calls including the evaluations of 50 initial randomly sampled data points for all algorithms, except for SMAC3 which determines its own initial sample size. The algorithms were run in single output setting (GBRT, CoCaBO and MixMOBO’s batch mode was not used for fair comparison). Each algorithm was run 10 times for every benchmark. Our metric for optimization is the ‘Normalized Reward’, defined as $(current\ optimum - random\ sampling\ optimum)/(global\ optimum - random\ sampling\ optimum)$. Figure 1.1 shows the Normalized Rewards versus the number of black-box function evaluations for MixMOBO and five other algorithms. The mean and standard deviation of the Normalized Rewards of the 10 runs for each algorithm, along with their standard deviations (S.D.), are plotted. The width of each of the translucent colored bands is equal to $1/5$ of their S.D.

MixMOBO outperforms all of the other baselines and is significantly better in dealing with mixed-variable problems. GBRT is the next best algorithm and performs better than MixMOBO on the Rastrigin function; however, note that the Rastrigin function does not

Table 1.1: Experimental values of the critical buckling P_c (minimization objective for MixMOBO), strain energy density at buckling and fracture, u_b and u_f respectively, elastic stiffness S , and ratio of normalized strain energy density compared to the Unblemished structure.

Structure	$P_c[\mu N]$	$u_b[MJm^{-3}]$	$u_f[MJm^{-3}]$	$S[MPa]$	$(u_{fi}/u_{bi})/(u_{f1}/u_{b1})$
Unblemished	3814.5	1.08	0.071	388.21	1
Random Sampling Optimal	996.2	0.08	2.85	347.19	526
MixMOBO Optimal	545.1	0.02	14.71	460.35	12030

include any categorical variables. For problems involving categorical variables, MixMOBO clearly outperforms the others. TPE and CoCaBO have similar performances, and SMAC3 has the poorest performance. All three are outdistanced by MixMOBO.

We then perform experiments to test the performance of our HedgeMO algorithm by comparing it to four different acquisition functions which make up the entirety of its portfolio. These acquisition functions, namely, EI, PI, UCB, and SMC, along with HedgeMO are tested on three different test functions: the Encrypted Amalgamated, Encrypted Syblinski-Tang, and Encrypted ZDT6. The latter is used as the multi-objective test function. The Normalized Reward for the multi-objective Encrypted ZDT6 is defined as $(current\ P\text{-optimum} - random\ sampling\ P\text{-optimum}) / (global\ P\text{-optimum} - random\ sampling\ P\text{-optimum})$. Here, $P\text{-optimum} = \frac{1}{N} \sum_{i=1}^N \exp(-\text{minimum Hamming distance in parameter space between } i^{th} \text{ global Pareto-optimal point and any point in the current Pareto-optimal set})$, where N is the number of global Pareto-optimal points.

The results of our acquisition function comparisons are shown in Figure 1.2, which shows that HedgeMO performs well across all three test functions. For single-objective test functions, PI outperforms HedgeMO for Encrypted Amalgamated test function. However, for the multi-objective Encrypted ZDT6 test function, PI performs significantly worse and is outperformed by both SMC and UCB. SMC performs well on multi-objective problems combinatorial problems and hence should be a part of portfolio for a hedging algorithm.

These results prove that for a range of different problems, acquisition function choice can play a huge role in the performance of the algorithm. For a black-box function, this information can not be known a priori, making hedging necessary. HedgeMO consistently performs well in all scenarios and ensures efficiency across a range of different problems. Thus, for unknown black-box functions, HedgeMO should be the hedge strategy of choice for multi-objective problems.

1.7 Application to Architected Materials

We applied our MixMOBO framework to the optimization of the design of architected, microlattice structures. Advances in modeling, fabrication, and testing of architected materials

have promulgated their utility in engineering applications, such as ultralight [56, 57, 58], reconfigurable [59], and high-energy-absorption materials [60], and in bio-implants [61]. The optimization of architected materials [62, 57, 59, 58] often requires searching huge combinatorial design spaces, where the evaluation of each design is expensive. [63, 64, 65]. The design space for the architected material we optimize here has ~ 8.5 billion possible

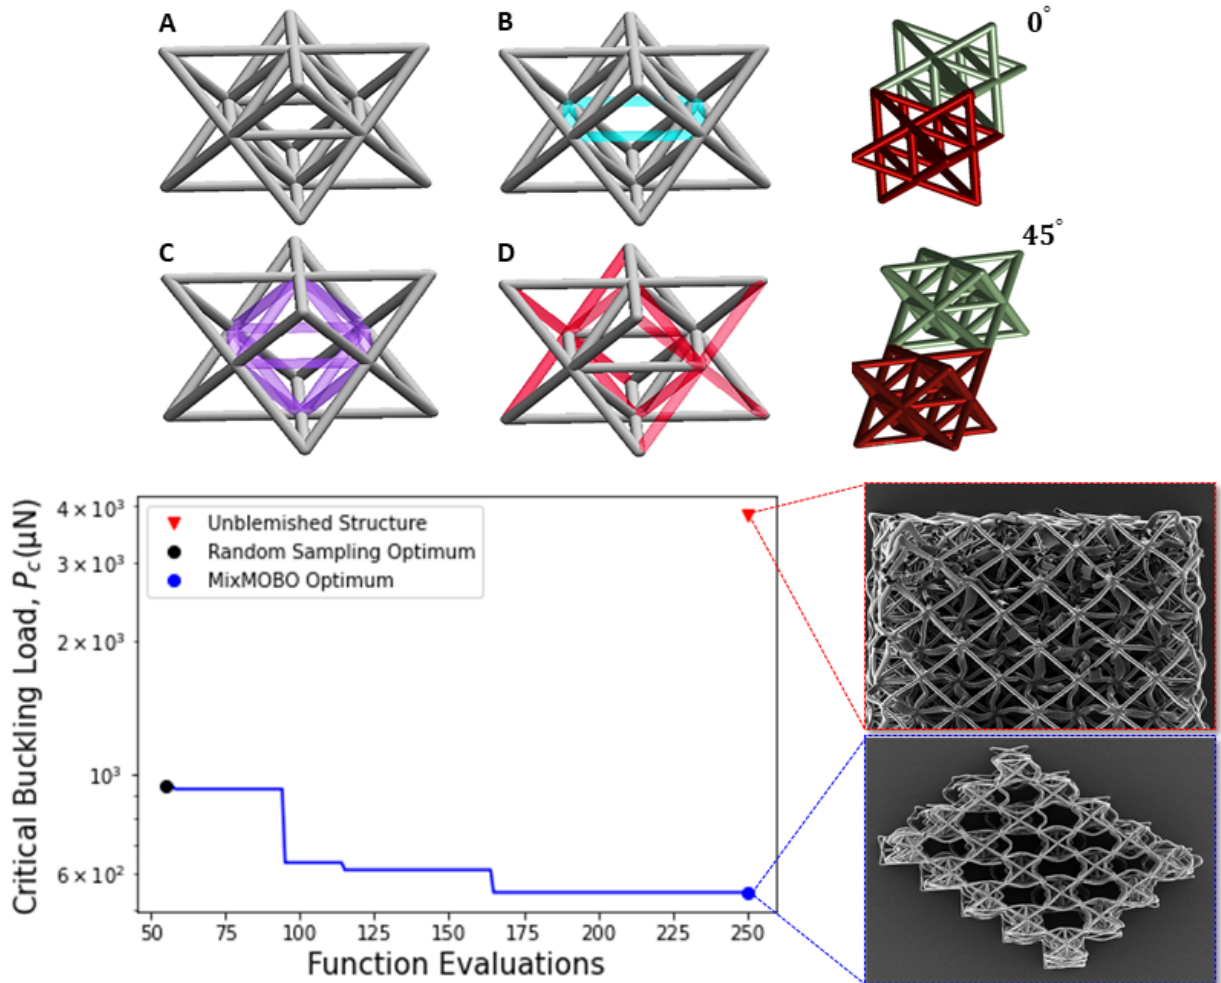


Figure 1.3: *Top Left:* The 4 unit cells, labelled **A** – **D**. *Top Right:* The 2 orientations in which they can be joined. *Bottom Left:* Optimization results using MixMOBO. *Bottom Right:* SEM images of Unblemished and Optimum structures.

combinations of its 17 categorical inputs (one with 2 possible states, and the other 16 with 4 possible states). Our goal is to maximize the strain energy density of a microlattice structure. We maximize the strain energy density (which is extremely expensive to compute, even for one design) by minimizing the buckling load P_c , while maintaining the lattice’s structural

integrity and stiffness before fracturing. Minimizing P_c (a proxy for maximizing the strain energy density by instigating buckling which leads to the densification of the deformed lattice members) is a more computationally tractable cost function to evaluate (but, it is still expensive and involves solving a numerical finite element code for each evaluation of the cost function.) The manufacturing and testing details of our methodology are included in Vangelatos, Sheikh, Marcus, Grigoropoulos, Lopez, Flamourakis, and Farsari [2], which focuses on the material aspects of the problem.

The design space consists of choosing one of four possible unit cells (shown in the upper left of Fig. 1.3, each with one or more defects (shown in color) in them, at each of the 16 independent lattice sites) creating 16 of the categorical inputs with 4 possible values; and the choice of whether the cells are connected along their faces or along their edges on 45°-diagonals (shown in the upper right panel of Fig. 1.3) creating the 17th categorical input with 2 possible values.

The minimization of P_c using MixMOBO was initialized with 50 random structures and the evaluation budget, including initial samples, was set at 250. The algorithm achieved a 42% improvement in the P_c of the lattice structure over the best structure obtained with the first 50 random samples (Figure 1.3). The optimal microlattice obtained using P_c as a proxy with MixMOBO has an experimentally measured normalized strain energy density that is 12,030 times greater than that of the unblemished microlattice structure with no defects that is cited in the literature to have the best strain energy density [66], a 4 orders of magnitude improvement. Table 1.1 shows the properties of the fabricated and experimentally measured design created with MixMOBO. The choices of the units cells in the optimally designed lattice that were determined by MixMOBO are not intuitive and have no obvious pattern or structure. Images of our optimized structures using Helium Ion Microscopy (HIM) 1.4, shows the comparison of the Unblemished structure from literature with our MixMOBO Optimal structure, before and after loading. It is evident that the MixMOBO Optimal structure due to its densification mechanism, can handle much higher loads without breaking [2].

1.8 Conclusions

The existing optimization literature does not offer an algorithm for optimizing multi-objective, mixed-variable problems with expensive black-box functions. We have introduced Mixed-variable Multi-Objective Bayesian Optimization (MixMOBO), the first BO based algorithm for optimizing such problems. MixMOBO is agnostic to the underlying kernel. It is compatible with modified kernels and other surrogate methods developed in previous studies for mixed-variable problems. MixMOBO allows for parallel batch updates without repeated evaluations of the surrogate surface, while maintaining diversification within the solution set. We presented the Hedge Multi-Objective (HedgeMO) algorithm, a novel Hedge strategy for which regret bounds hold for multi-objective problems. A new acquisition function, Stochastic Monte-Carlo (SMC) was also proposed as part of the HedgeMO portfolio. MixMOBO and HedgeMO were benchmarked and shown to be significantly better on a variety of test




Figure 1.4: HIM images of the loaded and unloaded unblemished and optimum structures. (a) Unloaded Unblemished structure (b) Unblemished structure after loading, showing severe fracture and collapse of many beam members. (c) Focused image revealing several fractured beams and the internal collapse of the upper layer that subsequently instigated the accumulation of damage in the underlying layers. (d) Unloaded MixMOBO Optimum structure. (e) MixMOBO Optimum structure after the structure was subjected to the same maximum compressive load as the structure shown in (b). Unloading of the optimum structure showed only excessive plastic deformation without catastrophic collapse and the manifestation of the buckling mode. (f) Focused revealing the effect of buckling that led to deformation but no fracture due to the occurrence of densification. (g) Side view of the unloaded optimum structure. (h) Side view of the unloaded optimum structure revealing that fracture was inhibited throughout the structure due to the densification precipitated by the low critical buckling load. **Scale:** Each scale bar is equal to $10^{-5}m$.

problems compared to existing mixed-variable optimization algorithms. MixMOBO was then applied to the real-world optimization of an architected micro-lattice, and we increased the structure’s strain-energy density by 10^4 compared to existing Unblemished structures in the literature reported to have highest strain energy density. Our future work entails further testing multi-objective and ‘ Q -batch’ settings. We have also applied MixMOBO for optimization of draft-tubes for hydrokinetic turbines [4] and Cauchy symmetric meta-material structures and are currently applying it for optimization of vertical-axis wind turbines.

1.9 Additional Details for MixMOBO

Benchmark Test Functions

In this section, we define the benchmark test functions, all of which are set to be maximized during our optimizations.

Contamination Problem

The contamination problem was introduced by Hu, Hu, Xu, Wang, and Cao [51] and is used to test categorical variables with binary categories. The problem aims to maximize the reward function for applying a preventative measure to stop contamination in a food supply chain with D stages. At each i^{th} stage, where $i \in [1, D]$, decontamination efforts can be applied. However, this effort comes at a cost c and will decrease the contamination by a random rate Γ_i . If no prevention effort is taken, the contamination spreads with a rate of Ω_i . At each stage i , the fraction of contaminated food is given by the recursive relation:

$$Z_i = \Omega_i(1 - w_i)(1 - Z_{i-1}) + (1 - \Sigma_i w_i)Z_{i-1} \quad (1.5)$$

here $w_i \in 0, 1$ and is the decision variable to determine if preventative measures are taken at i^{th} stage or not. The goal is to decide which stages i action should be taken to make sure Z_i does not exceed an upper limit U_i . Ω_i and Σ_i are determined by a uniform distribution. We consider the problem setup with Lagrangian relaxation [29]:

$$f(\vec{w}) = - \sum_{i=1}^D \left(cw_i + \frac{\rho}{T} \sum_{k=1}^T 1_{\{Z_k > U_i\}} \right) - \lambda \|\vec{w}\|_1 \quad (1.6)$$

Here violation of $Z_k < U_i$ is penalized by $\rho = 1$ and summing the contaminated stages if the limit is violated and our total stages or dimensions are $D = 21$. The cost c is set to be 0.2 and $Z_1 = 0.01$. As in the setup for [29], we use $T = 100$ stages, $U_i = 0.1$, $\lambda = 0.01$ and $\epsilon = 0.05$.

Encrypted Amalgamated

Analytic test functions generally cannot mimic mixed variables. To map the continuous output of a function into N discrete ordinal or categorical variables, the continuous range of the test

function's output is first discretized into N discrete subranges by selecting $(N - 1)$ break points, often equally spaced, within the bounds of the range. Then, the continuous output variable is assigned the integer round-off value of the subrange defined by its surrounding pair of break points. If necessary, the domain of the test function's output is first mapped into a larger domain so that each subrange has a unique integer value. To mimic ordinal variables, we are done, but for categorical variables, a random vector for each categorical variable is then generated which scrambles or 'encrypts' the indices of these values, thus creating random landscapes as is the case for categorical variables with a latent space. The optimization algorithm only sees the encrypted space and the random vector is only used when evaluating the black-box function.

We also define a new test function that we call the *Amalgamated function*, a piece-wise function formed from commonly used analytical test functions with different features (for more details on these functions we refer to Tušar, Brockhoff, and Hansen [52]). The Amalgamated function is non-convex and anisotropic, unlike conventional test functions where isotropy can be exploited.

For $i = 1 \dots n$, $k = \text{mod}(i - 1, 7)$:

$$f(\vec{w}) = \sum_{i=1}^D g(w_i) \quad (1.7)$$

where

$$g(w_i) = \begin{cases} \sin(w_i) & \text{if } k = 0, w_i \in (0, \pi) \\ -\frac{w_i^4 - 16w_i^2 + 5w_i}{2} & \text{if } k = 1, w_i \in (-5, 5) \\ -(w_i^2) & \text{if } k = 2, w_i \in (-10, 10) \\ -[10 + w_i^2 - 10\cos(2\pi w_i)] & \text{if } k = 3, w_i \in (-5, 5) \\ -[100(w_i - w_{i-1}^2)^2 + (1 - w_i)^2] & \text{if } k = 4, w_i \in (-2, 2) \\ \text{abs}(\cos(w_i)) & \text{if } k = 5, w_i \in (-\pi/2, \pi/2) \\ -w_i & \text{if } k = 6, w_i \in (-30, 30) \end{cases} \quad (1.8)$$

To create the Encrypted Amalgamated function, for categorical and ordinal variables, equally spaced points are taken within the bounds defined above. For our current work, we use a $D = 13$ with 8 categorical and 3 ordinal variables with 5 states each, and 2 continuous variables.

NK Landscapes

NK Landscapes were introduced by Kauffman and Levin [53] as a way of creating optimization problems with categorical variables. N describes the number of genes or number of dimensions D and K is the number of epistatic links of each gene to other genes, which describes the

‘ruggedness’ of the landscape. A large number of random landscapes can be created for given N and K values. The global optimum of a generated landscape for experimentation can only be computed through complete enumeration. The landscape cost for any vector is calculated as an average of each component cost. Each component cost is based on the random values generated for the categories, not only by its own alleles, but also by the alleles in the other genes connected through the random epistasis matrix, with K probability or ruggedness. A $K = 1$ ruggedness translates to a fully connected genome.

The NK Landscapes from Kauffman and Levin [53] were formulated only for binary variables. They were extended by Li, Emmerich, Eggermont, Bovenkamp, Bäck, Dijkstra, and Reiber [54] for multi-categorical problems, which is the formulation we use. Details of the NK Landscape test-functions we use can be found in Li, Emmerich, Eggermont, Bovenkamp, Bäck, Dijkstra, and Reiber [54]. For the current study, we use $N = 8$ with 4 categories each and ruggedness $K = 0.2$.

Rastringin

Rastringin function is a commonly used non-convex optimization function [52] with a large number of local optima. It is defined as:

$$f(\vec{w}) = -[10 + w_i^2 - 10\cos(2\pi w_i)], \quad w_i \in (-5, 5) \quad (1.9)$$

We use $D = 9$ for testing with 6 ordinal with 5 discrete states and 3 continuous variables. The ordinal variables are equally spaced within the bounds.

Encrypted Syblinski-Tang

We use the Syblinski-Tang function [52], an isotropic non-convex function. The function is considered difficult to optimize because many search algorithms get ‘stuck’ at a local optimum. For use with categorical variables, we encrypt it as described previously. The Syblinski-Tang function, in terms of input vector \vec{w} , is defined as:

$$f(\vec{w}) = -\frac{\sum_{i=1}^D w_i^4 - 16w_i^2 + 5w_i}{2}, \quad w_i \in (-5, 2.5) \quad (1.10)$$

For the current study, this function was tested with $D = 10$ categorical variables and 5 categories for each variable.

Encrypted ZDT6

ZDT benchmarks are a suite of multi-objective problems, suggested by Zitzler, Deb, and Thiele [55], and most commonly used for testing such problems. We use $ZDT6$, which is non-convex and non-uniform in its parameter space. We again modify the function by

encrypting it to work with categorical problems. *ZDT6* is defined as:

$$\begin{aligned}
 f_1(\vec{w}) &= \exp(-4w_1)\sin^6(6\pi w_1) - 1 \\
 f_2(\vec{w}) &= -g(\vec{w}) \left[1 - (f_1(\vec{w})/g(\vec{w}))^2\right] \\
 g(\vec{w}) &= 1 + 9 \left[\left(\sum_{i=2}^D w_i \right) / (n - 1) \right]^{1/4}
 \end{aligned} \tag{1.11}$$

Here $w_1 \in [0, 1]$ and $w_i = 0$ for $i = 2, \dots, D$. The function was tested for $D = 10$ with 5 categories each. We note that to evaluate the performance of MixMOBO, we compared it against the NSGA-II variant [44] that can deal with mixed variables (by running *ZDT4* in a mixed variable setting and *ZDT6* with categorical variables). No encryption is necessary for GAs. GAs required, on average, 10^2 more function calls compared to MixMOBO.

Chapter 2

Systematic Design of Cauchy Symmetric Structures Through Bayesian Optimization

2.1 Abstract

Using a new Bayesian Optimization algorithm to guide the design of the lattice structure of mechanical metamaterials, we design nonhomogeneous 3D structures possessing the Cauchy symmetry, which dictates the relationship between continuum and atomic deformations. The recent efforts to merge optimization techniques with the design of mechanical metamaterials has resulted in a concentrated effort to tailor their elastic and post elastic properties. Even though these properties of either individual unit cells or homogenized continua can be simulated using multi-physics solvers and well established optimization schemes, they are often computationally intensive and require many design iterations, rendering the validation stage a significant obstacle in the design of new metamaterial designs. This study aims to provide a framework on how to utilize miniscule computational cost to control the elastic properties of metamaterials such that specific symmetries can be accomplished. Using the Cauchy symmetry as a design objective, we engineer structures through the strategic arrangement of 5 different unit cells in a $5 \times 5 \times 5$ cubic symmetric microlattice structure. This lattice design, despite constituting a design space with 5^{10} 3D lattice configurations, can converge to an effective solution in only 69 function calls as a result of the efficiency of the new Bayesian optimization scheme. To validate the mechanical behavior of the design, the lattice structures were fabricated using multiphoton lithography and mechanically tested, revealing a close correlation between experiments and simulated results in the elastic regime. Ultimately, a similar methodology can be utilized to design metamaterials with other material properties of interest, aspiring to control properties at different length scales, an endeavor that requires inordinate computation cost.

2.2 Introduction

The continued development of advanced 3-dimensional manufacturing processes such as multiphoton polymerization has enabled the manufacture of complex geometries with sub-micron resolution at increasingly fast rate [67, 68, 69]. Furthermore, merging spatial light modulation technologies with multiphoton lithography has resulted in the production of larger arrays with the same feature resolution [70]. Consequently, these techniques have provided the avenue to design scalable architectures that yield bulk material properties that exceed those of natural materials [71, 72]. These architected materials, can be employed for various engineering applications such as ultralight [73, 74, 75, 76], ultraprecise [77, 78], reconfigurable [79], high energy absorption materials [60, 2], as well as different engineering domains such as wave mechanics [80] and optics [81]. Furthermore, the fabrication of flexible structures with tailored mechanical properties has proved important design bio inspired designs, inducing enhanced properties on significantly weaker bulk materials [82, 83]. Lessons from nature can be derived from the most unexpected cases, such as the flower beetle *Torynorrhina flammaea* [84] or the deep sea sponge *Euplectella aspergillum*, showing how the "optimization" of nature can lead to complex optomechanical designs that resemble the features and patterns of optomechanical or purely mechanical metamaterials [85, 86].

Evidently, from a structural standpoint, the lattice nature of microscale metamaterial structures has led to the predominant design strategy of different truss-like structures that mimic the crystalline structure and defect formation of metals and alloys [82, 57, 87, 88]. These architected lattice, plate or triply periodic minimal surface structures have been shown to provide high stiffness and increased strength, at extremely low relative densities [89, 90, 91]. This is also reflected in the plethora of results aiming specifically in the control of elastic properties, as they are determined by the stiffness tensor [92, 93, 94]. The vast majority of the literature focuses in the control of properties such as isotropy, leading to mechanical behavior independent of direction [95, 96, 97] and auxeticity, a riveting property leading to negative Poisson's ratio and resilience to collapse mechanisms such as necking and barrel shape formation [98, 99].

Apart from design approaches such as the addition of beam elements, and the reconfiguration of beam sizes [66, 57, 63, 56, 2], accelerating the design process and circumventing the inherent challenges of continuum modeling that drives this endeavor [100, 101] has been accomplished through optimization. Several studies have attempted employing optimization techniques to obtain the elastic properties of the structure. Techniques such as topology optimization [94, 93, 102], artificial neural networks [103, 104] and machine learning [105, 106, 107, 108] have risen in metamaterial design and investigation of controlled elastic properties. Nevertheless, in these publications, the required number of evaluations to obtain the optimum is thousands, tens of thousands or even millions of data to explore the design space [93], creating a high computational cost. As a consequence, the number and type of design variables (continuous, discrete, qualitative and ordered) has to be limited and in many cases isolate the study to 2D or uniform homogeneous 3D structures to sustain a feasible computational budget. Some techniques with low computational budgets for meta-material design have also been

proposed [109, 110, 111, 112, 113]. In our previous study we explored how these challenges can be successfully surpassed by the development of a novel Bayesian Optimization scheme [1]. We combined this algorithm with a new approach to model different building blocks in the metamaterial structure as states, which are discrete, qualitative design variables [2, 114, 66]. This approach required only 250 cost function evaluations, and resulted in defected mechanical metamaterials with architected defects, facilitating enhanced mechanical behavior which was orders of magnitude higher than the thoroughly studied monolithic metamaterial structures [115, 115, 60]. This study aims to expand use our improved Mixed-variable Multi-Objective Bayesian Optimization (MixMOBO) framework [1] on the following fronts: (a) decrease the computational cost by an order of magnitude even with a larger number of states, rendering such exorbitant problems more accessible even with meagre computational sources using MixMOBO and (b) design lattice structures that are nonhomogeneous in every direction, expanding the design of arbitrary metamaterials [116] to complex 3D problems.

Investigating the controlled elastic behavior of imperfect and nonhomogeneous lattices will elucidate a new approach to address such problems. However, instead of the thoroughly studied isotropy and auxeticity, this study will explore another intriguing property derived from crystalline materials, namely the Cauchy–Born Rule [117, 118, 119]. This model shows that elastically deformed states of the lattice model at the localized level are closely approximated by solutions of the continuum model. Even though this rule is based on the absence of lattice vibrations and ionic polarization, limiting its validity to describe complex phenomena in the lattice, this kinematic theory has been thoroughly employed as the constitutive behavior of continuum regions in multi-scale models. Thus the fine scale is proposed to depict the real behavior of crystalline structure when the continuum description is incapable to accomplish this [120]. As it will be explained later, this rule is referred with respect to the stiffness tensor of the material as "Cauchy Symmetry". Therefore, designing mechanical metamaterials with controlled auxeticity will provide new metamaterial families that can be used for homogenized continua and constitutive models and will enable the fabrication of structures with homogeneous behavior.

Thus, using Cauchy Symmetry as an example of a non-intuitive design objective, we present a method to systematically tailor the mechanical behavior of architected microstructures. By building a $5 \times 5 \times 5$ microlattice structure out of five different unit cells, we yield a large combinatorial design space. Such a design space with combinatorial or categorical variables, with a expensive-to-evaluate black-box cost function that would be prohibitively expensive to optimize with conventional optimization techniques. We use our Mixed-variable Multi-Objective Bayesian Optimization (MixMOBO) algorithm [1] to optimize this large and expensive combinatorial design space. Through finite element analysis simulations, we study the mechanical response and calculate the elastic constants of the lattice which are the data required by the algorithm. To validate our results, we employ multiphoton lithography to fabricate the structures resulted by the optimization, which enables the design of complex nonuniform 3D structures. Finally, with in-situ SEM-microindentation experiments, we evaluate the validity of the optimization algorithm. Our results illuminate how complex 3D structures can be successfully optimized with a low computational budget and can lead to

new types of metamaterials that can be employed for continuum modeling and controlled elastic behavior.

2.3 Problem Setup

Modeling of Cauchy-Symmetric Structures

The Cauchy-Born rule relates the position of atoms in a crystal to the overall strain of the medium [121]. Assuming that the potential energy of the lattice is a function of the distance between the atoms, or in lattice metamaterials, the nodes, through the second derivative on the potential energy with respect to the strain, the elastic stiffness tensor C can be obtained. Databases possessing millions of geometries to describe the mechanical properties of truss metamaterials as a function of the lattice node locations have been reported [93], indication that such an assumption can be valid as long as higher gradient elasticity is not dominant [101]. However, the particular form of the potential energy in the Cauchy-Born rule can lead to additional symmetries in the stiffness tensor, defined as “Cauchy Symmetries”. These additional symmetries reduce the number of independent elastic constants in the structure, leading in some cases to unphysical results as in a triclinic material. However, while the number elastic independent variables of bulk materials cannot be physically reduced, the arrangement of beam members in the metamaterial structure can result in the control of the elastic constants [122].

A particular case of interest can be metamaterials possessing cubic symmetry, since the vast majority of thoroughly studied structures such as the Diamond, the Octet Truss, the Kelvin have cubic symmetry [123]. In this reduced form, the effective stiffness tensor has only three independent elements: C_{11} , C_{12} , and C_{44} , where C_{11} and C_{12} are stiffness components of normal stress and C_{44} specifies the shear modulus of a cubically symmetric structure. This correlation between stiffness elements and loading mode are conveyed in Figure 2.1. However, Cauchy-Symmetry is obtained in a cubic material when $C_{12} = C_{44}$. This requirement will set the framework to define the problem and optimize it as it will be discussed in the next section.

Cauchy symmetry Problem Setup

The objective of this work is to design a nonmonolithic microlattice composed of discrete unit cells that is Cauchy Symmetric. Figure 2.2 depicts the process used to identify and validate a near-optimal structure. Throughout this process we use five different unit cells, shown in Figure 2.2a, labeled A, B, C, D or E. For simplification, we choose unit cells that have cubic symmetry so that they can be systematically joined together to create a bulk microlattice. Each unit cell is $10 \mu\text{m} \times 10 \mu\text{m} \times 10 \mu\text{m}$, which is compatible with the dimensions realizable by our multiphoton lithography (MPL) apparatus. These unit cells are positioned in a $5 \times 5 \times 5$ array. We strategically chose a layout scheme of unit cells depicted

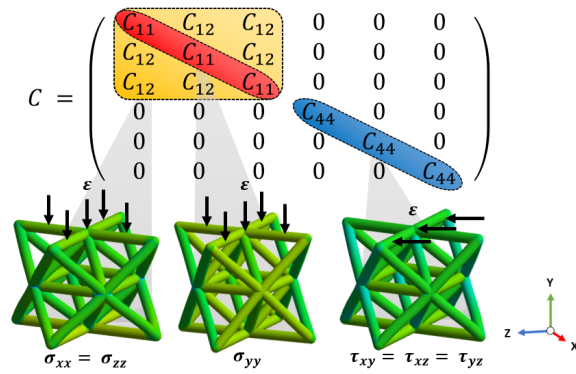


Figure 2.1: The stiffness matrix of a cubic symmetric unit cell can be described by three independent stiffness components which corresponded to the stress components orthogonal to the direction of strain (left) the stress components in the direction of loading (middle) and the shear modulus (right).

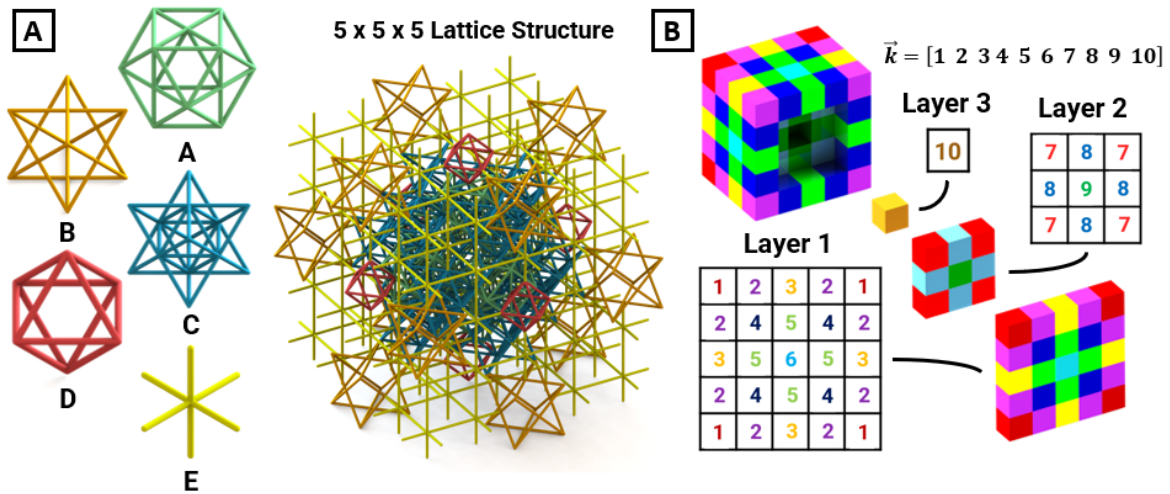


Figure 2.2: (a) Five types of cubic unit cells, labeled A-E, used to construct a 5 x 5 x 5 lattice. Unit cells are all face-centered cubic geometry to ensure compatible connectivity within the lattice (b) Unit cells are placed together into three layers that are cubic symmetric. Unit cells of the same type are placed into one of 10 positions, labeled 1-10, which the k-vector is composed of. Layer 1, located on the outside of the structure is composed of 5 x 5 unit cells. Layer 2 is the middle layer and is composed of 3 x 3 unit cells, and Layer 3 is a singular unit cell located in the center of the microlattice.

in Figure 2.2b. The micro lattices consist of an outer "shell" (Layer 1, 5×5 unit cells), a middle layer (Layer 2, 3×3 unit cells), and the inner core unit cell (Layer 3, 1×1 unit cell). The ten positions, 1-10 in the different layers, which can be equipped with unit cells, are a result of ensuring that the boundary condition of cubic symmetry is fulfilled. Accordingly, the same unit cells must be positioned at locations with the same letter.

This layout scheme was strategically chosen to provide two key benefits. The first benefit is that it roughly imposes isotropic bulk material behavior. In doing so, it reduces the stress tensor to three independent stiffness values which simplifies the number of simulations and mechanical tests required to probe the stiffness values. The second benefit is that it substantially reduces the number of possible combinations of lattices. With this restriction, the design space reduces from 2.3×10^{87} states to 9765625. Even though we search a design space that likely has multiple sufficiently low values of the cost function, the reduction in the dimensionality of the design space should still significantly reduce the number of function evaluations required. In addition, it is important to state that even though symmetries in the lattice structure will also decrease the design space, the optimization algorithm determines the overall number of required evaluations by the overall size of the design space.

Figure 2.2c shows the optimization utilized to find a Cauchy-Symmetric microlattice. Fabricating and physically testing structures is a labor-intensive process with a non-insignificant amount of noise, so finite element analysis is used as a proxy for calculating stiffness components for the cost function in our MixMOBO algorithm. The principal improvement of this algorithm is the utility of the hedge strategy which hedges acquisition functions. This approach uses a set of acquisition functions instead of a single one, rendering the algorithm much more efficient and versatile as it was been reported before [124]. To converge on an a near-optimal solution we employ finite element code to compute the static behavior of each lattice in its elastic domain. Each microlattice contains roughly 100,000-900,000 nodes, and the optimization of lattice properties requires that this computation be done repeatedly. In turn, even this more time-effective method remains prohibitively expensive to search the entire design space or employ other optimization techniques that still require thousands of function calls. Once our algorithm converged on an acceptable solution, we fabricated and mechanically tested (further information about the mechanical testing process are provided in Materials and Methods).

The material that was used was the hybrid organic-inorganic Zr-DMAEMA [(2-dimethylaminoethyl) methacrylate] (further details about the fabrication and the material properties are provided in Materials and Methods). Based on previously conducted work [114, 66], the stiffness values measured by mechanical testing can only be reliably related within the elastic regime, since the photoresist material is highly sensitive in the plastic domain due to variations in the fabrication parameters. To incorporate the noise caused by fabrication imperfections and the experimental measurement, a 0.005 noise variance (7% standard deviation) is built into the test function based on previously reported results [125]. Mechanical testing and structural analysis are conducted inside a scanning electron microscope (SEM).




Figure 2.3: Process diagram beginning with the selection of 50 initial random microlattices. Afterwards the process iteratively loops between (i) application of the MixMOBO algorithm to choose new parameters and (ii) the modeling, simulation, and evaluation of the Cauchy-symmetry cost function. Once a microlattice with sufficiently small cost function value is found, the iterative loop is exited and the optimal design is fabricated and tested to validate its properties.

Bayesian Optimization

Optimization for expensive black-box functions involving categorical or mixed variables is an area of active research. For such hard-to-evaluate problems, Bayesian optimization (BO) has proven to be efficient [18, 126] using small number of function calls compared to other optimization strategies. A large range of these problems require optimizing combinatorial or categorical design spaces. Architected material design [5, 6, 7, 8, 9, 2], hyper-parameter tuning for machine learning algorithms [10, 11, 12], drug design [13, 14], controller sensor placement [15] pose such problems with expensive black-box function and Bayesian optimization is a natural candidate for their optimization. However, for problems such as design of architected materials, the design space can include categorical or ordinal variables [2].

Much research has gone into Bayesian optimization (BO) for continuous design spaces with minimum number of expensive function calls [16, 17, 18]. Despite the success of continuous Bayesian optimization strategies, combinatorial problems remain an area of open research. The inherent continuous nature of Gaussian processes (GP) makes dealing with categorical variables challenging.

In our past work, detailed in [1], we developed a Mixed variable, Multi-Objective Bayesian Optimization (MixMOBO) algorithm, a generalized framework that can deal with categorical problems and can optimize a noisy black-box function with a small number of function calls. We applied a categorical variable Bayesian optimization algorithm (previously known as Evolutionary Monte Carlo Simulations, EMCS) for the design of an architected meta-material to have the maximum strain-energy density [2]. The fabrication and testing of this new material shows that it has a normalized strain energy density that is 10^4 times greater than existing unblemished microlattice structures in literature [2], generalized mixed-variable and multi-objective BO algorithm, for design of an Cauchy-Symmetric structure.

2.4 Materials and Methods

FEA Design

All FEA simulations were performed with the multipurpose finite element code ANSYS (Workbench 18.0). The following beam lattice properties were used in the FEA: 1.281 elastic modulus, 0.4999 Poisson's ratio. Only elastic material properties were used. The structures were meshed with 10-node, tetrahedral finite elements. Each structure roughly contained between 100,000 and 900,000 nodes. The average element size was selected by conducting a mesh convergence study. For uniaxial compressive testing, an infinitely stiff plane attached to the top surface of the structure was displaced downwards and C_{12} was calculated by averaging σ_y across all elements. For shear testing and a vertice on the structure was displaced laterally and C_{44} was calculated by averaging τ_{xy} across all elements. In both cases, the bottom face of the structure was attached to the substrate and fully constrained. All of the structures were designed using SOLIDWORKS. This methodology has been employed before for the precise calculation of the polar stiffness map of structure, leading to a close correlation between theoretical and experimental results [127, 128].

Mechanical Testing

Compression and shear tests were performed in situ with a nanoindentation apparatus (PI 87 SEM PicoIndenter, Hysitron) mounted inside the chamber of a scanning electron microscope (SEM) (FEI Scios 2, Thermo Fisher Scientific). A molybdenum flat tip indenter (model 72SC-D3/035 (407A-M) Probing Solutions, Inc.) with a diameter of 70 μm is applied in all mechanical tests. The glass substrates on which the microlattices are fabricated are fixed onto an SEM pin stub mount (TED PELLA) with PELCO $\text{\textcircled{R}}$ Pro C100 Cyanoacrylate Glue (TED PELLA). Test structures for compression testing are aligned such that the top face contacts the indenter tip while the side profile is imaged by the electron beam. Test structures for shear testing are aligned such that a side face with a flat bar attached on the side contacts the indenter and the top is imaged. Each structure was deformed at a rate of 250 nm/s, and to ensure repeatability, 4 tests were performed on each structure.

Fabrication

The microlattice structures were fabricated with a hybrid organic-inorganic resin Zr-DMAEMA (30 wt%). The resin is composed of 70 wt% zirconium propoxide and 10 wt% (2-dimethylaminoethyl) methacrylate (DMAEMA) (Sigma-Aldrich). The structures for mechanical testing were fabricated by sub-micron resolution direct femtosecond laser writing, which uses MPL and the aforementioned photoresist for high-resolution fabrication. The basic optical system consists of a FemtoFiber pro NIR laser emitting 780 nm wavelength, with a pulse width of 100 fs, and a repetition rate of 80 MHz. The local photopolymerization of the photosensitive material was accomplished with a 100 \times microscope objective lens (Plan-Apochromat 100 \times /1.40 Oil M27,

Zeiss). The laser output energy for the fabrication was measured before the objective lens at 4.2 mW, and the scanning speed used was set to 20 $\mu\text{m}/\text{s}$. The resin sample is attached to a stage with both piezo and servo elements with movements in the XYZ directions.

MixMOBO

For minimizing the objective, we pose the categorical variable problem as:

$$\vec{w}_{opt} = \operatorname{argmin}_{\vec{w} \in \mathcal{W}}(f(\vec{w})) \tag{2.1}$$

Here $f(\vec{w})$ is the objective to be minimized, and $\vec{w} \in \mathcal{W}$ is a categorical variable vector, defined as $\vec{z} = [w_1, \dots, w_m]$ for m total variables. \mathcal{W} is the combinatorial space. Each categorical variable takes a value $w_j \in \{C_1, \dots, C_j\}$ from C_j unordered categories (that cannot, by definition, be ordered on the real-number line).

For practical engineering problems such as design of architected meta-materials, $f(\vec{w})$ can be extremely expensive to evaluate. For combinatorial problems, gradient based algorithms cannot be used. The difficulty for optimization is also exacerbated due to the large number of possible combinations, approximately 1 billion possible combinations for our design space, making such problem intractable to optimize using conventional optimization techniques. Bayesian optimization is a sequential optimization technique for optimizing expensive noisy black-box function f with minimum number of function calls. For every iteration, a surrogate model g , is fit over the data set $\mathcal{D} = \{(w_1, f(w_1)), \dots, (w_i, f(w_i))\}$. Here i is the total number of points evaluated until the i^{th} iteration. Gaussian processes (GP) are usually used as the surrogate models. Once the surrogate surface has been fit to the data, the surrogate surface is explored to determine which point has the highest probability of finding the optimum point as well as improving the fit of the surrogate surface, exploitation and exploration respectively. An acquisition function is used to balance the exploitation and exploration trade-off and determine the next point \vec{w}_{i+1} for evaluation with f . \vec{w}_{i+1} is then evaluated using f and appended to the data set, $\mathcal{D} = \mathcal{D} \cup (w_{i+1}, f(w_{i+1}))$. The process is repeated until the evaluation budget for f or the global optimum is reached.




Figure 2.4: Benchmarks for MixMOBO.

Table 2.1: Benchmark Test Functions

Name	Objective Functions
Encrypt. Spherical	$f(\vec{w}) = w_i^2, w_i \in (-10, 10)$
Encrypt. Rastrigin	$f(\vec{w}) = [10 + w_i^2 - 10\cos(2\pi w_i)], w_i \in (-5.12, 5.12)$
Encrypt. Syblinski-Tang	$f(\vec{w}) = \frac{w_i^4 - 16w_i^2 + 5w_i}{2}, w_i \in (-5, 5)$
Encrypt. Amalgamated	$f(\vec{w}) = \sum_{i=1}^D g(w_i)$ $g(w_i) = \begin{cases} -\sin(w_i), & \text{if } k = 0, w_i \in (0, \pi) \\ \frac{w_i^4 - 16w_i^2 + 5w_i}{2}, & \text{if } k = 1, w_i \in (-5, 5) \\ w_i^2, & \text{if } k = 2, w_i \in (-10, 10) \\ [10 + w_i^2 - 10\cos(2\pi w_i)], & \text{if } k = 3, w_i \in (-5.12, 5.12) \\ [100(w_i - w_{i-1}^2)^2 + (1 - w_i)^2], & \text{if } k = 4, w_i \in (-2, 2) \\ - \cos(w_i) , & \text{if } k = 5, w_i \in (-\pi/2, \pi/2) \\ w_i, & \text{if } k = 6, w_i \in (-30, 30) \end{cases}$ $k = \text{mod}(i - 1, 7), i = 1, \dots, n$

Bayesian optimization has proven to be successful for optimizing expensive black-box functions with continuous design spaces. For combinatorial or mixed variable problems, conventional Bayesian optimization techniques cannot be used due to the smooth nature of GP. In our previous work [2], we introduced EMCS, a Bayesian optimization algorithm for combinatorial variables using Stochastic Monte-Carlo (SMC) acquisition function. In the current work, we use MixMOBO, a mixed variable algorithm we developed in [1], to optimize find a Cauchy-Symmetric structure within our design space. It also uses HedgeMO, a hedging strategy introduced as part of MixMOBO in [1], that uses a portfolio of acquisition functions rather than a single one. Hedge algorithms have proven to be efficient in dealing with a diverse set of problems since they don't rely on a single acquisition function which might affect the efficiency of the BO algorithm [43].

To test the efficacy of MixMOBO for our problem, we test MixMOBO on a range of test functions with the same design space as our problem, i.e. 10 categorical variables with 5

categories for each variable. These tests demonstrate the number of function calls that would be necessary to optimize a design space of this size. We run MixMOBO 5 times for each test function to ensure reproducibility with 50 initial random data points. We use Encrypted Amalgamated, Encrypted Rastringin, Encrypted Syblinski-Tang, and Encrypted Spherical as our test functions. These test functions are commonly used for performance benchmarking of optimization algorithms [52] and have been modified for testing categorical variables.

All test functions are defined to be minimized. Analytical test functions are generally not able to mimic mixed variables. We describe a method to ‘Encrypt’ the test functions, which is then suitable for mixed variable problems. The categorical variable dimension of the test function is first discretized into the required number of categories by equally spaced points within the bounds. A random vector for each categorical variable is generated which scrambles or ‘encrypts’ the indices of these values thus creating random landscapes as is the case for categorical variables with a latent space. The optimization algorithm only sees the encrypted space and the random vector is only used when evaluating the black-box function. For ordinal variables, the design space is discretized.

We also define a new test function that we call the *Amalgamated* function, a piece-wise function formed from commonly used analytical test functions with different features. For more details on these functions we refer to Tušar, Brockhoff, and Hansen [52]. Amalgamated function is non-convex and anisotropic, unlike conventional test functions where isotropy can be exploited. The other test functions used in Encrypted mode are commonly used for testing optimization algorithms [52] and are detailed in Table ??.

For categorical variables, equally spaced points are taken within the bounds defined above. In our current work, we use a $D = 10$ categorical variables with 5 states each similar to our lattice optimization problem. The results of our benchmarks are depicted in 4.8. The plots show Normalized Cost, defined as $(global\ optimum - current\ optimum)/(global\ optimum - random\ sampling\ optimum)$ versus number of black-box function evaluations, with the Normalized global minima at 0. The mean Normalized Cost of 5 runs for each test function along with the S.D. is plotted. For each test function run, global optimum is found within 150 function calls for each run, demonstrating the efficacy of our MixMOBO setup.

2.5 Results

Convergence of the Algorithm

The convergence of the algorithm is summarized in Figure 2.5. It is noted that despite the fact the design space is comprised of almost 10^7 structures, the algorithm converged after 69 function evaluations, which is at least 3 orders of magnitude lower compared to other optimization schemes as it was reported elsewhere [1]. The structure, as it is mathematically described in Figure 2.2b is shown in Figure 2.5 and has the following sequence $w_{opt} = [EEBEEDECAE]$.

It is noted that this structure does not have any specific pattern or layer arrangement that shows some periodicity or uniformity in specific directions. To experimentally validate the optimized structure, it was fabricated by the MPL apparatus provided in the previous section. Specifically, to control the uniformity of the cross section of the structures, the setup and fabrication protocol had the following modifications.




Figure 2.5: Optimization progress for Cauchy-Symmetric structures

Properties of Optimum Structure

To extract the elastic constants from the experimental results, the directional elastic properties were obtained following the methodology that is provided elsewhere [129]. Specifically, the direction Young’s Modulus, Poisson’s ratio and Shear Modulus are provided by the following equations:

$$E(\mathbf{n}) = 1/S'_{11} \tag{2.2}$$

$$G(\mathbf{n}, \mathbf{m}) = 1/S'_{44} \tag{2.3}$$

$$\nu(\mathbf{n}, \mathbf{m}) = -S'_{12}/S'_{11} \quad (2.4)$$

where \mathbf{n} is the direction normal to the loading, \mathbf{m} is the perpendicular direction in the shear plane, S'_{11} , S'_{12} and S'_{44} are the tensor products of the vectors \mathbf{n} and \mathbf{m} with the compliance tensor C^{-1} . Based on these definitions, for the direction $\mathbf{n} = [100]$ and $\mathbf{m} = [010]$, the ratio of the elastic constants C_{12} and C_{44} is given by the following equation:

$$C_{12}/C_{44} = \frac{F\varepsilon_{22}\varepsilon_{12}}{\varepsilon_{11}^2 T(1 - \varepsilon_{22}/\varepsilon_{11})(1 + 2\varepsilon_{22}/\varepsilon_{11})} \quad (2.5)$$

where F is the measured force during compression, T is the measured force during shear, ε_{11} is the strain in loading direction during compression, ε_{22} is the strain in the perpendicular direction during compression and ε_{12} is the strain during shear. To obtain the measurements for the aforementioned quantities, compression and shear test were conducted. The fabricated optimal structure is presented in Figure 2.6(a) (isometric view) and 2.6(b) (top view), captured in the Helium Ion Microscope, enabling high resolution imaging and large depth of focus to observe the internal nonuniformity.




Figure 2.6: (a) Orthogonal view of the Cauchy Symmetric structure taken in a helium ion microscope. (b) Top view of structure highlighting with a large depth of focus to resolve internal beam members. Scale bars are both $20 \mu m$.

Characteristic force-displacement curves for the two types of experiments are shown in Figure 2.7(a). Through the experimental curves and the video recordings provided in the supplementary material (Video A for compression and Video B for shear), the parameters of equation 2.5 can be obtained. A representative frame of the compression measurement and

a shear measurement are shown in Figure 2.7(b) and 2.7(c) respectively. For a set of shear tests and compression tests, the ratio C_{12}/C_{44} was found equal to 0.9752 ± 0.0035 , a result close to the theoretical estimate of the FEA simulations which is 1.




Figure 2.7: (a) Mechanical loading response curves of the optimal Cauchy Symmetric lattice structure measured using a PI-87 picoindenter. Points are added to both curves to identify slopes that can be employed for the validation of the cost function. (b) Representative frame of the compression measurement. The indenter applies a compressive load to the structure in the lateral direction. (c) Representative frame of the shear measurement. The indenter is attached to the bar that is fabricated on the edge of the structure, applying a shear load to the geometry. Scale bars are both $20 \mu\text{m}$.

2.6 Discussion

It was demonstrated how the control of the stiffness tensor of the lattice structure such that the Cauchy symmetry rule was accomplished using the MixMOBO scheme. The versatility and efficiency of the algorithm is reflected on the minuscule number of simulations required to obtain the optimum. Specifically, the optimum structure was obtained using 69 data points simulations (i.e. 207 linear elastic FEA simulations) to calculate the elastic behavior of the lattice. This is a consequence of the incorporation of the hedge strategy, which enabled the decrease of required simulations compared to previously conducted work [2] and two or three orders of magnitude less compared to other problems using other optimization schemes [130, 131, 132]. Furthermore, this algorithm is able to work with mixed variables (a mixture of categorical, discrete and continuous variables) and multi-objectives, expanding the framework of this work to different types of problems. Currently we are applying this for optimization of draft-tubes for hydrokinetic turbines and vertical-axis wind turbine shapes.

The structure possessing the Cauchy symmetry rule from a sample of approximately 10^7 geometries does not show any inherent pattern that could be predicted by intuition or purely mechanics principles. This result highlights the necessity of optimization for the design of nonmonolithic structures, but also the versatility of MixMOBO for black box and extortionate

problems. Future work should focus on the utility of mixed variables, such as continuous and discrete, which will lead to non monolithic structures that also possess nonuniform thickness [93].

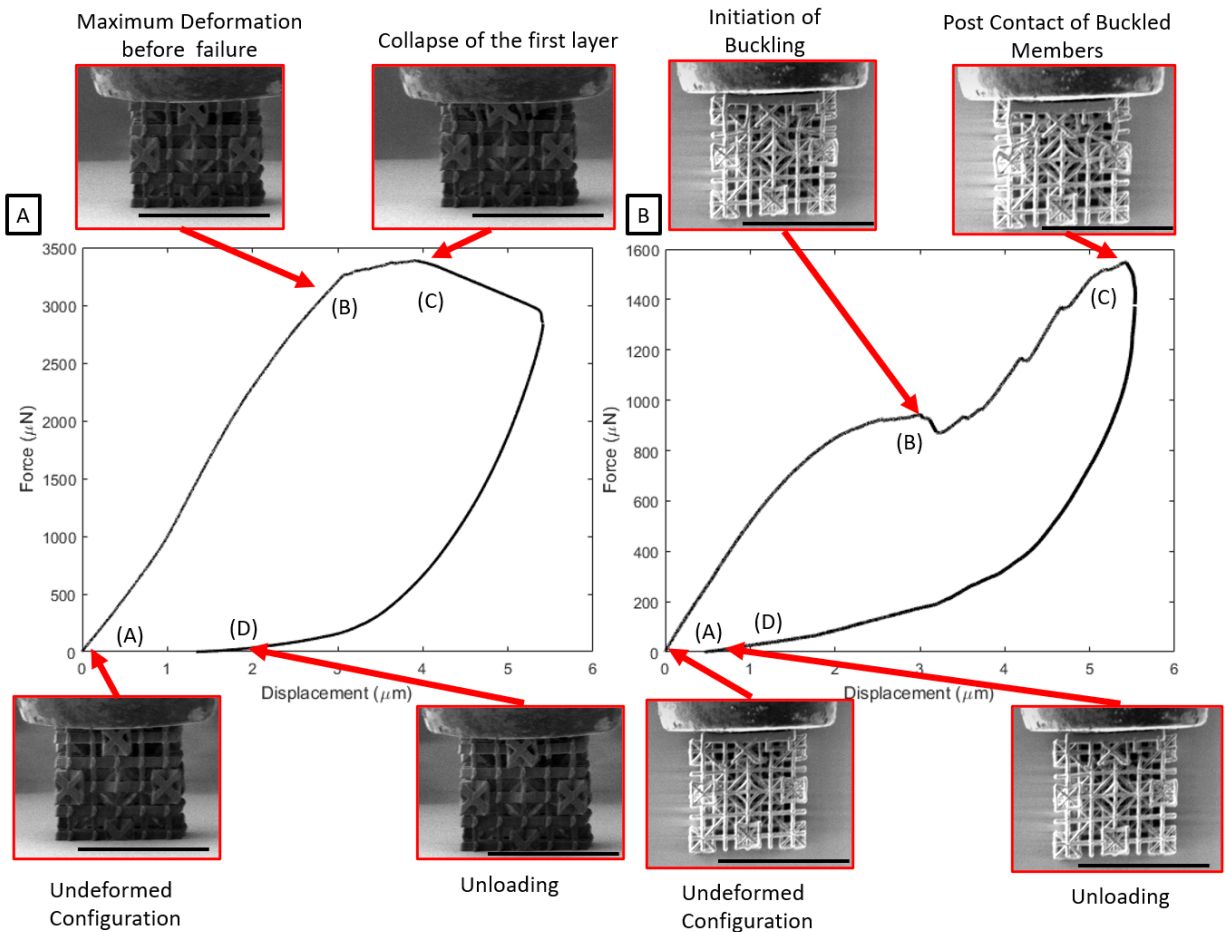


Figure 2.8: In situ analysis of the force displacement curves for compression and shear. (a) The response of the structure during compression shows a smooth linear response without discernible boundary effects (stage A-B), until the critical failure load is reached and the first layer collapses (stage C), leading to permanent deformation (stage D). (b) The response of the structure during shows the transition from linear behavior to instability and buckling of the lattice members (stage A-B), which result the initiation of microbuckling phenomena and densification, creating a serrated profile (stage B-C) without extensive permanent deformation (stage C-D). The black scale bar is $50 \mu\text{m}$.

Regarding the mechanical performance of the optimal structure, a juxtaposition of the force-displacement curve with the video caption also demonstrates the reverberations of Cauchy Symmetry to the mechanical performance. These results are depicted in Figure

2.8. In previous work the introduction of architected defects led to controlled densification and tailored buckling, but the small number of unit cells that can be realized by the MPL also leads to boundary effects [75, 87, 92]. However, a frame by frame comparison of the compression measurement shows that whole array of lattices follows the overall strain of the medium (Figure 2.8a), without substantial localised failure events as they had been observed before [66]. While this is one of the necessary assumptions for the Cauchy-Born rule to be valid in a crystalline material, the observation of this effect to the optimal array signifies the potential utility of such structures for the design of homogenised microscale lattices. In addition, during loading mode such as shear at the edge of the structure, localized buckling and post contact phenomena are shown (Figure 2.8b), leading to a serrated profile in the force-displacement curve and densification. The buckling deformation is dominant on $\pm 45^\circ$ plane. This same mechanism was observed under compression in lattice structures with architected defects or different types of unit cells [66, 2] and resembles the formation of mechanical twinning during shear that has been investigated when the Cauchy-Born rule is satisfied [133]. The consequence of this mechanism in shear loading mode can result in architected materials with improved tribological properties, enhanced post linear elastic mechanical performance and strain hardening [114].

Furthermore, future work should focus on the exploration of the design space with different types of unit cells. Different unit cells will significantly increase the design space. For instance, utilizing 6 or 7 types of different unit cells instead of 5 leads to a design space comprised of 60466176 and 282475249 lattice structures respectively. Moreover, increasing the size of the lattice increases the size of the design space. While a $5 \times 5 \times 5$ 3D cubic lattice has 10 design variables, it can be shown that the number of design variables increases each two steps. A $6 \times 6 \times 6$ array still has 10 design variables, while a $7 \times 7 \times 7$ and $8 \times 8 \times 8$ have 20 design variables. Utilizing 5 different types of unit cells will lead to a design space of approximately 9.53×10^{13} lattices. Therefore, different strategies should also be explored to decrease the computational cost for such design spaces and identify different approaches to define the design space. To this end, the generation of data-bases to easily extract specific geometries and use for the optimization algorithm would be a versatile tool to automate this process.

Future work should also focus on the effect of such lattice structures on the vibrational behavior of structure and the effect of the Cauchy symmetry on the dispersion curves of the lattice. This investigation will reveal how the tailored stiffness tensor will affect the wave propagation to the medium and control the energy absorption in the materials and create band-gaps. For the experimental validation of these effects, experimental techniques such as heterodyne interferometry and laser Doppler vibrometry can be employed to record a reflection of a laser beam through time and calculate the velocity of the vibration through the medium [134]. Finally, these results illuminate the necessity of larger scale printing that can capture the required feature resolution of the lattice members fabricated by MPL and can facilitate a tailored mechanical performance in bulk material. To achieve this, more advanced additive manufacturing techniques such as projection lithography [70] can increase the printing area and also reduce the fabrication time to efficiently produce mesoscale samples.

2.7 Conclusions

In this work, a new Bayesian optimization scheme, MixMOBO, was used for the design of nonmonolithic architected materials described by discrete and qualitative design variables. Utilizing 69 data points, the optimum of the structure that possesses Cauchy symmetry was obtained in a design space of 10^7 structures. Utilizing MPL and in-situ mechanical testing for compression and shear, it was revealed that the the structure's mechanical response shows the formation of shear planes and controlled buckling and a smooth strain, alluding how the assumptions of Cauchy-Born rule in the crystal lattice can be harnessed for the controlled mechanical performance of mechanical metamaterials. These results elucidate how the mitigation of the computational cost can pave the way for the exploration of exorbitant design spaces and the strategic design of architected nonmonolithic materials with tailored mechanical performance.

Chapter 3

Airfoil Optimization using Design-by-Morphing

3.1 Abstract

We present Design-by-Morphing (DbM), a novel design methodology to create a search space for topology optimization of 2D airfoils. Most design techniques impose geometric constraints or designers' bias on the design space itself, thus restricting the novelty of the designs created, and only allowing for small local changes. We show that DbM methodology doesn't impose any such restrictions on the design space, and allows for extrapolation from the search space, thus allowing for truly radical and large search space with only a few parameters. We apply DbM to create a search space for 2D airfoils, and optimize this shape design space for maximizing the lift-over-drag ratio, CLD_{max} , and stall angle tolerance, $\Delta\alpha$. Using a genetic algorithm to optimize the DbM space, we show that we create a Pareto-front of radical airfoils that exhibit remarkable properties for both objectives.

3.2 Introduction

Optimizing the shape of an airfoil is an integral design stage for aerodynamic components like aircraft wings[135, 136, 137, 138] and wind-turbine blades[139, 140, 141, 142, 143, 144]. A typical airfoil optimization process contains three main components: shape parameterization, airfoil evaluation, and optimization, among which the parameterization method determines both the design space and the complexity of the optimization problem. In this sense, a desirable parameterization technique must cover a wide design space within a limited number of design parameters[145, 146, 147], which is especially important during the early design stage when minimum geometric constraints are placed and radical changes during the optimization process are welcomed.

Different shape parameterization methods offer different fidelity and ranges of control[148, 147, 145], and, according to the scope of the design parameters, one can place these methods

on a spectrum where, on one end, the change of one parameter affects only a local section of the airfoil shape thus offering a finer control of the shape, and, on the other end, each design parameter affects the airfoil's global contour[145].

On the local end of the spectrum is the **discrete method**[149], whose design parameters are exactly the discrete surface points that define the airfoil shapes. Because the displacement of each point can be adjusted, the design space is potentially limitless[150], and very fine local control as well as high fidelity can be achieved. However, to describe an airfoil shape accurately, a large number of surface points are needed, which increases the complexity of the optimization problem. Furthermore, to accommodate the large number of design variables, one usually uses gradient-based method to guide the optimization which is limited to small local changes and can easily get stuck at a local optimum.

Increasing the geometrical extent of each parameter's influence, one would find classical methods that determine the airfoil shape based on the regional features or the control points and perform curve-fittings of some kind. For example, the popular **parametric section (PARSEC) method**[151] uses eleven parameters that represent specific sectional features such as leading edge radius and upper and lower crest locations and approximate the airfoil surface using a 6th order polynomial. Another popular method would be the **Bézier parameterization**[152], which forms the upper and the lower surfaces of the airfoil through the Bézier curves defined by the pre-selected control points. Additionally, a combination of the two techniques, **Bézier-PARSEC parameterization**[153], also exists, which creates Bézier curves using the parameters of the PARSEC method and combines these curves to form the shape contours. One main issue with the above methods is their inability or inefficiency to include high-fidelity features: the PARSEC and the Bézier-PARSEC method have fixed number of parameters and offer very limited range of fidelity, while the Bézier parameterization requires higher-degree Bézier curves to describe complex shapes, which become inefficient to calculate as the order increases[150].

To include high-fidelity features, or, equivalently, represent more complex curves, **B-splines**[154, 155], including **nonuniform rational B-spline (NURBS)**[156], can be used, which form curves by connecting low-order Bézier segments defined by the control points. Naturally, with denser control points, these methods move to the local end of the spectrum and are able to represent high-fidelity features, but the computation complexity also increases. In an effort to reduce the number of the design parameters, the control points can be grouped together, and global transformations such as twisting and thickening can be used as the parameters instead. This is known as the **free-form deformation (FFD) method**[157, 158] and is closer to the global end of the spectrum. A similar method, namely **Radial Basis Function Domain Element (RBF) approach**[159, 160], also exists and uses radial basis function to exert deformation on the airfoil.

To the global end of the spectrum includes methods that use spectral construction of some basis functions or modes to form or deform the airfoil shapes. One typical way of determining the basis modes is through **proper orthogonal decomposition (POD)** of a set of airfoil data, and dimensionality can be reduced by using only the dominant modes[161, 162]. Other methods include the **Hicks-Henne's approach**[163], which uses a linear combination of sine

functions to deform the airfoil surface, and **class/shape function transformation (CST) method** proposed by Kulfan[164, 165], which represents an airfoil shape as the product of a class function and a shape function formed by a linear combination of Bernstein polynomials. Similar to all the other methods on the spectrum, in order to resemble high-fidelity features, more basis modes have to be included, which again falls into the so-called the curse of dimensionality.

There have been attempts that aim to reduce the number of parameters needed while capturing a large enough design space[166, 167, 168]. A most recent one was conducted by Chen et al.[169], who used a deep generative model, called Bézier-GAN, to parameterize airfoil shapes by learning from the major shape variations in an existing database. Furthermore, they preserve the minor features of the airfoil shapes via a noise space, which allows them to separate the major and the minor features hence leading to a faster design space exploration. However, this study, like many other dimension reduction methods, relies on the assumption that the optimum design is not far from an existing database, which is not always true. For our study, we mainly consider the shape parameterization technique for the early design stage and prefer not to make the same assumption. In particular, We are interested in a method that would contain high-order features while keeping a limited number of design parameters and allowing radical change from the initial airfoil shapes.

In this paper, we apply the Design-by-Morphing (DbM) parameterization technique, a novel design strategy that was introduced by Oh et al.[170] and has been used in recent years for geometry optimization of different problems[170, 171, 172], to the airfoil optimization problem. Specifically, our DbM method ‘morphs’ the baseline shapes together to create new shapes and can interpolate as well as extrapolate the design space, which allows for both the high-fidelity representation of shapes without the curse of dimensionality and radical improvements in the shapes without any geometric constraints[170, 172]. Our paper makes the following scientific contributions:

- A DbM parameterization technique designed for the two-dimensional airfoil shape optimization allowing both accurate reconstruction of the existing airfoil database and radical change of airfoil shapes while being free of geometric constraints and designers’ biases.
- An optimization strategy using the DbM parameterization technique and the genetic algorithm that is able to create the Pareto-front of multi-objective airfoil optimizations.

3.3 Design-by-Morphing

Design-by-Morphing (DbM) works by morphing homeomorphic, i.e. topologically equivalent, shapes to create a continuous and constraint-free design search space that can produce radical extrapolated shapes, something which is unique from existing design strategies. The details of DbM are presented in the subsequent subsections.

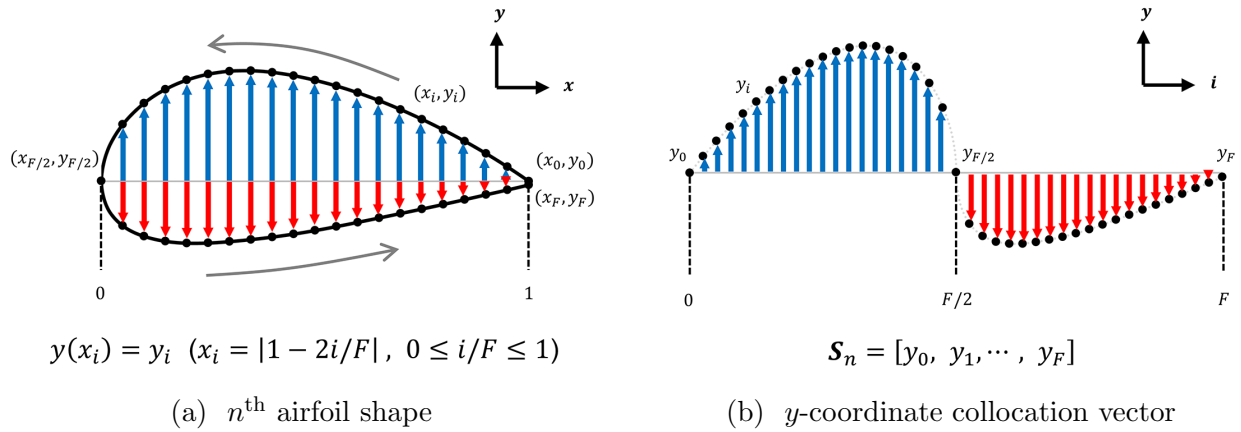


Figure 3.1: An example of DbM. The coordinates of the baseline shapes are weighted, summed, and normalized to form the coordinates of a morphed shape.




Figure 3.2: An Example of DbM. Column 1 shows the baseline shapes. Column 2 depicts the elements of the collocation vectors of the baseline shapes plotted as a function of the index i of the collocation vector. Column 3 shows the weighted elements of the collocation vector plotted as a function of the index i of the collocation vector. Column 4 shows the resultant collocation vector of the morphed shape and the morphed shape itself.

Baseline Shapes and Morphing

The DbM technique generally requires two or more ‘baseline shapes’, usually chosen from pre-existing designs in the literature, for the design search space creation. In order to be ‘morphed’ together, these baseline shapes must be homeomorphic, which can be achieved

by establishing a one-to-one correspondence between the shapes via some systematic shape collocation methods in either the functional[170] or geometric space[171, 172]. The new shapes can then be generated by applying weights to the baseline shape collocation vectors and summing them together.

The DbM method is valid for shapes of any dimensions, and because radically different baseline shapes can be morphed together, exotic shapes can be created. Furthermore, in addition to ‘interpolation’ between the shapes, applying negative weights during morphing allows ‘extrapolation’ from the search space spanned by the baseline shapes, which can create truly novel and unusual shapes. Lastly, DbM is completely free from any geometric parameter constraints and the only implicit constraints are the selections of the ‘baseline shapes’ themselves.

For 2D airfoils, the closed shapes can be collocated in the Euclidean coordinate system. We note here that all 2D shapes bounded by a single surface are homeomorphic to one another. Using the leading edge of each airfoil as origin, each shape can be collocated by taking fixed and uniformly spaced points on the x -axis, which creates a one-to-one correspondence between the shapes. This collocation strategy is demonstrated in Figure 3.1, and the baseline shapes used in this paper are chosen from various airfoils from literature, which are detailed later. Morphing is performed by multiplying a specific airfoil shape with a scalar weight, summing up the weighted vectors, and then normalizing them, which is given by Eq. 3.1:

$$\vec{P}(\vec{x}) = \frac{1}{\sum_{m=1}^N w_m} \sum_{n=1}^N w_n \vec{S}_n(\vec{x}). \quad (3.1)$$

Here $\vec{S}_n(\vec{x})$ is the y -coordinate collocation vector determining the n^{th} baseline shape, collocated at $\vec{x} = [x_0, \dots, x_F]$ where the i th x -coordinate $x_i = |1 - 2i/F|$ and F is the number of collocation points. Accordingly, the first half elements of \vec{S}_n represents the top surface of the airfoil and the second half elements of \vec{S}_n renders the bottom surface of the airfoil. N is equal to the total number of baseline shapes. $w_n \in [-1, 1]$ is the morphing weight applied to the y -coordinate vector of the n^{th} baseline shape. A visual demonstration of our strategy is presented in Figure 3.2.

Intersection Control

For smooth baseline shapes, applying positive weights, i.e. interpolation, will always create smooth shapes without applying any geometric constraint. However, because the DbM imposes no geometric parameter restraints, extrapolation, i.e. applying negative weights, may produce non-physical geometries such as self-intersections, which have ‘zero-area’ regions as shown in Figure 3.3(a). One may discard the morphed airfoil shapes with self-intersections during the optimization but that diminishes the size of our design space. Instead, we recover new shapes by removing the intersections. This is accomplished by first locating within the morphed coordinate vector where the intersection occurs and restructuring the coordinate vector by ‘flipping’ it between the intersection points as shown in Figure 3.3(c). The vector

is then ‘stiffened’ to remove the zero area between the intersections by removing the points in their neighborhoods and then linearly interpolating between the broken coordinate vectors. As seen in Figure 3.3(d), this removes the ‘zero-area’ space and gives some physical area to the shape at the point of intersection. The above process is repeated until all intersections are removed, e.g. both intersections in Figure 3.3 are successfully removed, and, finally, a moving-average smoothing filter is applied to smooth out the sharp edges.

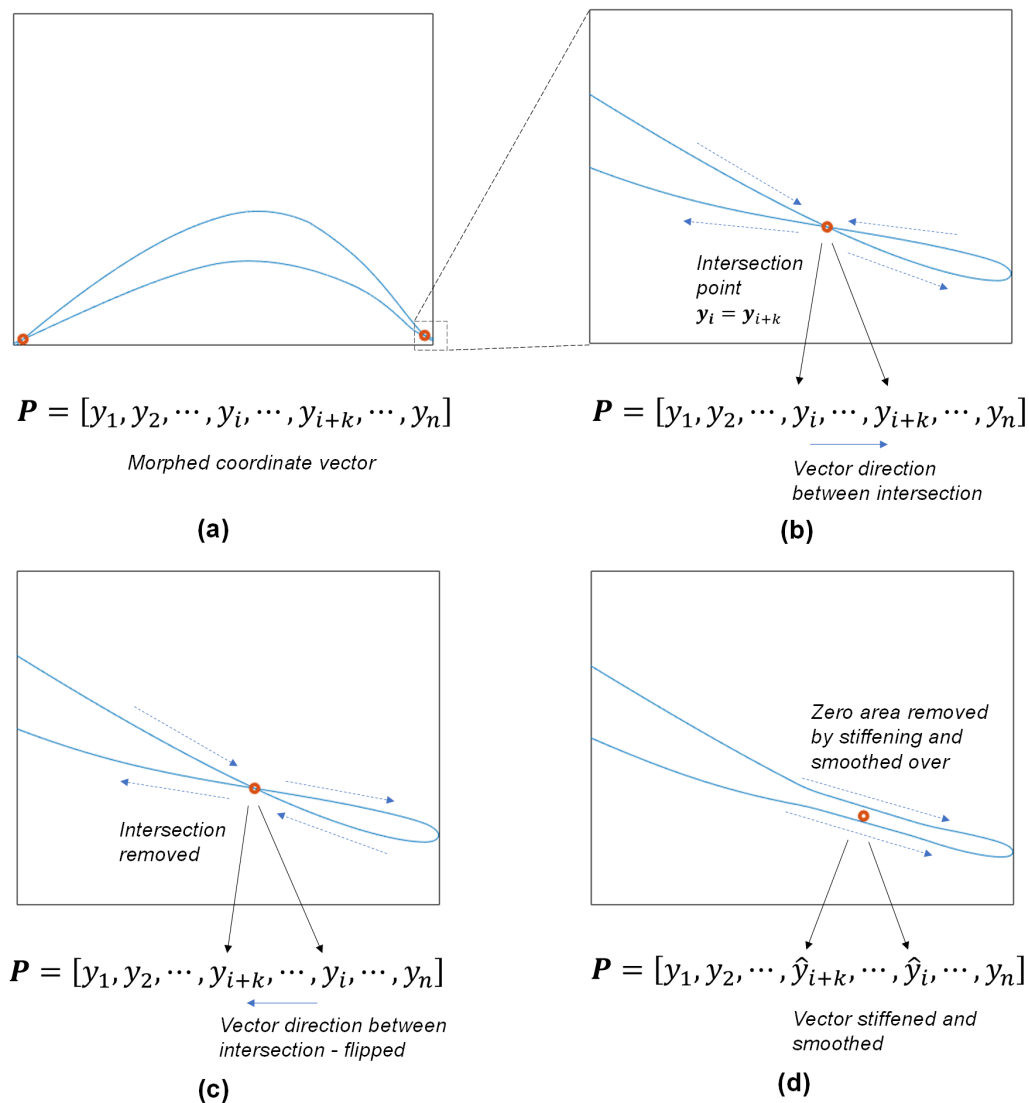


Figure 3.3: Conditioning for intersection removal. (a) Intersections are detected; (b) Blown up image of one intersection. Shape coordinates direction is depicted by arrows; (c) Intersection removed by flipping vector between intersection; (d) Zero area removed by linear interpolation to remove the intersecting area and then smoothed over, shown by hat coordinates

Baseline Shape Selection

The selection of baseline shapes is an important component of DbM strategy and ultimately determines the size and the novelty of our search space. Metaphorically, the selection of the baseline airfoil shapes serves as the gene pool for the morphed airfoils, and its diversity is important for creating a large design space. Our baseline shape selection contains good ones with either high lift-to-drag ratio or good stall performance, bad ones with poor aerodynamic performance, commonly used airfoil shapes in the literature or industry, and airfoils with irregular shapes to provide novelty to the design space. It is worth noting that, contrary to the conventional airfoil optimization processes[173], we deliberately include the bad performers so that our optimization can suppress their features by assigning negative weights to the corresponding baseline shapes. Our results in later sections will demonstrate this in greater detail.

In this paper, we select 25 baseline shapes (see Figure 3.4) from the UIUC airfoil coordinates database[174]. The airfoils are selected to ensure diversity and to introduce radical features in the design space. Their model names and characteristics are attached in Appendix 3.8. Each airfoil shape is represented by 4000 coordinates that span from the first surface trailing edge around the leading edge to the second surface trailing edge with equally distributed x -coordinates parallel to the airfoil chord line of a unit length.

Representation Capacity

To examine the robustness and the extent of our design space generated by the morphing of only 25 airfoil baselines, we reconstructed the pre-existing 2D airfoil shapes archived in the UIUC airfoil database[174] via DbM. A total of 184 randomly-chosen airfoils are tested, which accounts for approximately 11% of the UIUC database as of 2022. All 184 airfoils were reconstructed to the same accuracy, and a random selection of 100 out of the 184 airfoils are shown in Figure 3.5.

For each airfoil, the shape was reconstructed by running a global optimization of the weight vector that minimizes the total area of the original and morphed shapes where one shape does not overlap with the other, e.g. the geometric XOR of 2 closed shapes. Note that the geometric XOR serves as a good measure of the similarity between 2 shapes since it gradually goes to zero as the shapes become identical to one another. As a result, all the test airfoils were successfully re-created by the morphing of only 25 baseline shapes with the areal difference of less than 0.01. It means the average error in y -coordinate is $\pm 0.5\%$ since the airfoil chord is normalized as a unit length and the area is formed by 2 airfoil surfaces. This affirms that the current 25 baseline shapes are diverse enough to span the design space via DbM to explore airfoils in a universal manner.

Overall, by using only 25 dimensions, our DbM method is able to recreate the UIUC database airfoils with high fidelity, and it does not suffer from the ‘curse of the dimensionality’ compared to other techniques where the fidelity of the parametrization depends on the number of independent dimensions used. Moreover, DbM’s capability of creating exotic shapes via its




Figure 3.4: Twenty-five baseline shapes picked from the UIUC airfoil coordinates database[174]. See Appendix 3.8 for more details.

extrapolation feature increases the chance to find novel solutions that are deviated from the previously-established space like the UIUC database. This exploration is essential especially for the airfoil design where the correlation between the geometric feature and the aerodynamic performance of an airfoil can be very non-intuitive, thus necessitating exploratory design spaces.

3.4 Optimization Methodology

Our airfoil optimization methodology is built around the DbM technique introduced in Sec. 3.3. As shown by the flowchart in Figure 3.6, the optimization starts from the selection of the baseline shapes and then evaluates and optimizes the airfoils formed by morphing these baseline shapes using DbM. Our methodology does not rely on one specific airfoil evaluation tool or one specific optimizer, and discussions on their choices are provided in Sec. 3.4 and




Figure 3.5: Reconstruction of randomly chosen 100 pre-existing airfoil shapes via DbM using twenty-five baseline shapes in Figure 3.4.




Figure 3.6: General flowchart of airfoil optimization via DbM

Sec. 3.4 respectively.

Airfoil Evaluation

Our optimization methodology is not limited to one particular airfoil performance analyzer. One can use any CFD or experimental methods. For the optimization of airfoil shapes using CFD-based solvers, the evaluation of the objective functions (aerodynamic properties) typically falls into two categories: the full Reynolds-averaged Navier-Stokes (RANS) based approach and the interacted viscous/inviscid zonal approach. The RANS-based approach is computationally expensive and demands the optimizer to be highly efficient, and, to

accommodate the large number of design variables as often seen in the aerodynamic designs, gradient-based optimizers coupled with adjoint methods for computing the derivatives are deemed the most feasible[175, 176, 177]. On the other hand, the viscous/inviscid zonal approach, which combines separate solutions for the inviscid external flow and the viscous shear layer flow iteratively to form a continuous profile, is faster and less expensive. Among the many inviscid/viscous airfoil analysis codes, the XFOIL program[178] has been the most dominant and widely adopted one[179, 180, 181, 182, 183, 184, 185, 186]. It couples a vorticity panel method for exterior flow with an integral boundary-layer method for viscous boundary layers and uses an e^9 -type amplification formulation to determine the transition point[178]. Its applicability to airfoil designs has been demonstrated in the past literature, where its predictions of aerodynamic properties are shown in good agreement with the wind-tunnel experiment data[187, 188] and the RANS-based simulation results[189].

The specific choice of the evaluation tool used in this paper is not essential to manifest the power of the DbM parameterization technique, which is the main focus of our paper. For this work, we opt for XFOIL because of its acceptable accuracy under our flow condition as well as its low computation cost. Its wide usage also allows quick reproduction of our optimization results. It is used in a black-box manner so that any other commercial or in-house airfoil analysis tools can be incorporated into our optimization framework if necessary. Our detailed airfoil evaluation setup is attached in the appendix 3.8.

Optimization

When given a set of solutions, for single objective optimization problems, the most optimal solution within the set can be determined. However, for multi-objective optimization, multiple and potentially conflicting objectives must be considered simultaneously to determine the optimal answer in the solution set[190, 191]. If the designer has a quantitative ranking of the objectives, these objectives can be combined together to formulate a single objective problem, but when no such ranking exists, constructing a Pareto front is the most common methodology[192, 193, 194], which is applicable to real-world problems such as the design of architected materials[195, 196, 2], turbo-machinery[197, 198, 199], process-engineering[200, 201], shape design[202], and structural engineering[203, 204] when multiple objectives that cannot be quantitatively ranked are involved.

We pose the multi-objective optimization problem as

$$\vec{w}_{opt} = \operatorname{argmax}_{\vec{w} \in \mathcal{W}}(\vec{f}(\vec{w})), \quad (3.2)$$

where $\vec{f}(\vec{w}) = [f_1(\vec{w}), f_2(\vec{w}), \dots, f_K(\vec{w})]$. Here f_1, \dots, f_K are the K objectives to be maximized and \vec{w} is the design variable vector. Generally \vec{w} is a d -dimensional vector defined over a bounded set $\mathcal{W} \subset \mathbb{R}^d$ representing d continuous variables. $\{\vec{w}_{opt}\}$ is a set of Pareto-optimal solution vectors, i.e., a vector which is not Pareto-dominated by any other vector. For the reader's convenience, it is noted that a design variable vector \vec{w} is Pareto-dominated by another design variable vector \vec{w}' if $f_k(\vec{w}') \leq f_k(\vec{w})$ for all $k \in \{1, \dots, K\}$. To obtain the

Pareto-front, especially when objectives cannot be weighted or when a non-convex black-box function is considered, evolutionary or genetic algorithms are a natural choice[205, 206]. In fact, they have been commonly implemented in many previous aerodynamic optimization studies due to their gradient-free nature and wide region of the search domain[165, 207, 208, 209]. On the other hand, when the cost functions are expensive to compute (e.g. when using experiments as an evaluation tool), Bayesian optimization methods have proven to be efficient[1].

Our study considers a bi-objective ($K = 2$) two-dimensional airfoil shape optimization. In particular, we optimize the shape of a subsonic airfoil operating in an incompressible flow with $Re \equiv Uc/\nu$ of 1×10^6 , where U and ν are the free-stream flow speed and fluid kinematic viscosity, respectively, and c is the airfoil chord length. The parameter to be optimized is the morphing weight vector for the DbM technique:

$$\vec{w} \equiv (w_1, \dots, w_{25}) \in \mathcal{D}^{25}, \quad (3.3)$$

where $\mathcal{D} = [-1, 1] \subset \mathbb{R}$ and w_i ($i = 1, 2, \dots, 25$) is the weight applied to the i^{th} baseline shape. The design objectives to be maximized are the maximum lift-drag ratio at any angle of attack α , i.e. $f_1(\vec{w}) = CLD_{max}(\vec{w})$, and the difference between the stall angle α_s and the angle where the maximum lift-drag ratio occurs, i.e. $f_2(\vec{w}) = \Delta\alpha(\vec{w})$, where $\Delta\alpha$ is often called the stall angle tolerance. Precise definitions of these design objectives are explicated in Appendix 3.8, and both objectives are evaluated using the XFOIL simulations, which are efficient enough to be used with the genetic algorithm.

We use a MATLAB-based variant of the popular NSGA-II[44] algorithm, which is a controlled, elitist genetic algorithm. Our initial population consists of the single-objective optimums of each design target as well as random samples in the design space. The population size of 372 is used with a total of 3,000 maximum generations, and the solutions are actively ranked within each generation so as to maintain diversity and avoid over-crowding in the Pareto-optimal solution set. Our setup was tested on the commonly used set of ‘ZDT’ benchmark problems for multi-objective problems, suggested by Zeidtlar et al.[55]. The details of the test problems and the validation results are provided in Appendix 3.8.

3.5 Results

The Pareto front on the $\Delta\alpha - CLD_{max}$ objective plane, which resulted from the 3,000 generation genetic algorithm (GA) runs, is depicted in Figure 3.7. The convergence of the front is confirmed by the large generation number with the population size of 372, involving around 1.1 million XFOIL evaluations of CLD_{max} and $\Delta\alpha$. After non-dominant or duplicate individuals are removed in the final generation of the population, we are able to identify 208 Pareto-dominant airfoil shapes composed via DbM using 25 baseline airfoil shapes. For comparison, these 25 baseline shape cases are evaluated and plotted as red hollow circles in Figure 3.7 together. The reason why baseline #19 has zero CLD_{max} and $\Delta\alpha$ is that we inverted the shape intentionally and therefore XFOIL failed to evaluate its aerodynamic

performance. We assigned the objective functions zero values for such failing cases because they represented airfoil geometries that are not aerodynamically viable in the XFOIL space. The GA optimization successfully developed the Pareto front, where two ends are posed at $(CLD_{max}, \Delta\alpha) = (30.63, 40^\circ)$ and $(CLD_{max}, \Delta\alpha) = (273.39, 10^\circ)$. Even in the largest maximum lift-drag ratio case, the angle of attack gap between stall and design point is found to be 10° , giving the airfoil a tolerant range for off-design operations.

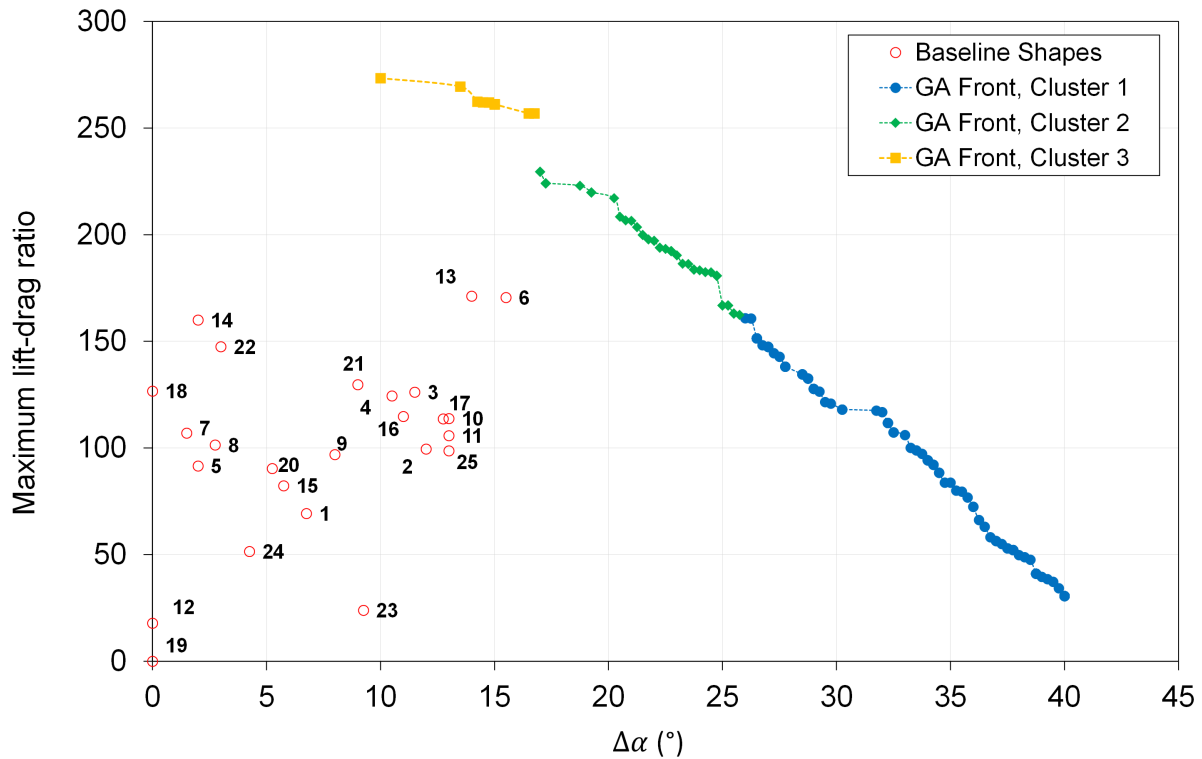


Figure 3.7: The Pareto front consisting of the optimal airfoil shapes as a result of the 3,000 generation runs of NSGA-II. Twenty-five red hollow circles with indices indicate twenty-five baseline airfoil shape evaluations. See Appendix 3.8 to understand how the clustering is performed.

The front is divided into 3 different clusters, each of which constitutes a segment of the front which does not overlapping one another. It is worth noting that the non-overlapping division of the front is a consequence of clustering through the Principal Component Analysis (PCA), rather than arbitrary. The detail of the clustering is provided in Appendix 3.8.

Figure 3.8 shows nine representative optimal airfoil shapes on the Pareto front in ascending order of CLD_{max} . In each cluster, three airfoil shapes having considerably different objective function values were selected to be presented. Also, note that Figure 3.8a illustrates the

extreme case of the smallest CLD_{max} and largest $\Delta\alpha$ while Figure 3.8i depicts the other extreme of the largest CLD_{max} and smallest $\Delta\alpha$. It can be seen that within the cluster the overall shape remains identical and only a gradual decrease in the airfoil thickness is observed as CLD_{max} increases. Since thin airfoils such as bird-like airfoils[210], which we take as part of the baseline shapes, e.g. #13 and #14, are known to have high CLD performance, the trend of airfoil thickness observed in the Pareto front appears to be reasonable.

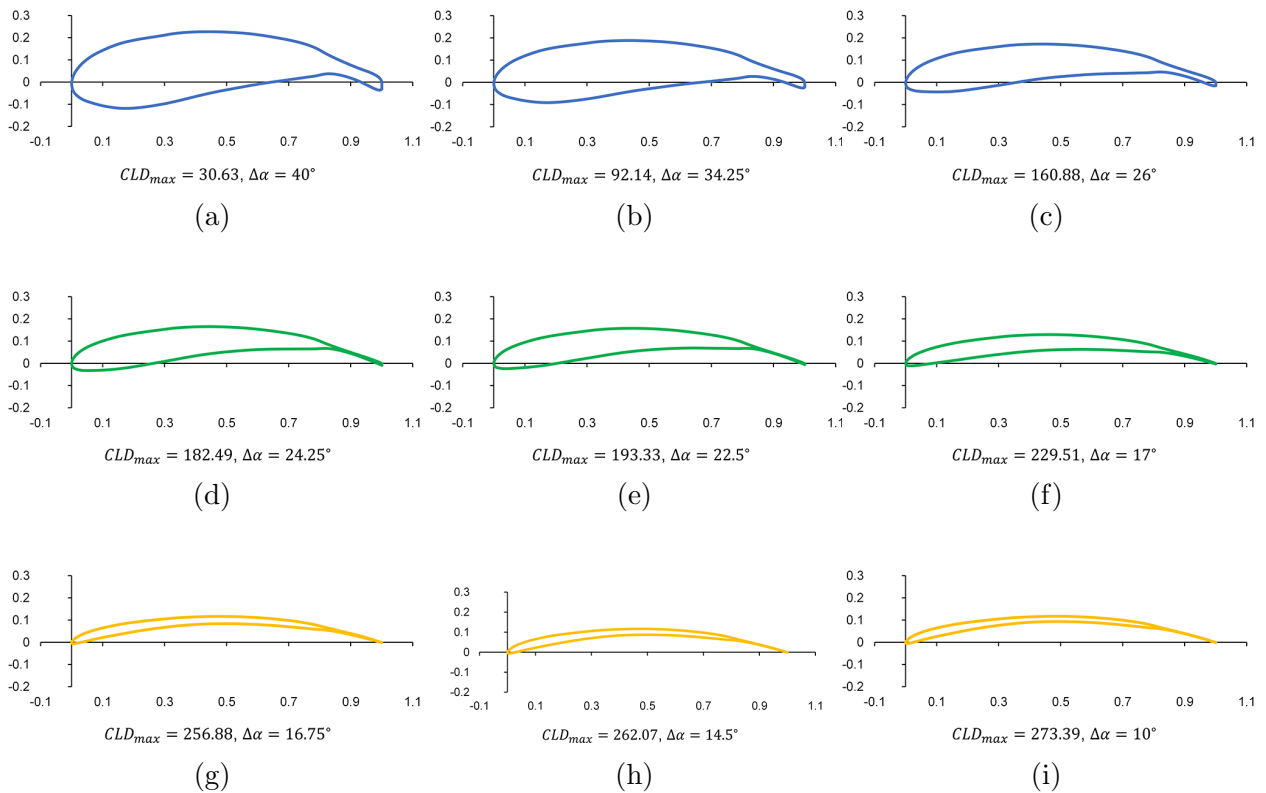


Figure 3.8: Nine representative Pareto-optimal airfoil shapes. a-c are in cluster 1, d-f are in cluster 2 and g-i are in cluster 3.

In cluster 1 where 112 optimal airfoil shapes exist, it is found that they mostly look similar to the total mean airfoil shape (see Figure 3.10a). This makes sense because they account for the majority of the airfoil shapes located on the front. Moreover, this cluster is located near the origin in the PCA-projected weight space (see Figure 3.13 in Appendix 3.8), indicating that there was no radical morphing of the airfoil shape taking place from the mean shape.

Next, cluster 2 contains 83 optimal airfoil shapes. Compared to the airfoil shapes in cluster 1, the most distinguishing feature is that the trailing edge region becomes narrow. This creates the sharp trailing edge, which is generally favorable to obtain a lift increase.

However, they are still not deviated far from the origin in the PCA-projected weight space, and their common shape mostly resembles the total mean airfoil shape.

Finally, 13 optimal airfoil shapes are discovered in cluster 3 from the optimization. This cluster includes the airfoil shapes experiencing more drastic morphing than those of in other cluster. It can be confirmed that they are the thinnest airfoils where the leading edge region’s thickness is also diminished.

The mean weight distributions with respect to 25 original baseline shapes are given in Figure 3.9. Overall, the weight distributions of 3 clusters comply with the weight distribution of the total mean. It turned out that baseline shape #13 (model name: AS6097) is commonly the most significant one for morphing. Since this baseline shape is the best in CLD_{max} and the second best in $\Delta\alpha$ among 25 baseline shapes (see Figure 3.7), it was likely to survive in the GA runs over the generations against the selection pressure that only sorts out dominant individuals in terms of both CLD_{max} and $\Delta\alpha$. However, excellence in the objectives of an individual baseline shape does not necessarily guarantee its survival, which is the case for the globally best baseline shape #6 (model name: AH 79-100C), as an individual’s superior ‘phenotype’ may be no longer revealed, or even suppressed after the morphing is done and all ‘genes’ are mixed with each other.




Figure 3.9: Mean weight distributions of the Pareto-optimal airfoil shapes with respect to twenty-five baseline airfoil shapes.

As we discussed from the examination of the morphed airfoil shapes, both cluster 1 and 2’s mean weight distributions show no considerable difference from the total mean weight distribution. Through small shape variation from the total mean airfoil shape as in

Figure 3.10a, it is possible to reach these optima relatively easily. In contrast, cluster 3 has a number of weights that are quite different from the mean (e.g., #6 and #11) and substantial morphing would be required if one starts with the total mean airfoil shape.

In the context of the present study, each axis obtained by the PCA can be represented by a unique form of morphed airfoil shapes because 25 PCA coefficient vectors defined in the weight space \mathcal{D}^{25} are orthogonal to each other. These 25 new morphed airfoils span the whole design space and therefore serve as alternative baseline shapes in lieu of the original ones. More importantly, the dominance of the first 2 PCA axes with respect to the data point variance suggests that the major geometric feature of 208 airfoil shapes we found via the optimization is virtually generated by morphing of these two new airfoils. Small variance of a PCA axis indicates that the data points are not considerably deviated from their mean on the axis. In other words, the baseline shape corresponding to this PCA axis has an marginal impact on morphing the airfoil shape for optimization. Once we pick two baseline shapes from the first two dominant PCA axes, whose associated collocation vectors are say \vec{P}_1 and \vec{P}_2 , and use them to morph the airfoil shape obtained from the total mean of the Pareto-optimal weight vector set, which corresponds to the mean collocation vector \vec{P}_{mean} , we get better understanding of how the morphing, especially along each PCA axis, has an influence on major geometric changes in the optimal airfoil shapes. These airfoil shapes are depicted in Figure 3.10, where the black and red surfaces are generated from the first and second half of the collocation points, respectively. For example, we note that the orientation of two surface of \vec{P}_1 is flipped in comparison to that of \vec{P}_{mean} , meaning that the stronger the weight of PCA axis 1 in the positive direction is, the narrower a morphed airfoil shape is.

3.6 Discussion

Most parameterization strategies depend upon careful selection of constraints and parameters, which determines their probability of success. The fidelity offered by such methods is very dependent on the number of the parameters chosen. Moreover, these designs are limited by the parametric constraints and the implicit designer’s bias, making extrapolation or radical global changes difficult. Data driven methods typically rely on the assumption that the optimum solutions are not far from the training data-set, which again prevents radical shape changes.

Design-by-Morphing, on the other hand, creates a design space that is uninhibited by any geometric constraints and also allows extrapolation from the design space. It doesn’t suffer from the curse of dimensionality when parameterizing airfoils and allows high-fidelity representation of airfoils without increasing the number of independent parameters in the problem. Using only 25 baseline shapes from the UIUC database, we were able to recreate the UIUC database with 0.5% error. We also showed that radical, global changes are possible using DbM. Applying that for the bi-objective shape optimization with objectives of maximizing CLD_{max} and $\Delta\alpha$, we were able to achieve significant results compared to our baseline shapes.

We posit that for design parametrization of airfoils and for other 2D/3D shapes, DbM

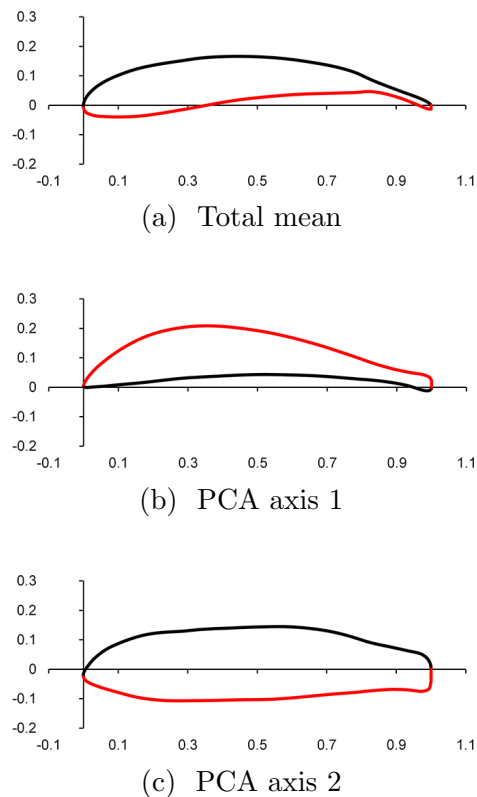


Figure 3.10: Morphed airfoil shapes generated by the optimal weight vectors, representing a the total mean of all optimal airfoils' weights, b the coefficients of the PCA axis having the most variance and c the coefficients of the PCA axis having the second-most variance. The black and red surfaces correspond to the first and second half of the collocation points, respectively.

should be the method of choice for creating an unconstrained, unbiased and non-data intensive design space that allows radical modifications, which can often be non-intuitive shapes.

3.7 Conclusion

DbM methodology creates a design space for radical 2D airfoils. We show that the space creates novel airfoils that are not constrained by geometric parameters or designer bias. Optimizing the design space created for dual objectives of CLD and $\Delta\alpha$, we show remarkable improvements in both objectives and provide a Pareto-front of optimal airfoil designs. Our final airfoils show remarkable improvements over our existing baseline shapes. For optimizing 2D or 3D airfoils, DbM should be used as the method of choice for design space creation.

Moreover, our methodology is flexible to be used for optimizing shapes for other fluid machinery as well. Currently we are applying DbM in tandem with Bayesian optimization for the optimization of 3D airfoils and vertical-axis wind turbines.

3.8 Additional Details for Airfoil Optimization and Design-by-Morphing

Aerodynamic Optimization Objectives

Airfoil optimization has become common in aerodynamic design problems involving maximization of one or more performance parameters of an airfoils. We mainly consider the following 2 performance parameters: the lift-drag ratio and stall angle. Given the flow speed U , fluid density ρ and airfoil chord length c , the lift and drag coefficients of an airfoil per unit span at an angle of attack α , C_l and C_d , are expressed as

$$C_l(\alpha) \equiv \frac{l(\alpha)}{\frac{1}{2}\rho U^2 c}, \quad C_d(\alpha) \equiv \frac{d(\alpha)}{\frac{1}{2}\rho U^2 c} \quad (3.4)$$

where l and d are lift and drag force per unit span, respectively, both of which change with respect to α . In this paper, these parameters are predicted via XFOIL[178], an program for analysis of subsonic isolated 2D airfoils, with varying α and then used for the optimization. Based on C_l and C_d , the lift-drag ratio CLD is calculated as:

$$CLD(\alpha) = \frac{C_l(\alpha)}{C_d(\alpha)}. \quad (3.5)$$

On the other hand, we define the stall angle α_s as an angle of attack where C_l reaches the first local maximum when we increase the angle starting from 0° , or

$$\alpha_s \equiv \min_{\alpha \geq 0} \alpha \quad \text{where } \exists \delta > 0 \quad \text{such that} \quad (3.6)$$

$$C_l(\alpha) \geq C_l(x) \quad \forall x \in [\alpha - \delta, \alpha + \delta]$$

Note that this definition is more conservative than the typical definition of the stalling in practice, where flow at the rear region begins to fully separate and C_l is globally maximized. α_s is occasionally smaller than the global maximum of C_l . Nonetheless, this approach helps avoid overestimation of the stall angle, which is expected to happen in XFOIL because of the nature of its flow solver having a limited accuracy in stall and post-stall conditions.

CLD and α_s have been typically considered to be significant to characterize the airfoil performance. For example, when it comes to lift-type wind turbines, the point where CLD is maximized may be commonly chosen as the design point. Since a wind turbine cannot always operate in the design condition, however, α_s needs to be additionally considered to evaluate how far the turbines run under an increasing-lift condition. For well-designed airfoils, α_s

generally occurs later than the design point, which yields tolerance in operation beyond the design point. Consequently, the stall angle tolerance, i.e. the range between these two angles of attack $\Delta\alpha$, which is expressed as

$$\Delta\alpha \equiv \max(0, \alpha_s - \operatorname{argmax}_{\alpha \in \mathbb{R}} \operatorname{CLD}(\alpha)), \quad (3.7)$$

can be a proper choice to evaluate the off-design performance [211]. Figure 3.11 depicts a scheme of how CLD and $\Delta\alpha$ are determined on performance curves of an airfoil.

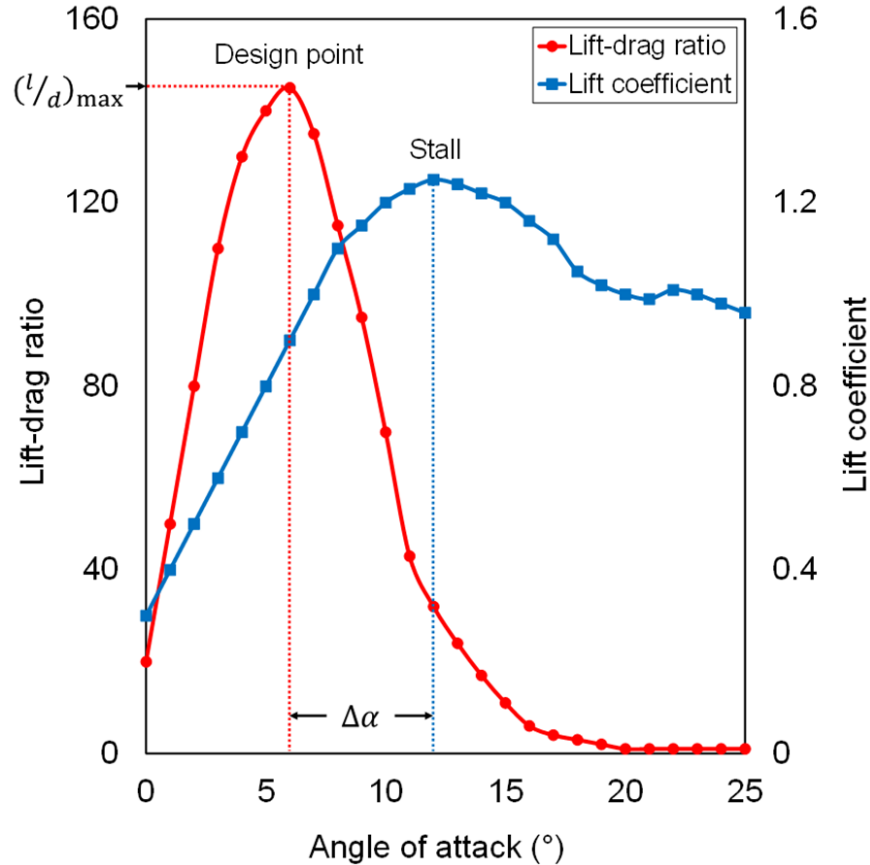


Figure 3.11: Airfoil performance curves

Baseline Airfoil Shapes and Validation

Our optimization methodology does not rely on one specific airfoil evaluation tool. To compare our results with the previous literature and to help future researchers quickly reproduce our results, we use XFOIL[178] in the present study. The two design objectives, CLD_{max} and $\Delta\alpha$, are obtained from the C_l and C_d data calculated by the XFOIL at different angle of attacks (see Figure 3.11).




Figure 3.12: Multi-objective optimization of benchmark test functions using GA

Optimization Test Functions and Validation

We use the multi-objective problems, suggested by Zeidtlr et al.[55], for testing our GA setup. The details of the test functions are given in Table 3.2. All the test functions were minimized with 25 variables in the design space.

MATLAB’s NSGA-II genetic algorithm, a fast sorting and elitist multi-objective genetic algorithm, was used in the current study. Single objective optimization for each objective

and random sampling were used for initialization. The population size of 372 was used with a total of 3,000 maximum generations. A ‘phenotype’ crowding distance metric was used. This setup was validated on the test functions described above. All the problems were benchmarked with 25 variables ($d = 25$) and two objective functions ($K = 2$) as with the present airfoil optimization problem. The results of our setup on four benchmark problems are shown in Figure 3.12. It was found that the algorithm could capture $ZDT1$, $ZDT2$, and $ZDT4$ accurately and predicts $ZDT6$, which is not only non-convex but also non-uniform, reasonably well.

Airfoil Shape Clustering

To analyze characteristics of the optimized airfoil shapes in detail, the airfoil shapes on the Pareto front are classified into 3 clusters using k -means clustering based on the Euclidean distance with $k = 3$. The clustering is performed in the design variable space, or weight space, of \mathcal{D}^{25} rather than in the objective plane because the purpose of clustering is to identify common geometric features over different airfoil shapes as a result of the optimization. The selection of the cluster size is based on the PCA of the optimal weight vector set.

Figure 3.13 shows the projection of the 25-dimensional weight vector set to the 2-dimensional subspace spanned by 2 PCA axes having the first- and second-most variance. The explained variance ratios of PCA axes 1 and 2 are 77.8% and 14.6%, respectively. On the other hand, the PCA axis of the third-most variance accounts for only 2.5% of the variance, affirming that the 2-dimensional projection in Figure 3.13 adequately scatters the clusters. From this observation, $k = 3$ is thought to be the most appropriate cluster size.

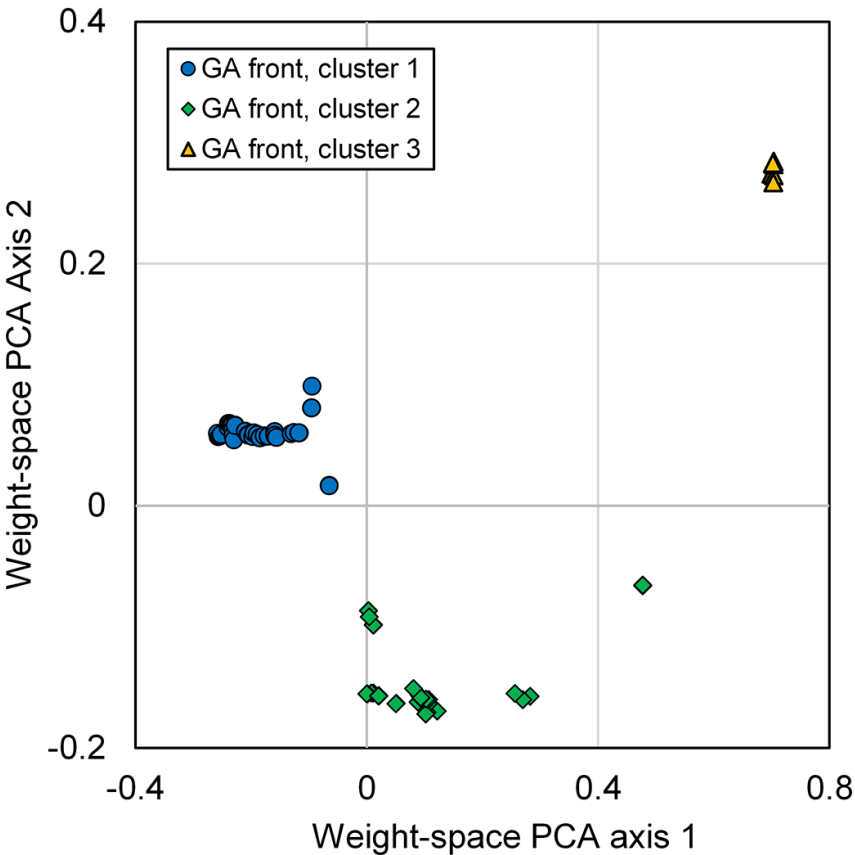


Figure 3.13: Projection of the 25-dimensional optimal weight vectors to the 2-dimensional subspace spanned by 2 PCA axes of the dominant variance. *k*-means clustering with the cluster size of 3 is used to identify the clusters.

Table 3.1: The model names, features, shape outlines, and XFOIL evaluation results of the 25 baseline shapes used by DbM in this paper. The coordinates of the baseline shapes are obtained from the UIUC airfoil coordinates database[174]. The airfoil evaluation results are obtained for an incompressible outer-flow of $Re = 1 \times 10^6$. The reference evaluation results are interpolated from the Airfoil Tools online database[212], where N/A indicates that there is no data available for this airfoil.


























Index	Model Name	Series (Features)	Airfoil Shape	Reference[212]		Present	
				CLD_{max}	$\Delta\alpha$	CLD_{max}	$\Delta\alpha$
1	NACA 0012	NACA (4-digit)		75.6	8.50	69.3	6.75
2	NACA 2412	NACA (4-digit)		101.4	12.00	99.5	12.00
3	NACA 4412	NACA (4-digit)		129.4	1.75	126.2	11.50
4	E 205	Eppler		128.3	8.50	124.4	10.50
5	AH 81-K-144 W-F Klappe	Althaus		89.7	2.00	91.6	2.00
6	AH 79-100 C	Althaus		183.0	14.75	170.6	15.50
7	AH 79-K-143/18	Althaus		110.9	1.50	107.0	1.50
8	AH 94-W-301	Althaus		103.0	4.00	101.4	2.75
9	NACA 23112	NACA (5-digit)		98.6	6.75	96.9	8.00
10	NACA 64(2)-415	NACA (6-digit)		120.6	12.50	113.8	13.00
11	NACA 747(A)-315	NACA (7-digit)		111.5	12.00	105.8	13.00
12	Griffith 30% Suc- tion	Griffith (Suction)		17.3	0.00	17.9	0.00
13	AS 6097	Selig (Bird-like)		N/A	N/A	171.2	14.00
14	E 379	Eppler (Bird-like)		N/A	N/A	160.0	2.00
15	Clark YS	Clark		85.7	5.25	82.3	5.75
16	Clark W	Clark		116.1	11.00	114.8	11.00
17	Clark Y	Clark		114.8	11.75	113.7	12.75
18	Chen	Chen		125.4	0.00	126.7	0.00
19	S2027 Flipped	Selig (Flipped)		N/A	N/A	0.00	0.00
20	GOE 417A	Gottingen (Thin plate)		86.7	5.25	90.4	5.25
21	GOE 611	Gottingen (Flat bottom)		125.6	9.00	129.7	9.00
22	Dragonfly Canard	Dragonfly		144.6	2.50	147.5	3.00
23	FX 79-W-470A	Wortmann (Fat)		N/A	N/A	23.9	9.25
24	Sikorsky DBLN-526	Sikorsky (Fat)		53.3	4.75	51.5	4.25
25	FX 82-512	Wortmann		99.1	14.75	98.7	13.00

Table 3.2: Benchmark Test Functions. All of the test functions are bi-objective with extended to n -dimensional constrained search space.

Problem	Bounds	Objective Functions	Optima	Note
$ZDT1$	$w_i \in [0, 1],$ $i = 1, \dots, n$	$f_1(\vec{w}) = w_1$ $f_2(\vec{w}) = g(\vec{w}) [1 - (f_1(\vec{w})/g(\vec{w}))^{1/2}]$ $g(\vec{w}) = 1 + 9 \left(\sum_{i=2}^n w_i \right) / (n - 1)$	$w_1 \in [0, 1]$ $w_i = 0,$ $i = 2, \dots, n$	convex
$ZDT2$	$w_i \in [0, 1],$ $i = 1, \dots, n$	$f_1(\vec{w}) = w_1$ $f_2(\vec{w}) = g(\vec{w}) [1 - (f_1(\vec{w})/g(\vec{w}))^2]$ $g(\vec{w}) = 1 + 9 \left(\sum_{i=2}^n w_i \right) / (n - 1)$	$w_1 \in [0, 1]$ $w_i = 0,$ $i = 2, \dots, n$	non-convex
$ZDT4$	$w_1 \in [0, 1]$ $w_i \in [-5, 5],$ $i = 2, \dots, n$	$f_1(\vec{w}) = w_1$ $f_2(\vec{w}) = g(\vec{w}) [1 - (f_1(\vec{w})/g(\vec{w}))^{1/2}]$ $g(\vec{w}) = 10n + \sum_{i=2}^n (w_i^2 - 10 \cos(4\pi w_i)) - 9$	$w_1 \in [0, 1]$ $w_i = 0,$ $i = 2, \dots, n$	non-convex
$ZDT6$	$w_i \in [0, 1],$ $i = 1, \dots, n$	$f_1(\vec{w}) = 1 - \exp(-4w_1) \sin^6(6\pi w_1)$ $f_2(\vec{w}) = g(\vec{w}) [1 - (f_1(\vec{w})/g(\vec{w}))^2]$ $g(\vec{w}) = 1 + 9 \left[\left(\sum_{i=2}^n w_i \right) / (n - 1) \right]^{1/4}$	$w_1 \in [0, 1]$ $w_i = 0,$ $i = 2, \dots, n$	non-convex, non-uniform

Chapter 4

Optimization of the Shape of a Hydrokinetic Turbine's Draft Tube and Hub Assembly Using Design-by-Morphing with Bayesian Optimization

4.1 Abstract

Finding the optimal design of a hydrodynamic or aerodynamic surface is often impossible due to the expense of evaluating the cost functions (say, with computational fluid dynamics) needed to determine the performances of the flows that the surface controls. In addition, inherent limitations of the design space itself due to imposed geometric constraints, conventional parameterization methods, and user bias can restrict *all* of the designs within a chosen design space regardless of whether traditional optimization methods or newer, data-driven design algorithms with machine learning are used to search the design space. We present a 2-pronged attack to address these difficulties: we propose (1) a methodology to create the design space using morphing that we call *Design-by-Morphing* (DbM); and (2) an optimization algorithm to search that space that uses a novel Bayesian Optimization (BO) strategy that we call *Mixed variable, Multi-Objective Bayesian Optimization* (MixMOBO). We apply this shape optimization strategy to maximize the power output of a hydrokinetic turbine. Applying these two strategies in tandem, we demonstrate that we can create a novel, geometrically-unconstrained, design space of a draft tube and hub shape and then optimize them simultaneously with a *minimum* number of cost function calls. Our framework is versatile and can be applied to the shape optimization of a variety of fluid problems.

4.2 Introduction

Motivation and Background

Renewable energy is fundamental to meeting our energy demands in a sustainable fashion. Although there exist several renewable sources for clean energy (wind, wave, solar, etc.), hydroelectric power is perhaps the most pertinent as a viable replacement for petroleum, natural gas and other fossil fuels. In fact, hydropower was one of the largest renewable energy sources in the United States in 2021 [213]. As such, it is important to consider ways in which the hydroelectric power plant (HEPP) can be more efficient, and thereby more cost-effective.

The relative size of HEPP's is usually classified by the amount of power they generate. Although there exists no unique definition of "small" hydropower, it is often accepted as a generating capacity at or below 10 megawatt electrical (MWe) [214]. The demand for "small" hydropower is steadily increasing [215] despite concerns about its potential adverse environmental impacts (but impact studies and how the impacts scale with the size of the hydropower plant remain controversial [216, 217]). Currently, there are numerous undeveloped sites across the globe that have large potentials for efficient and sustained power generation via small hydropower [218]. The sites with the most potential and that are easiest to exploit are those with low-impact stream-reaches, existing non-powered dams, and sites with existing conduits [214]. Motivated by this potential for inexpensive and sustainable energy, we propose here a new methodology for the design of some of the parts of a small hydrokinetic turbine. As a specific demonstration of our methodology, we provide a new, more efficient design of the draft tube and hub assembly of a small, low-head Kaplan turbine.

The turbine component of the HEPP works by converting the potential and kinetic energy of the water entering the turbine into mechanical work, and then producing electricity via a generator [219]. As one of the oldest and largest sources of renewable energy, there exists a wide variety of hydrokinetic turbines, but a concept common to all of them is the *dynamic pressure* or *hydraulic head* $P + \rho\mathbf{v}^2/2$ (where P is the static pressure, ρ the density and \mathbf{v} the velocity) of the water entering the turbine. Using Bernoulli's principle (see § 2) we can relate this head to the amount of power an ideal turbine can produce and the physical properties of the hydropower system such as the relative heights of the dam, penstock, and tail water discharge [220]. For high head ranges, impulse turbines (e.g., Pelton) exploit only the velocity of the fluid across the runner (see Fig. 4.1) to create the mechanical of the turbine blade; whereas in medium and low head ranges, reaction turbines (e.g., Francis and, with later developments for low head applications, Kaplan) exploit both the fluid's velocity *and* the fluid pressure or enthalpy across the runner. It is the conversion of enthalpy that allows the draft tube and hub assembly to enhance the performance of a reaction turbine [221, 222]. The draft tube and hub assembly make up only part of the overall hydropower system's performance, but they are most important for low-head turbines, which we consider here to be those under 20 m.

The draft tube is a diffuser, or several diffusers joined together, that sits beneath the runner of the turbine and directs the flow downstream of the turbine blades to the tailwater

pool. It therefore has a large role in determining the dimensions of the lower section of the power plant [223]. The draft tube increase the efficiency of the turbine by adjusting the dynamic pressure ($P + \rho\mathbf{v}^2/2$) such that the static pressure P just downstream of the turbine blade is decreased. The adjustment is done by decelerating the velocity of the fluid passing through the draft tube, and it is this property (hereafter referred to as the “pressure recovery”) that affects the power-generating capacity of Kaplan and other reaction turbines [224, 225].

The *hub* is a conically shaped part that extends just past the inlet of the draft tube, centered about the axis turbine’s of rotation and connecting its blades. Water flows into the draft tube through an annular-shaped region between the hub and the draft tube wall. The hub rotates with the same angular velocity as the turbine blade, while the outer boundary of the draft tube is non-rotating. (See Fig. 4.1.) It is important to optimize the hub because it modifies the inlet flow to the draft tube, and therefore has a large effect on the pressure recovery.

Poorly designed parts of the draft tube or hub can significantly decrease a turbine’s efficiency by exacerbating turbulence and increasing friction losses. For example, the draft tube’s elbow (Fig.4.1), which is necessary for redirecting the tailwater flow, can promote flow separation due to excessive centrifugal force at its inner radius. Similarly, a poorly designed hub can allow the swirl flow at the inlet due to the turbine blades create instabilities in the flow that lead to noise, vibration (prompting failure due to fatigue), and even the reversal of flow through the center of the draft tube (causing sudden changes in power output of the turbine) [226]. A well-designed hub can help prevent these problems, and, in addition, allow a larger opening angle of the diffuser (and therefore higher pressure recovery). As such, we not only optimize the draft tube, but also the hub, which is usually neglected in prior optimization studies [227, 228, 229].

Much of the optimization effort of Kaplan draft tubes in recent years was focused on improving the sharp-heeled elbow draft tube [229]. This shape was first installed in 1949 at the Hölleforsen Hydro Power Station in Vattenfall, northern Sweden, utilising a 25m head and with a power generating capacity of 150 MWe [230]. Gubin [229] and Dahlbäck [231] independently argued that this draft tube shape needed improvement, and subsequently there have been many proposed design changes. Marjavaara and Lundström [232] and Marjavaara [227] use a Response Surface Method (RSM) surrogate modeling strategy, as well as a commercial CFD code (ANSYS CFX4.4) to create new designs with different parametrizations (circular and elliptical, respectively) of the elbow section. More recent improvements include those of Daniels, Rahat, Tabor, Fieldsend, and Everson [233] and Daniels, Rahat, Everson, Tabor, and Fieldsend [234] who use multi-objective Bayesian optimization to maximize pressure recovery using a series of subdividing curves, optimizing over the inflow cone, outer-heel, and secondary straight diffuser. Other improved designs focus on the optimization of the draft tube for low-head applications while retaining the sharp-heel shape [235, 236].

Design-by-Morphing

To address the design challenges for improved hydrodynamic and aerodynamic surfaces, we introduced a new framework, Design-by-Morphing (DbM), for creating a design space that is versatile enough to include old and new designs and that is sufficiently free of human bias that radical and counter-intuitive designs are also included. DbM was first used by Oh et al. [170] and later used in other optimization problems [3, 171]. An N -dimensional design space for a shape is created by choosing N baseline shapes. The shape within our design space is specified by the choice of the N weights of these baselines from which the design is morphed. The bounds of the weights are sufficiently large that the morphed shapes can be not only interpolations among the shapes, but also extrapolations. Furthermore, any of the weights can be negative so that features of poorly performing baseline shapes can either be suppressed or entirely avoided. Negative weights and large positive weights allow for unintuitive shapes and for extrapolations. Many design techniques only allow for small departures from existing designs [237] or allow only local changes at one or a few specific locations, rather than global changes to the overall shape [238, 239]. This is especially problematic with most methods that use control points. CAD, and/or NURBS. Generally, methods with parametric control [148, 238, 240, 241, 242] and the adjoint method [237, 239, 243] limit the amount of change that can be made to a design so that radical new designs are not possible. DbM also allows a more extensive design space where both spatially local and global changes can be made to the design, and those changes can be subtle and/or radical [3]. This allows us to find an optimum which may be a completely new, unconventional shape [170].

Improved Bayesian Optimization

Once a design space is chosen, many engineering optimization problems require the repeated numerical (or laboratory) evaluation of an expensive multi-modal black-box function to determine the performance or cost of a particular design. In optimizations of a surface that interacts with a fluid, generally the fluid flow must be computed with an expensive Computational Fluid Dynamics (CFD) program to determine the quantitative performance of each candidate design. Those quantitative results are then fed into an optimizing algorithm to determine the best performing design. The expense of computing the performance function with CFD makes many optimization problems intractable. Furthermore, the highly nonlinear behaviour of fluid flows often leads to a design space where the performance function has many *local* maxima, and it is difficult for the search algorithm to find the *global* maximum.

Optimizing draft tubes is an example of a search requiring an expensive multi-modal, black-box, performance function because the efficiency and pressure recovery of each candidate draft tube must be computed with CFD. Bayesian Optimization (BO) is an efficient search method for this type of optimization because it requires far fewer evaluations of the performance function to find an optimum than most other optimization algorithms [18, 126]. BO techniques have been successful in the design of architected meta-materials [5, 6, 7, 8, 9, 2], hyperparameter tuning for machine learning algorithms [10, 11, 12], drug design [13, 14], and

controller sensor placement [15]. In this study, we use an improvement that we made to BO, that we call Mixed-variable Multi-Objective Bayesian Optimization (MixMOBO) [1]. Previously, we used MixMOBO to design an architected meta-material that has maximum strain-energy density [2].

4.3 Preliminaries

Any kinetic energy or potential energy that is not converted into mechanical energy of the turbine shaft (and then into electrical energy) is discarded (i.e., wasted) when it leaves the turbine blades. A well-designed draft tube minimizes the wasted energy by converting dynamic head into static head. The schematic of a hydrokinetic turbine assembly is shown in Figure 4.1. The overall drop in pressure across the turbine blade is $(\Delta P) \equiv P_0 - P_1$, where P_0 is the average fluid pressure just upstream of the turbine blade and is assumed to be fixed and independent of the designs of the hub and draft tube. P_1 is the average fluid pressure just downstream of the turbine blade, which is dependent on the designs of the hub and draft, and, in general, must be computed numerically and cannot be estimated using control volume analysis or a Bernoulli equation. However, control volume analysis does allow us to estimate the theoretically available power, \dot{W} , that can drive the turbine blade and produce electricity. It is

$$\dot{W} = A_0 v_0 (\Delta P) \quad (4.1)$$

where $A_0 = A_1$ are the cross-sectional areas of the flow upstream and downstream of the turbine blade (where the subscripts refer to the location in Fig. 4.1) and $v_0 = v_1$ are the characteristic velocities at these same locations.¹ We assume that along with P_0 , $A_0 = A_1$ and $v_0 = v_1$ are fixed, given parameters, and are not affected by the hub and draft tube designs. The goal of a well-designed hub assembly and draft tube is to maximize \dot{W} .

Traditionally, [223, 225, 244], the *mean pressure recovery* is defined as $P_2 - P_1$ and is the performance function used to the power output \dot{W} . We define the *dimensionless mean pressure recovery coefficient* as:

$$C_{prm} \equiv \frac{P_2 - P_1}{\frac{1}{2}\rho v_1^2} \quad (4.2)$$

where ρ is the density of the fluid, P_1 and P_2 are the average pressures, averaged over their respective cross-sections, and the subscripts refer to the location in Fig. 4.1. Using this definition and eq. (4.1), we see that \dot{W} and C_{prm} are related by

$$\dot{W} = A_0 v_0 [P_0 - P_2 + C_{prm}(\rho v_1^2/2)]. \quad (4.3)$$

¹Although v_0 can be thought of as an average streamwise velocity of the fluid upstream of the turbine blades, the amount of power that the blades can extract depends on the detailed flow interaction of the fluid velocity with the blade, which among other things depends on the swirl of the upstream velocity. Therefore, we leave the definition of this ‘characteristic’ velocity purposefully undefined.

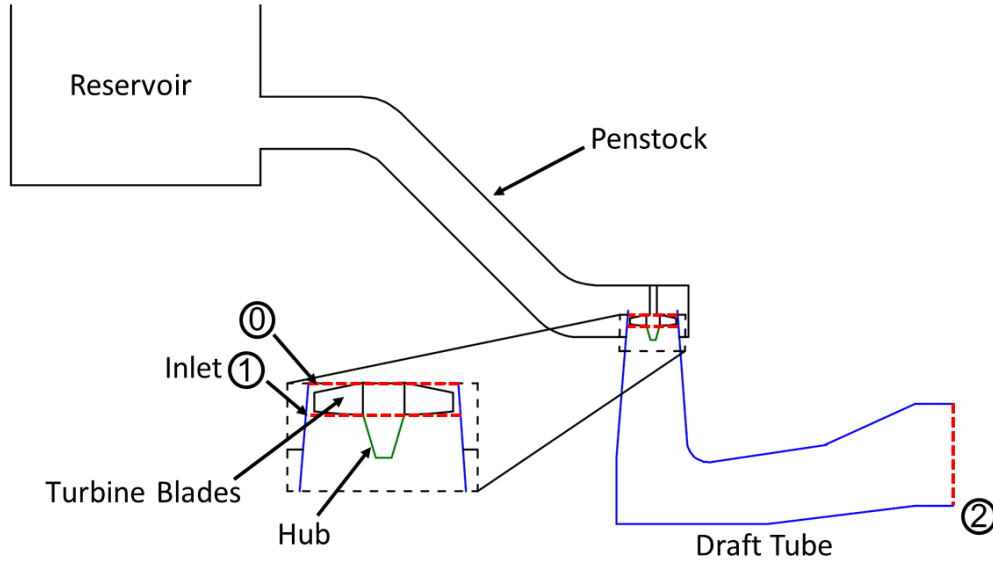


Figure 4.1: Simplified schematic of a hydroelectric power plant. The enumerated cross-sections in red are: (0) the entrance to the turbine; (1) the inlet of the hub/draft tube assembly; and (2) the outlet of the draft tube. The relative size of the draft tube (depicted in red) and hub (depicted in green) has been enlarged for clarity. The turbine blades are between points (0) and (1). The reservoir surface is open to atmosphere. The pressure at the draft tube exit (2) is P_2 and is a given, fixed reference or gauge pressure.

Because $(A_0 v_0 \rho v_1^2)$ is positive (and because A_0 , v_0 , ρ , P_0 , P_2 , and v_1 are assumed to be fixed, and independent of the designs of the hub and draft tube), maximizing C_{prm} and maximizing \dot{W} are equivalent. We therefore have chosen C_{prm} to be the performance function that is maximized in this study.

Note that in order to prevent back flow without a draft tube, P_1 would need to be greater than or equal to the given reference or gauge pressure P_2 . Because P_0 is assumed to be given (and the same value with or without a draft tube), the maximum value of \dot{W} without a draft tube would be constrained by

$$\dot{W} \leq A_0 v_0 (P_0 - P_2). \quad (4.4)$$

The draft tube allows $P_1 < P_2$, and therefore allows an extra amount of power, $A_0 v_0 C_{prm}(\rho v_1^2/2)$, to be generated. Note that because ρ , P_2 , and v_1 are assumed to be fixed, maximizing C_{prm} minimizes P_1 . Also note that because the fluid must be discharged from the draft tube at (2) with a finite velocity, and therefore the discharged fluid must necessarily contain some kinetic energy that cannot be recovered or converted into shaft mechanical energy, C_{prm} can never sufficiently maximized, and P_1 never sufficiently minimized (even in the inviscid case) to make the generating system 100% efficient [219, 220, 221].

4.4 Methodology

An overview of our procedure is demonstrated in the schematic in Fig. 4.2. Our method starts with five different draft tube and two different hub baseline shapes to create the DbM search space. This search space is then initially sampled randomly for 50 data points. This data was then used to determine the next epoch or batch of designs or test points to evaluate using the MixMOBO algorithm. Each batch is a set of 5 data points or designs. As each new batch of designs is evaluated, their C_{prm} 's are added to the data base that MixMOBO uses to compute the next batch of designs to evaluate. This procedure continues until the evaluation budget is reached.

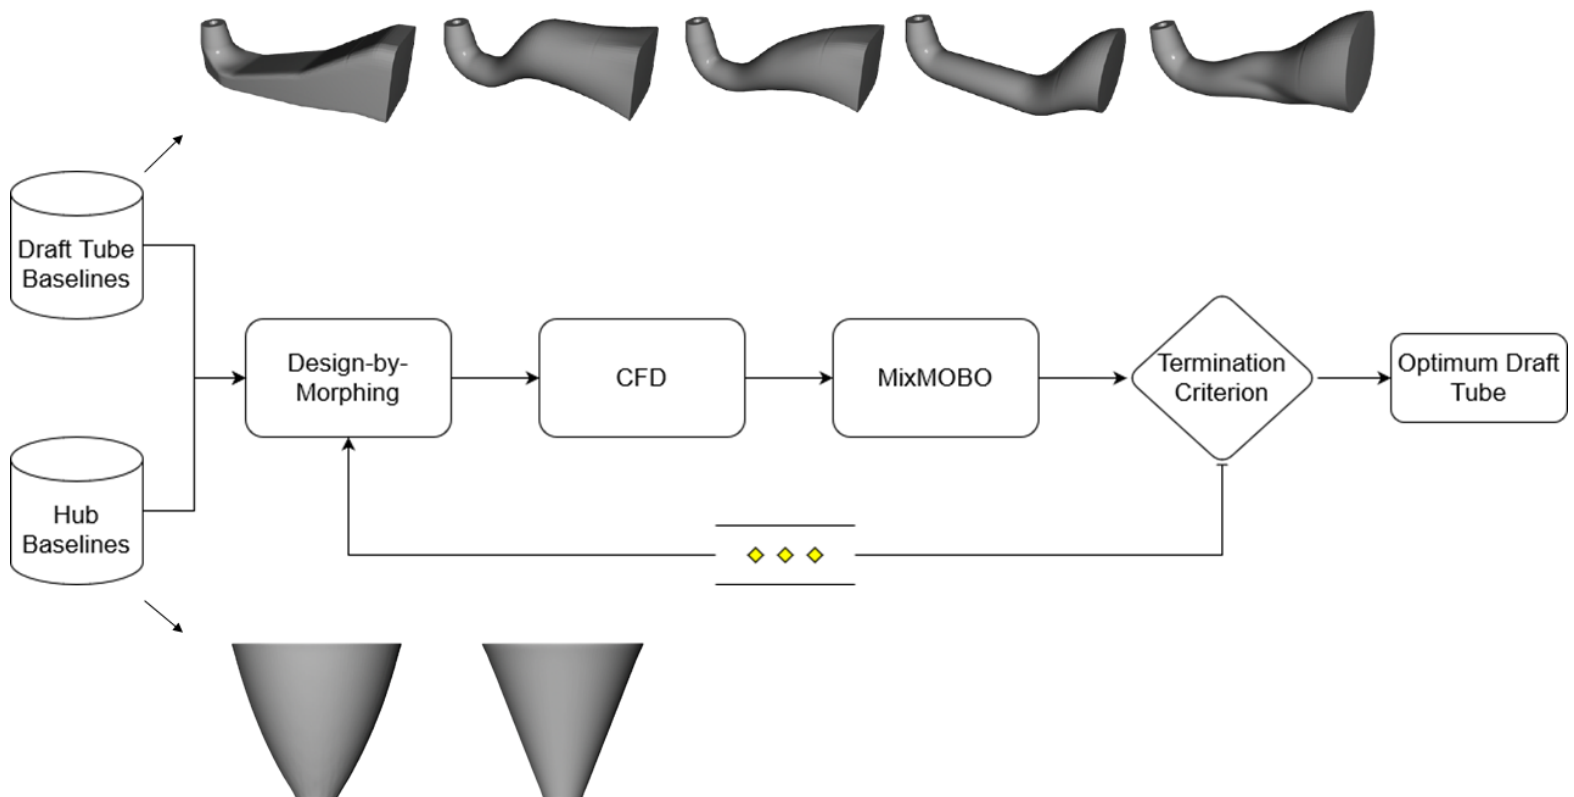


Figure 4.2: Optimization flowchart. The baselines for our draft tube and hub shapes are shown. The baselines are morphed together to create new shapes that are evaluated for their performance using CFD. This search space is sequentially optimized using parallel batches (represented by diamonds) of 5 shapes using MixMOBO, until the evaluation budget is reached. Note that the hub and the draft tube shapes are optimized simultaneously.

Baseline Shapes

The design space is created from five baseline draft tube shapes and two baseline hub shapes. The shapes are morphed with the weights chosen by MixMOBO to create a new draft tube and hub shape. These baselines are shown in Figure 4.2. The five baseline draft tube shapes are homeomorphic to each other as are the two baseline hub shapes. Some of the baselines that we chose were used previously in literature so that we were able to exploit any inherent advantageous features they might have, while other baselines were chosen to have non-intuitive features that would lead to radical designs, rather than incremental improvements. In the past, the fundamental design features of the Kaplan draft tube were formed through experimental observation and quasi-empirical formulae derived from geometries already installed and in use in HEPPs [233]. The works of Gubin [229], Cervantes [223], Mulu [245] and Nilsson et. al. [226] [246]) provide insights into draft tube geometries.

Our first baseline shape for the draft tube is the sharp heel draft tube, which has been the focus of extensive optimization attempts. It is also the subject of extensive experimental and numerical studies, the majority of which were completed through the European Research Community On Flow, Turbulence And Combustion (ERCOFTAC) Turbine-99 Workshop series [247, 248, 244]. For this reason, this shape also provides excellent validation of our CFD setup.

The second and third baseline draft tube shapes are based on designs cited by Gubin [229] and are particularly well-suited to low-head applications (i.e., for Kaplan turbines). The fourth and fifth baseline tube shapes were both devised in order to create features which would expand the design space. These shapes have radically different features such as a rounded outlet and a circular diffuser. All the baseline draft tubes have the same shape at their inlets, and in all cases the planes containing inlet and outlet are perpendicular to each other.

Both baseline hubs have the same radius at the inlet, the same radius at the end of the hub, and the same length. The first baseline hub shape is based on one currently used in a Kaplan turbine [245, 249]. The second baseline hub is a cone shape, and which was historically used in a wide variety of low to medium head reaction turbines [229]. Morphing these two hubs allows for a variation in the inlet geometry, which in turn modifies the flow profile near the inlet and therefore the pressure recovery. The baseline hub shapes are shown in Figure 4.2.

Design-By-Morphing

Design-by-Morphing (DbM) works by creating a one-to-one correspondence between the set of baseline shapes. The shapes are then “morphed” i.e. linearly combined together using “weights”. These weights are the independent parameters that form the search space for optimizing the shape. DbM requires that the set of shapes be homeomorphic or topologically equivalent so that a one-to-one correspondence can be created between the surface elements. This one-to-one correspondence is created by defining a collocation strategy for the baseline

shapes that ensures all constraints for our design are fulfilled (for example inlet and outlet orientation) even if individual shapes are radically different from each other.




Figure 4.3: The origin curve is an arbitrarily chosen curve, shown in the figure with a dot-dash, that starts at the inlet and ends at the outlet. Its only constraint is that it must lie within all of the baseline draft tubes. Because all of the baseline draft tubes used in this study are reflection symmetric about the same x - z plane, we chose to embed the origin curve within this symmetry plane. A unit vector \hat{s} lies along the origin curve pointing from the inlet to the outlet. The arc-length, or coordinate, along the origin curve is denoted as s , with a value of $s = 0$ at the entrance and $s = S$ at the exit. We choose N_s equally-spaced collocation points $\{s_i\}$, $i = 0, 1, 2, \dots, N_s - 1$ along the origin curve, with s_0 in the inlet plane, and s_{N_s-1} in the outlet plane. Planes are defined at each value of s_i that are locally perpendicular to the origin curve, with the planes at s_0 and s_{N_s-1} , at the inlet and outlet, respectively. The figure shows three of these planes as dashed lines at s_0 , s_j , and s_{N_s-1} , where $1 \leq j \leq N_s - 2$. The solid lines in each of these planes show the local polar coordinates (r, φ) within each plane. In each of the N_s planes, we chose the origin of φ to lie in the x - z symmetry plane, and the angle $\varphi = 0$ is indicated in each of the three illustrated planes as broken lines.

For the draft tube baseline shapes (shown in Figs. 4.3 and 4.4), the inlet is constrained to a plane parallel to the horizontal (x - y) plane and the outlet is constrained to a plane parallel to the vertical (y - z) plane for all the baseline shapes. Note that this ensures that all the morphed shape inlets and outlets are similarly constrained i.e. inlets are in the same horizontal plane and the outlets are in the same vertical plane. Fig 4.3 explains how we construct a single coordinate system (r, φ, s) for all of the baseline and morphed draft tubes that is mapped from the cylindrical coordinates (r, φ, z) . The z axis is mapped to the origin curve in the figure with with arc length s and local unit vector \hat{s} .

The boundary of the draft tube is defined by its intersections with each of the N_s perpendicular planes shown in Fig. 4.3. Within each perpendicular plane, the intersection of the plane with the draft tube boundary is defined by the closed curve $r_j(\varphi)$, the radial

distance of the draft tube boundary from the origin curve in the j^{th} perpendicular plane. We note here that, because the same origin line is used for all of the baseline shapes, the start and end points of the origin line are always constrained to the same location in the inlet and outlet planes. The draft tube shape, however, changes around the origin line, meaning that the average path length for a fluid particle to travel from inlet to outlet, can vary for each draft tube.

In each of the N_s perpendicular planes, we discretize φ into N_φ equally-spaced collocation angles φ_k with $k = 0, 1, 2, \dots, N_\varphi - 1$, with $\varphi_0 = 0$ and $\varphi_{N_\varphi} = 2\pi(N_\varphi - 1)/N_\varphi$.

The radius of the p^{th} draft tube as a function of the arc length s and polar angle φ is completely defined by the radial location matrix $R_{k,j}^p \equiv r^p(\varphi_k, s_j)$, with $k = 0, 1, 2, \dots, N_\varphi - 1$, and $j = 0, 1, 2, \dots, N_s - 1$. Note that the s_j perpendicular planes and the φ_k collocation angles are the same for all of the the baseline and the morphed draft tubes.

For hub shapes, the same collocation strategy is used as we used for the draft tube shapes, with the origin curve of the hubs passing through the center of both of the baseline hubs because the baseline hub are axi-symmetric.

After morphing, each morphed radial matrix can be projected back into a 3-D shape as in Figure 4.4.

Because all of the five baseline draft tube shapes are homeomorphic, they can easily be combined into a new morphed shape, given by $r^{morph}(\varphi_k, s_j) \equiv R_{k,j}^{morph}$, once the weights w_p , $p = 1, 2, \dots, 5$ of each baseline are chosen (see Fig. 4.4):

$$R_{k,j}^{morph} \equiv \frac{1}{\sum_{p=1}^5 w_p} \left| \sum_{p=1}^5 [w_p R_{k,j}^p] \right| \quad \forall \sum_{p=1}^5 w_p \neq 0. \quad (4.5)$$

Note that negatives weights and weights greater than unity allow for extrapolations; negative weights also allow us to ‘‘avoid’’ some baselines. This means that the only existing constraints are those imposed by choosing the baselines shapes themselves. That said, the lack of CAD parameterization means that DbM may produce non-physical, self-intersecting shapes when negative weights are applied. If the morphed shape is not physical because it has intersecting radial boundaries, we set the pressure recovery function of that morphed draft tube function to be zero so that the optimization method avoids that region of design space. For draft tube optimization, we limit the allowable range of the weights of each of the draft tubes such that: $w_p \in [-0.5, 1.0]$.

The morphed shape of the hub is given by an equation similar to eq. (4.5), but the sum is over only 2 baseline radii. We denote the weights for hubs with β to differentiate it from the draft tube weights. Furthermore, the sum weight for the second hub baseline is constrained by $\sum_{p=1}^2 \beta_l = 0.5.$, and for the morphed hub shape, $\beta_1 \in [-0.5, 1]$. Due to the fact that sums of the weights of the draft tube baseline shapes and of the hub baseline shapes are fixed, there are only 6 degrees of freedom in choosing the values of the weights of the baseline shapes, so our design space has 6 dimensions.

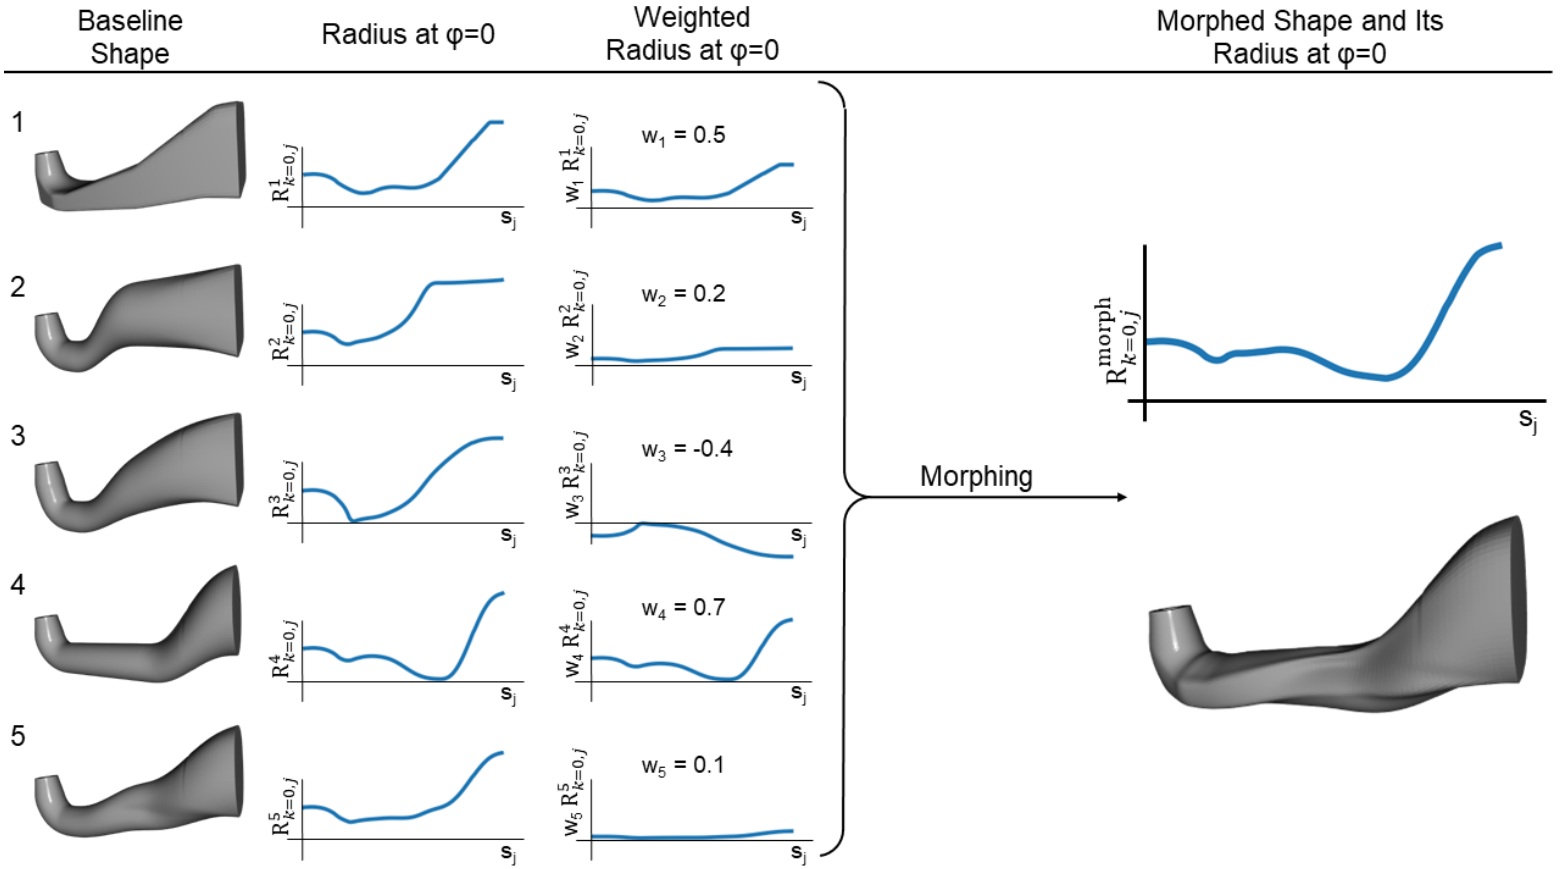


Figure 4.4: An example of design-by-morphing (DbM). Column 1 shows the baseline shapes. Column 2 shows the radius $R_{k=0,j}^p \equiv r^p(\varphi_0 \equiv 0, s_j)$ of each baseline shape as a function of j . Note that for ease of visualization, we are only plotting the radius at the “top” of the baseline shape, rather than for all φ . Column 3 shows the weighted radius of each baseline at its top (i.e., the product of the baseline’s weight with its radius at the top). Column 4 shows the morphed shape produced from the weights as well as the radius of the top of the morphed draft tube as a function of s .

CFD Setup and Validation

The coefficient of pressure recovery C_{prm} of each morphed draft tube shape that is physically allowable is determined with CFD. We validated the CFD code by comparing our simulations of the sharp heel draft tube with previously published results. In order to compute the flow in a morphed hub and draft tube shape, a mesh of that shape is created from a surface point cloud of the shape using Gmsh [250]. The average statistics of a typical mesh is given in Table 4.1, where D is the inlet diameter. The software used is OpenFOAM, and the solver used is pimpleFoam.

Table 4.1: Draft Tube Mesh Statistics

Max Cell Size / D	0.638
Number of Elements	7.9e+05
Number of Nodes	1.3e+05
Meshing Algorithm	Delaunay (3D)

For turbulence modeling, we use the $k-\omega$ *SST* turbulence model. This choice was made based on its success in previous draft tube studies [225, 251]. Based on the kinematic viscosity of water, the distance between the outer edge of the hub and the inlet wall, and the average streamwise velocity at the inlet, the Reynolds number is $5.56e+05$, and based on the azimuthal velocity of the hub, it is $9.48e+05$. No-slip conditions are applied at the rotating, inner hub wall (62.3 rad s^{-1}) and at the non-rotating draft tube boundary.² The inlet velocity is axisymmetric with non-zero values of the azimuthal, radial, and streamwise components. Our CFD simulations use the experimentally measured inlet velocities found by Engström, Gustavsson, and Karlsson [248]. The outlet pressure is the pressure at location (2) in Fig. 4.1, and is fixed gauge pressure, but the pressure at the inlet, needed to determine the C_{prm} , is computed by the CFD code as a function of r , φ , and time. We run the CFD for 128,571 time steps for each morphed hub/draft tube shape. Our choice of the number of time steps is based on how long it takes the solution to reach a statistical steady state – see Fig. 4.5. Note that C_{prm} , given by eq. (4.2), is not determined at a single time step, but rather it is averaged over the final 28,500 time steps of the computation (and note that as shown in Fig. 4.5, the solution has converged to a statistically-steady equilibrium during those final time steps). The time step was chosen to be $0.00134 D/u_{avg}$, where D is the outer diameter of the hub and u_{avg} is the average streamwise inlet velocity. This time step was chosen based on a time resolution study (see below).

The CFD is validated in two ways. The first compares our CFD-computed pressures with the experimentally-measured values along the top and bottom center lines of the baseline 1, Sharp-Heel, draft tube. The second uses baseline 1 to test the convergence of the CFD code by refining the numerical time step and grid size. We note here that the experimental studies did not provide the C_{prm} values, thus no C_{prm} comparison could be made.

The local pressure recovery along the top and bottom lines of the first baseline shape (i.e., the sharp heel draft tube) is shown (and defined) in Figure 4.6. The figure compares our numerically-computed local pressure recovery values to the experimentally-measured values [248]. Figures similar to Fig. 4.6 appear throughout the ERCOFTAC Turbine-99 Workshop series, and the deviation between our CFD and experimental values are consistent with the previous numerical studies.

²The optimization studies in the past for draft tubes [235, 236], did not consider a rotating hub inlet condition, which makes the CFD boundary conditions less close to the actual operating or experimental conditions




Figure 4.5: Numerical evaluation for all five baseline draft tube shapes using the first baseline hub shape. In our study, each simulation was run for 128,571 time steps. We chose this number of time steps to insure that the computed flow has come to a statistically steady state. Note that it requires approximately 50,000 time steps for the C_{prm} to come to a statistically steady state. (We assume that the flow itself comes to a statistically steady state some time between 50,000 and 128,571 time steps). The sharp heel draft tube (baseline shape **1**), has the highest C_{prm} among the tube baseline shapes.

The second way we validated our use of our OpenFoam code and gridding scheme was by time step and mesh refinement. We chose our time step size by decreasing its value in the CFD code until the late time-averaged C_{prm} values of the sharp heel draft tube with the first hub baseline shape changed by less than 0.01%. In particular, once we set the value of the time step as $0.00134 D/u_{avg}$, we repeated the calculation of the local pressure recovery $C_p(j)$ for both the top and bottom as a function of j as defined in Fig. 4.6 – once with a time step 20% greater than that used in the figure and once with it 20% lower. Plots of these new curves computed with these time steps are indistinguishable from the curves (error less than 0.5%) shown in Fig. 4.6 (at the resolution at which we created the figure), validating that our time step is sufficiently small. We determined the size of the spatial resolution of the grid for the CFD in a similar manner: we decreased the grid size (or equivalently, increased the square of the number of grid points) until the late time-averaged values C_{prm} values of the sharp heel draft tube changed by less than 0.4%. We also recomputed the two curves of the $C_p(j)$ in Fig. 4.6 with a spatial resolution of the spatial gridding mesh with a size that was 20% greater than that used in the figure and once with it 20% lower. Again, we found that the plots of these new curves computed with these grids are indistinguishable (error less than 0.5%) from the curves shown in Fig. 4.6, validating that our grid is sufficiently fine.

It is important to note that Fig. 4.5 shows that it takes more than 50,000 time steps for our numerically-computed C_{prm} to reach a statistical equilibrium. The 50,000 time steps

corresponds to a time that is approximately equal to 10 “advective times”, where the latter is the streamwise length of the draft tube divided by the average streamwise velocity at the inlet u_{avg} . Previous numerical studies of draft tubes were often validated by comparing the numerically-computed local pressure recovery $C_p(j)$ at the bottom or top of the draft tube with the experimentally measured values. However, in many of those studies the codes were run for only *one* advective time (5000 of our time steps) before the pressure recovery factor C_{prm} was evaluated [248]. Figure 4.5 clearly shows that although that time be sufficient to numerically capture the final statistically steady state of $C_p(j)$, it is insufficient for the full flow and C_{prm} to have come to a statistical steady state.

Bayesian Optimization

The search in our 6-dimensional design space (consisting of the 5 weights of the baseline draft tubes and the 1 independent weight of the two baseline hubs) for the draft tube/hub design with the maximum C_{prm} is an example of an optimization problem that can generically be posed as:

$$\vec{w}_{opt} = \underset{\vec{w} \in \mathcal{W}}{\operatorname{argmax}} [f(\vec{w})], \quad (4.6)$$

where $f(\vec{w})$ is the objective to be maximized (in this case C_{prm}), and \vec{w} is a D -dimensional (in this case 6-dimensional) variable vector, defined over a bounded set $\mathcal{W} \subset \mathbb{R}^D$ (in this case, the weights of the baselines). For many practical engineering problems, $f(\vec{w})$ is expensive to




Figure 4.6: Pressure Recovery at the “Top” or “Bottom” of the Sharp-Heel draft tube and Hub 1 (in the x - z symmetry plane). According to our collocation defined in Section 4.4, the “Top” corresponds to $R_{k=0,j}^1$ and “Bottom” corresponds to $R_{k=N_\varphi/2,j}^1$ as a function of j , where $j = 0$ is at the inlet. We define $C_p(j) \equiv \frac{P_{wall}(j) - P_{wall}(0)}{\frac{1}{2}\rho\left(\frac{Q}{A_1}\right)^2}$, where $P_{wall}(j)$ is the pressure at the draft tube wall. Solid circles are the experimental data; continuous curves are our numerically computations. The time step is $0.00134 D/u_{avg}$.

evaluate (as is the case here, which requires repeated CFD evaluations of a complex draft tube designs). In such cases, a Bayesian optimization algorithm is often the best choice [18]. Bayesian optimization is a sequential optimization technique, specifically designed to find the optimal solution of a noisy black-box function f with the fewest possible evaluations or function calls to f . At every iteration of the algorithm, a surrogate model g , usually a Gaussian process (GP), is fit over the data set $\mathcal{D} = \{[w_1, f(w_1)], \dots, [w_i, f(w_i)]\}$. Here i is the total number of points (in this case, morphed draft tube/hub designs) evaluated until the i^{th} iteration. Once the surrogate surface has been determined, one of two required things must be done: the point in the design space that is the most likely optimum (e.g., the design with the largest C_{prm}) must be determined, and the point whose evaluation will most likely best improve the fit of the surrogate surface must be determined. These two determinations are not the same, and the former is called “exploitation” and the latter “exploration”. Bayesian optimization is a balance of exploitation and exploration that continues until there is either evidence that a global optimum has been found or a maximum proscribed number of iterations of the algorithm is reached. An acquisition function is used to determine the next point (or points) \vec{w}_{i+1} in the design space to be evaluated with f (in this case, with the costly CFD), by balancing the competing needs of exploitation and exploration. Once the point (or points) has been determined, f is evaluated for that point (or points) and is then appended to the data set, $\mathcal{D} = \mathcal{D} \cup (w_{i+1}, f(w_{i+1}))$. In our case, the process is repeated until the evaluation budget is reached. Our choice for the evaluation budget was based on several optimization experiments on test functions and design spaces, described below, that we believe to be representative of the draft tube/hub optimization.

In our previous work, we developed a Mixed variable, Multi-Objective Bayesian Optimization (MixMOBO) algorithm [1], a framework for optimizing mixed-variable multi-objective problems with noisy black-box function using parallel batch updates of the data set, that is, letting the acquisition function pick several points to evaluate at each iteration to allow their CFD evaluations to be carried out in parallel. MixMOBO was proven to be more efficient in terms of number of black-box function evaluations compared to other algorithms in small data settings. In our previous studies, we used MixMOBO to design a microlattice structure with the objective of maximizing its strain-energy density [2]. We are also applying MixMOBO for the optimization of vertical axis wind turbines [171]. We note here that with continuous variables and a single objective function, our current study is a specialized case for MixMOBO algorithm since it can optimize mixed-variable and multi-objective problems.

We apply our Bayesian optimization algorithm in tandem with DbM here for the shape optimization of a draft tube/hub. As an additional refinement to our optimization algorithm for the draft tube/hub problem, we use a hedge strategy that we call HedgeMO [1] in which multiple acquisition functions are used [43]. We use the Upper Confidence Bound, Expected Improvement, Probability of Improvement, and Stochastic Monte-Carlo [18, 1] acquisition functions.

To determine the number of iterations of the Bayesian optimization algorithm, or epochs (where each epoch determines the C_{prm} of 5 morphed draft tube/hub designs in parallel), required to likely find the optimal design within the design space, we carried out optimizations

of a suite of test functions with different properties and whose maximum values could be determined analytically. We optimized the Spherical, Rastringin, Syblinski-Tang and Amalgamated functions, all of which are standard functions used to test optimization schemes [52] with the exception of the Amalgamated function, which is novel and created by us to mimic some of the properties of the draft tube/hub design space. The test functions are given in Table 4.4 of the Appendix, and the results of our optimization method are shown in Fig. 4.8 of the same Appendix.

MixMOBO, like most Bayesian optimization schemes, needs to be initialized with evaluations of random designs. Based on the results shown in the Appendix on Test Functions, we use 50 random evaluations, and 75 epochs with 5 parallel batch evaluations per epoch to optimize our 6D design space.

4.5 Results

We searched for the morphed draft tube/hub design with the maximum C_{prm} in our 6-dimensional design space using MixMOBO with 50 initial random designs and 75 epochs with 5 design evaluations per epoch. The results are shown in Fig. 4.7 and Table 4.2. The Normalized C_{prm} used in the figure is defined as:
$$\frac{[(C_{prm} \text{ of Current Epoch's Optimum}) - (C_{prm} \text{ of Sharp heel draft tube with the baseline 1 hub})]}{[(C_{prm} \text{ of the Optimum design found by MixMOBO}) - (C_{prm} \text{ of Sharp heel with baseline 1 hub})]}.$$

We note here that our CFD simulations are run to convergence, as explained in Section 4.4. If we ran our simulations for only one advective time, we would potentially have gotten a higher C_{prm} value, which would not be converged and erroneous, as shown in Fig. 4.5.

Figure 4.7 and the Table show that the C_{prm} of the sharp heel draft tube with the baseline hub 1 is significantly lower than the C_{prm} of the best (and second and third best) design found with MixMOBO. In fact, the C_{prm} of the best of the 50 random designs that were used to initialize MixMOBO was better than that of the sharp heel draft tube, which shows the strength of our Design-by-Morphing approach. The sharp heel draft tube had the best C_{prm} (when coupled with hub shape 1) of all the baseline draft tubes. We note that the C_{prm} of the optimal design found by MixMOBO is significantly better than the Sharp-Heel draft tube, which is the draft tube of choice for Kaplan turbines around the world [229].

Note that all three of the morphed draft tube/hub designs in Table 4.2 have at least one negative weight, meaning that they are extrapolations, rather than interpolations of the baseline shapes. Generally, extrapolation is not possible with conventional design techniques. The Design-by-Morphing (DbM) response surface is highly sensitive to any changes in the values of the weights, which makes the DbM design space highly non-convex. We found that with each new epoch, small changes in the weights led to significant changes in the overall shape of the morphed draft tube/hub. This characteristic means that the DbM space needs to be optimized with high precision where very small changes in the weights need to be considered. In particular, if we were to consider the value of each weight to be a

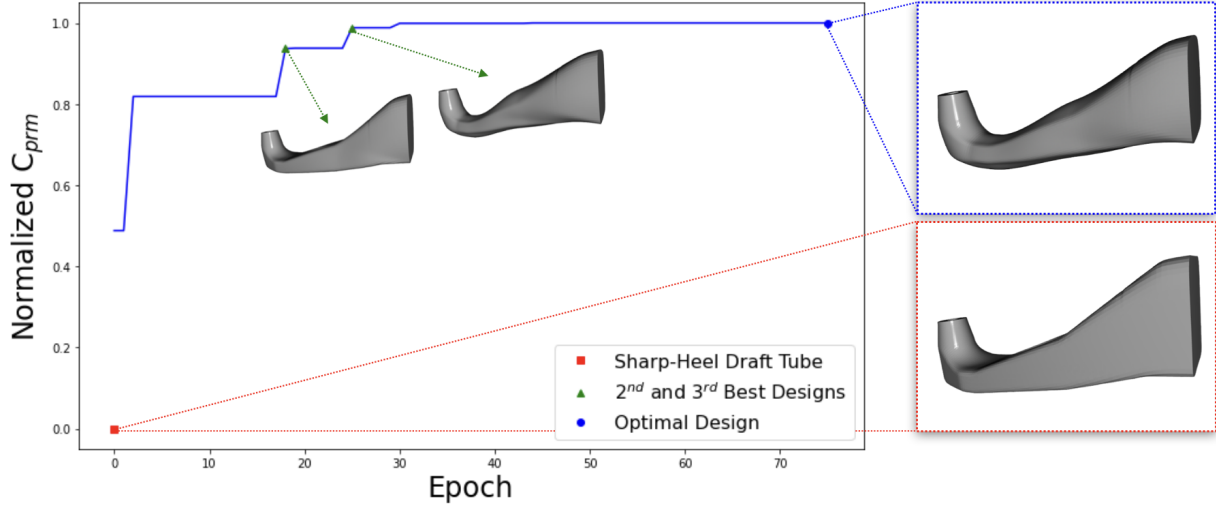


Figure 4.7: Optimization progress showing the Normalized C_{prm} of the current epoch's optimum as a function of the epoch number. The Normalized C_{prm} is defined in the text. By definition, the Normalized C_{prm} must increase monotonically with the epoch number. The morphed draft tube/hub with the second highest, and third highest C_{prm} 's are shown within the main figure. The designs that are illustrated in the panels to the right of the main figure are the shape of the morphed draft tube/hub design with the highest C_{prm} (framed with a blue box) and also the sharp heel draft tube with the first baseline hub (framed with a red box). The main figure begins at epoch zero and the Normalized C_{prm} at epoch 0 is that of the design of 50 random designs used to initialize MixMOBO with the highest C_{prm} .

Design	Epoch	ω_1	ω_2	ω_3	ω_4	ω_5	β_1	C_{prm}
Sharp-Heel		1	0	0	0	0	0.5	0.9370
3 rd Best	19	0.86	-0.17	0.40	-0.21	0.98	0.98	0.9607
2 rd Best	26	0.88	-0.24	0.26	-0.17	0.99	0.92	0.9617
Optimal	30	0.90	-0.14	0.22	-0.25	0.99	0.99	0.9620

Table 4.2: C_{prm} and DbM weights of draft tube/hub designs. The designs listed are those with the best, second best, and third best C_{prm} 's found with MixMOBO. We also show the epoch number when the design was found by MixMOBO. (See Fig. 4.7.) In addition, the sharp heel draft tube with the baseline 1 hub is listed.

discrete, rather than a continuous variable, then the design space would need to have a very large number of discrete variables. Since we are treating the weights here as continuous variables, the sensitive dependence of the morphed shape on the weight values means that most conventional optimization schemes, the search may get “stuck” at a local maximum and will fail to find the global maximum. Bayesian Optimization algorithms, including MixMOBO, that are designed to search for the global maximum of C_{prm} , will tend to over-explore the region of the design space around a local maximum. This effect, however, is less pronounced for BO based schemes since they do not depend upon the gradient of the surrogate surfaces. MixMOBO also uses “mutation” to get out of the local maxima. Our numerical experiments with MixMOBO show that it tends to not get stuck at local maxima, and it finds the global maximum [1].

4.6 Conclusion

We have introduced a novel systematic shape optimization framework using the Design-by-Morphing (DbM) technique with a Mixed-variable Multi-Objective Bayesian Optimization (MixMOBO) algorithm. As a proof-of-concept of this optimization framework, we found the shape of a draft tube/hub for a hydrokinetic turbine such that the coefficient of pressure recovery C_{prm} was maximized. Our Design-by-Morphing (DbM) technique creates novel shapes from interpolations and extrapolations of pre-existing designs and/or new ones, and the shapes we create are free from designer’s biases. Most existing design strategies cannot extrapolate among designs, and this limits their abilities to create radically different designs.

The design space created with the Design-by-Morphing framework was searched by our novel MixBOBO Bayesian Optimization. This search method is especially useful when the property that is being optimized, in this case the C_{prm} , is costly to compute (or find experimentally). In these cases, it is necessary for the search algorithm to find the optimum design with as few evaluations of C_{prm} as possible. Using MixMOBO, we successfully found a design with large C_{prm} (certainly a local maximum in the design space, and possibly the global maximum) with 30 epochs, or 200 evaluations of a design’s C_{prm} .

The draft tube/hub design found here has a significantly better coefficient of mean pressure recovery than the sharp-heel draft tube, which is the most commonly used draft tube for hydrokinetic turbines. More significantly, we have shown that the design framework used here, which can easily be generalized to a number of engineering design optimization problems, with the current study providing a proof-of-concept for our optimization framework. DbM provides a methodology to create a bias and constraint free design space which also allows extrapolation from the existing designs and can yield radical shapes. MixMOBO provides a global optimization strategy to optimize design spaces where evaluating each candidate design is extremely expensive, the design space contains mixed variables and/or multiple objectives, rendering conventional optimization techniques intractable. The framework can be used for optimizing expensive black-box shape optimization problems

such as vertical-axis wind turbines [171], high-performance airfoils [3] and architected meta-materials [1].

4.7 Additional Details for Draft-Tube and Hub Optimization

MixMOBO Convergence Tests

The test functions that we used to determine the number of epochs that are needed for MixMOBO to likely find the maximum C_{prm} are defined in Table ???. Similar to the design space of the draft tube/hub, each test function was tested with 6 dimensions.

Figure 4.8 shows how MixMOBO approaches the global maximum of each of the four test functions as a function of epoch. Two of the four test functions find the global maximum within 75 epochs.

The optimization was terminated after 75 epochs because the test functions reached reasonable convergence, other than Rastringin function, which is know to have a a multitude of local optima, as shown in Fig. 4.8. We could only afford 75 epochs worth of CFD evaluations of candidate draft tube designs in our computational budget so the evaluations were not carried past 75 epochs. As depicted in our Results, the draft tube optimization converged to the final design in 30 epochs.

Table 4.3: Benchmark test functions

Name	Objective Functions	Notes
Spherical	$f(\vec{w}) = -w_i^2, w_i \in (-10, 10)$	convex
Rastrigin	$f(\vec{w}) = -[10 + w_i^2 - 10 \cos(2\pi w_i)], w_i \in (-5.12, 5.12)$	non-convex
Syblinski-Tang	$f(\vec{w}) = -\frac{w_i^4 - 16w_i^2 + 5w_i}{2}, w_i \in (-5, 5)$	non-convex
Amalgamated	$f(\vec{w}) = \sum_{i=1}^D -g(w_i)$ $g(w_i) = \begin{cases} -\sin(w_i), & \text{if } k = 0, w_i \in (0, \pi) \\ \frac{w_i^4 - 16w_i^2 + 5w_i}{2}, & \text{if } k = 1, w_i \in (-5, 5) \\ w_i^2, & \text{if } k = 2, w_i \in (-10, 10) \\ [10 + w_i^2 - 10 \cos(2\pi w_i)], & \text{if } k = 3, w_i \in (-5.12, 5.12) \\ [100(w_i - w_{i-1}^2)^2 + (1 - w_i)^2], & \text{if } k = 4, w_i \in (-2, 2) \\ - \cos(w_i) , & \text{if } k = 5, w_i \in (-\pi/2, \pi/2) \\ w_i, & \text{if } k = 6, w_i \in (-30, 30) \end{cases}$ $k = \text{mod}(i - 1, 7), i = 1, \dots, n$	non-convex, non-uniform, anisotropic

Table 4.4: Details about the test functions, other than the Amalgamated function, are given in Tušar, Brockhoff, and Hansen [52]. All of the test functions have known global maxima. We created the Amalgamated function, a piece-wise function formed from commonly used analytical test functions with different features. The Amalgamated function is non-convex and anisotropic (as is the design space of the draft tube/hub), unlike the other test functions listed here, which are isotropic. These other test functions are commonly used for testing optimization algorithms. Similar to the optimization of the draft tube/hub, each test function here has 6 dimensions.




Figure 4.8: Benchmarks for determining the number of epochs of MixMOBO needed to find the likely global maximum. The number of random evaluations of the function to be maximize that are needed to initialize MixMOBO, was set at 50. Our MixMOBO algorithm was then run with 5 parallel batch evaluations per epoch. The evaluation budget was 425 black-box function evaluations including the 50 initial random evaluations. Thus, we ran the code for $75 = (425 - 50)/5$ epochs. Because each optimization begins with random evaluations, the algorithm was run 5 times for each test function to determine the average and standard deviation of the number epochs needed to find the global maximum. The mean Normalized Reward, defined as $(current\ optimum - random\ sampling\ optimum) / (global\ optimum - random\ sampling\ optimum)$ is plotted along with the $0.2 \times$ standard deviation plotted as colored bands. A Normalized Reward of unity means that the algorithm has successfully found the global maximum.

Bibliography

- [1] Haris Moazam Sheikh and Philip S. Marcus. Bayesian optimization for multi-objective mixed-variable problems, 2022. arXiv: 2201.12767 [cs.LG].
- [2] Zacharias Vangelatos, Haris Moazam Sheikh, Philip S. Marcus, Costas P. Grigoropoulos, Victor Z. Lopez, George Flamourakis, and Maria Farsari. Strength through defects: a novel bayesian approach for the optimization of architected materials. *Science Advances*, 7(41), 2021. DOI: 10.1126/sciadv.abk2218.
- [3] Haris Moazam Sheikh, Sangjoon Lee, Jinge Wang, and Philip S. Marcus. Airfoil optimization using design-by-morphing, 2022. DOI: 10.48550/ARXIV.2207.11448. URL: <https://arxiv.org/abs/2207.11448>.
- [4] Haris Moazam Sheikh, Tess A. Callan, Kealan J. Hennessy, and Philip S. Marcus. Optimization of the shape of a hydrokinetic turbine’s draft tube and hub assembly using design-by-morphing with bayesian optimization, 2022. DOI: 10.48550/ARXIV.2207.11451. URL: <https://arxiv.org/abs/2207.11451>.
- [5] Peter I. Frazier and Jialei Wang. Bayesian optimization for materials design. *Springer Series in Materials Science*:45–75, 2015. ISSN: 2196-2812. DOI: 10.1007/978-3-319-23871-5_3.
- [6] Desai Chen, Mélina Skouras, Bo Zhu, and Wojciech Matusik. Computational discovery of extremal microstructure families. *Science Advances*, 4(1):eaao7005, 2018.
- [7] Wen Chen, Seth Watts, Julie A Jackson, William L Smith, Daniel A Tortorelli, and Christopher M Spadaccini. Stiff isotropic lattices beyond the maxwell criterion. *Science Advances*, 5(9):eaaw1937, 2019.
- [8] Lucas A Shaw, Frederick Sun, Carlos M Portela, Rodolfo I Barranco, Julia R Greer, and Jonathan B Hopkins. Computationally efficient design of directionally compliant metamaterials. *Nature Communications*, 10(1):1–13, 2019.
- [9] Jian Song, Yuejiao Wang, Wenzhao Zhou, Rong Fan, Bin Yu, Yang Lu, and Lixiao Li. Topology optimization-guided lattice composites and their mechanical characterizations. *Composites Part B: Engineering*, 160:402–411, 2019.
- [10] Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical bayesian optimization of machine learning algorithms, 2012. arXiv: 1206.2944 [stat.ML].

- [11] Yutian Chen, Aja Huang, Ziyu Wang, Ioannis Antonoglou, Julian Schrittwieser, David Silver, and Nando de Freitas. Bayesian optimization in alphago. *CoRR*, abs/1812.06855, 2018. arXiv: 1812.06855.
- [12] ChangYong Oh, Efstratios Gavves, and Max Welling. BOCK : Bayesian optimization with cylindrical kernels. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 3868–3877. PMLR, 2018.
- [13] Edward Pyzer-Knapp. Bayesian optimization for accelerated drug discovery. *IBM Journal of Research and Development*, PP:1–1, November 2018. DOI: 10.1147/JRD.2018.2881731.
- [14] Ksenia Korovina, Sailun Xu, Kirthevasan Kandasamy, Willie Neiswanger, Barnabas Poczos, Jeff Schneider, and Eric Xing. Chembo: bayesian optimization of small organic molecules with synthesizable recommendations. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 3393–3403. PMLR, 2020.
- [15] Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes: theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9(8):235–284, 2008.
- [16] J. Mockus. Application of bayesian approach to numerical methods of global and stochastic optimization. *Journal of Global Optimization*, 4:347–365, 1994.
- [17] Christopher K. I. Williams Carl Edward Rasmussen. *Gaussian processes for machine learning*. MIT Press, 2005.
- [18] Eric Brochu, Vlad M. Cora, and Nando de Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning, 2010. arXiv: 1012.2599 [cs.LG].
- [19] Daniel Golovin, Benjamin Solnik, Subhdeep Moitra, Greg Kochanski, John Karro, and D. Sculley. Google vizier: a service for black-box optimization. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '17, pages 1487–1495, Halifax, NS, Canada. Association for Computing Machinery, 2017. ISBN: 9781450348874. DOI: 10.1145/3097983.3098043.
- [20] Binxin Ru, Ahsan S. Alvi, Vu Nguyen, Michael A. Osborne, and Stephen J Roberts. Bayesian optimisation over multiple continuous and categorical inputs, 2020. arXiv: 1906.08878 [stat.ML].
- [21] Shivapratap Gopakumar, Sunil Gupta, Santu Rana, Vu Nguyen, and Svetha Venkatesh. Algorithmic assurance: an active approach to algorithmic testing using bayesian optimisation. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

- [22] Dang Nguyen, Sunil Gupta, Santu Rana, Alistair Shilton, and Svetha Venkatesh. Bayesian optimization for categorical and category-specific continuous inputs, 2019. arXiv: 1911.12473 [cs.LG].
- [23] Peter Z G Qian, Huaqing Wu, and CF Jeff Wu. Gaussian process models for computer experiments with qualitative and quantitative factors. *Technometrics*, 50(3):383–396, 2008.
- [24] Qiang Zhou, Peter ZG Qian, and Shiyu Zhou. A simple approach to emulation for computer models with qualitative and quantitative factors. *Technometrics*, 53(3):266–273, 2011.
- [25] Yichi Zhang, Daniel W Apley, and Wei Chen. Bayesian optimization for materials design with mixed quantitative and qualitative variables. *Scientific Reports*, 10(1):1–13, 2020.
- [26] Aryan Deshwal and Janardhan Rao Doppa. Combining latent space and structured kernels for bayesian optimization over combinatorial spaces. *CoRR*, abs/2111.01186, 2021. arXiv: 2111.01186.
- [27] Aryan Deshwal, Syrine Belakaria, and Janardhan Rao Doppa. Bayesian optimization over hybrid spaces, 2021. arXiv: 2106.04682 [cs.LG].
- [28] Changyong Oh, Efstratios Gavves, and Max Welling. Mixed variable bayesian optimization with frequency modulated kernels, 2021. arXiv: 2102.12792 [stat.ML].
- [29] Ricardo Baptista and Matthias Poloczek. Bayesian optimization of combinatorial structures, 2018. arXiv: 1806.08838 [stat.ML].
- [30] Julien Pelamatti, Loïc Brevault, Mathieu Balesdent, El-Ghazali Talbi, and Yannick Guerin. Efficient global optimization of constrained mixed variable problems, 2018. arXiv: 1806.03975 [math.OC].
- [31] Changyong Oh, Jakub M. Tomczak, Efstratios Gavves, and Max Welling. Combinatorial bayesian optimization using the graph cartesian product, 2019. arXiv: 1902.00448 [stat.ML].
- [32] Eduardo C. Garrido-Merchán and Daniel Hernández-Lobato. Dealing with categorical and integer-valued variables in bayesian optimization with gaussian processes. *Neurocomputing*, 380:20–35, 2020. ISSN: 0925-2312. DOI: 10.1016/j.neucom.2019.11.004.
- [33] Marius Lindauer, Katharina Eggensperger, Matthias Feurer, André Biedenkapp, Difan Deng, Carolin Benjamins, René Sass, and Frank Hutter. Smac3: a versatile bayesian optimization package for hyperparameter optimization, 2021. arXiv: 2109.09831 [cs.LG].

- [34] James Bergstra, Daniel Yamins, and David Cox. Making a science of model search: hyperparameter optimization in hundreds of dimensions for vision architectures. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 115–123, Atlanta, Georgia, USA. PMLR, 2013.
- [35] Erik Daxberger, Anastasia Makarova, Matteo Turchetta, and Andreas Krause. Mixed-variable bayesian optimization. *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, 2020. DOI: 10.24963/ijcai.2020/365.
- [36] Louis C. Tiao, Aaron Klein, Matthias Seeger, Edwin V. Bonilla, Cedric Archambeau, and Fabio Ramos. Bore: bayesian optimization by density-ratio estimation, 2021. arXiv: 2102.09009 [cs.LG].
- [37] Maximilian Balandat, Brian Karrer, Daniel R. Jiang, Samuel Daulton, Benjamin Letham, Andrew Gordon Wilson, and Eytan Bakshy. Botorch: a framework for efficient monte-carlo bayesian optimization, 2020. arXiv: 1910.06403 [cs.LG].
- [38] C.M. Fonseca, L. Paquete, and M. Lopez-Ibanez. An improved dimension-sweep algorithm for the hypervolume indicator. In *2006 IEEE International Conference on Evolutionary Computation*, pages 1157–1163, 2006. DOI: 10.1109/CEC.2006.1688440.
- [39] Samuel Daulton, Maximilian Balandat, and Eytan Bakshy. Differentiable expected hypervolume improvement for parallel multi-objective bayesian optimization, 2020. arXiv: 2006.05078 [stat.ML].
- [40] Samuel Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy. Multi-objective bayesian optimization over high-dimensional search spaces, 2021. arXiv: 2109.10964 [cs.LG].
- [41] Shinya Suzuki, Shion Takeno, Tomoyuki Tamura, Kazuki Shitara, and Masayuki Karasuyama. Multi-objective bayesian optimization using pareto-frontier entropy, 2020. arXiv: 1906.00127 [cs.LG].
- [42] Leshi Shu, Ping Jiang, Xinyu Shao, and Yan Wang. A New Multi-Objective Bayesian Optimization Formulation With the Acquisition Function for Convergence and Diversity. *Journal of Mechanical Design*, 142(9), March 2020. ISSN: 1050-0472. DOI: 10.1115/1.4046508. 091703.
- [43] Eric Brochu, Matthew W. Hoffman, and Nando de Freitas. Portfolio allocation for bayesian optimization, 2011. arXiv: 1009.5419 [cs.LG].
- [44] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: nsga-ii. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002. DOI: 10.1109/4235.996017.
- [45] K. P. Murphy. *Machine learning: a probabilistic perspective*. MIT Press, 2012.

- [46] Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias W. Seeger. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 58(5):3250–3265, 2012. ISSN: 1557-9654. DOI: 10.1109/tit.2011.2182033.
- [47] GPyOptAuthors. GPyOpt: a bayesian optimization framework in python. <http://github.com/SheffieldML/GPyOpt>, 2016.
- [48] Scikit-learn. Scikit-optimize. <https://scikit-optimize.github.io/stable/>, 2021.
- [49] James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In *Proceedings of the 24th International Conference on Neural Information Processing Systems, NIPS’11*, pages 2546–2554, Granada, Spain. Curran Associates Inc., 2011. ISBN: 9781618395993.
- [50] Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In Carlos A. Coello Coello, editor, *Learning and Intelligent Optimization*, pages 507–523, Berlin, Heidelberg. Springer Berlin Heidelberg, 2011. ISBN: 978-3-642-25566-3.
- [51] Yingjie Hu, Jianqiang Hu, Yifan Xu, Fengchun Wang, and Rong Cao. Contamination control in food supply chain. In pages 2678–2681, January 2011. DOI: 10.1109/WSC.2010.5678963.
- [52] Tea Tušar, Dimo Brockhoff, and Nikolaus Hansen. Mixed-integer benchmark problems for single- and bi-objective optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO ’19*, pages 718–726, Prague, Czech Republic. Association for Computing Machinery, 2019. ISBN: 9781450361118. DOI: 10.1145/3321707.3321868.
- [53] Stuart Kauffman and Simon Levin. Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology*, 128(1):11–45, 1987. ISSN: 0022-5193. DOI: [https://doi.org/10.1016/S0022-5193\(87\)80029-2](https://doi.org/10.1016/S0022-5193(87)80029-2).
- [54] Rui Li, Michael Emmerich, Jeroen Eggermont, Ernst Bovenkamp, Thomas Bäck, Jouke Dijkstra, and Johan Reiber. Mixed-integer nk landscapes. In volume 4193, pages 42–51, January 2006. ISBN: 978-3-540-38990-3. DOI: 10.1007/11844297_5.
- [55] Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele. Comparison of multiobjective evolutionary algorithms: empirical results. *Evol. Comput.*, 8(2):173–195, 2000. ISSN: 1063-6560. DOI: 10.1162/106365600568202.
- [56] Xiaoyu Zheng, Howon Lee, Todd H Weisgraber, Maxim Shusteff, Joshua DeOtte, Eric B Duoss, Joshua D Kuntz, Monika M Biener, Qi Ge, Julie A Jackson, et al. Ultralight, ultrastiff mechanical metamaterials. *Science*, 344(6190):1373–1377, 2014.
- [57] Minh-Son Pham, Chen Liu, Iain Todd, and Jedsada Lertthanasarn. Damage-tolerant architected materials inspired by crystal microstructure. *Nature*, 565(7739):305–311, 2019.

- [58] Xuan Zhang, Andrey Vyatskikh, Huajian Gao, Julia R Greer, and Xiaoyan Li. Lightweight, flaw-tolerant, and ultrastrong nanoarchitected carbon. *Proceedings of the National Academy of Sciences*, 116(14):6665–6672, 2019.
- [59] Xiaoxing Xia, Arman Afshar, Heng Yang, Carlos M Portela, Dennis M Kochmann, Claudio V Di Leo, and Julia R Greer. Electrochemically reconfigurable architected materials. *Nature*, 573(7773):205–213, 2019.
- [60] Jian Song, Wenzhao Zhou, Yuejiao Wang, Rong Fan, Yinchu Wang, Junying Chen, Yang Lu, and Lixiao Li. Octet-truss cellular materials for improved mechanical properties and specific energy absorption. *Materials & Design*, 173:107773, 2019.
- [61] Jiayi Song, Christos Michas, Christopher S Chen, Alice E White, and Mark W Grinstaff. From simple to architecturally complex hydrogel scaffolds for cell and tissue engineering applications: opportunities presented by two-photon polymerization. *Advanced Healthcare Materials*, 9(1):1901217, 2020.
- [62] Jens Bauer, Almut Schroer, Ruth Schwaiger, and Oliver Kraft. Approaching theoretical strength in glassy carbon nanolattices. *Nature Materials*, 15(4):438–443, 2016.
- [63] Lucas R Meza, Satyajit Das, and Julia R Greer. Strong, lightweight, and recoverable three-dimensional ceramic nanolattices. *Science*, 345(6202):1322–1326, 2014.
- [64] JB Berger, HNG Wadley, and RM McMeeking. Mechanical metamaterials at the theoretical limit of isotropic elastic stiffness. *Nature*, 543(7646):533–537, 2017.
- [65] Thomas Tancogne-Dejean, Marianna Diamantopoulou, Maysam B Gorji, Colin Bonatti, and Dirk Mohr. 3d plate-lattices: an emerging class of low-density metamaterial exhibiting optimal isotropic stiffness. *Advanced Materials*, 30(45):1803334, 2018.
- [66] Z Vangelatos, K Komvopoulos, and C Grigoropoulos. Regulating the mechanical behavior of metamaterial microlattices by tactical structure modification. *Journal of the Mechanics and Physics of Solids*:104112, 2020.
- [67] Qiang Geng, Dien Wang, Pengfei Chen, and Shih-Chi Chen. Ultrafast multi-focus 3-D nano-fabrication based on two-photon polymerization. *Nature Communications*, 10:2179, May 2019. ISSN: 2041-1723.
- [68] Jordan R Raney and Jennifer A Lewis. Printing mesoscale architectures. *Mrs Bulletin*, 40(11):943–950, 2015.
- [69] Linas Jonušauskas, Darius Gailevičius, Sima Rekštytė, Tommaso Baldacchini, Saulius Juodkazis, and Mangirdas Malinauskas. Mesoscale laser 3d printing. *Optics express*, 27(11):15205–15221, 2019.
- [70] Paul Somers, Zihao Liang, Jason E Johnson, Bryan W Boudouris, Liang Pan, and Xianfan Xu. Rapid, continuous projection multi-photon 3d printing enabled by spatiotemporal focusing of femtosecond pulses. *Light: Science & Applications*, 10(1):1–11, 2021.

- [71] Jae-Hwang Lee, Jonathan P Singer, and Edwin L Thomas. Micro-/nanostructured mechanical metamaterials. *Advanced materials*, 24(36):4782–4810, 2012.
- [72] James Utama Surjadi, Libo Gao, Huifeng Du, Xiang Li, Xiang Xiong, Nicholas Xuanlai Fang, and Yang Lu. Mechanical metamaterials and their engineering applications. *Advanced Engineering Materials*, 21(3):1800864, 2019.
- [73] Minh-Son Pham, Chen Liu, Iain Todd, and Jedsada Lertthanasarn. Damage-tolerant architected materials inspired by crystal microstructure. *Nature*, 565(7739):305–311, January 2019. ISSN: 1476-4687.
- [74] Xuan Zhang, Andrey Vyatskikh, Huajian Gao, Julia R. Greer, and Xiaoyan Li. Lightweight, flaw-tolerant, and ultrastrong nanoarchitected carbon. en. *Proceedings of the National Academy of Sciences*, 116(14), April 2019. ISSN: 0027-8424, 1091-6490. (Visited on 01/21/2022).
- [75] Lucas R. Meza, Satyajit Das, and Julia R. Greer. Strong, lightweight, and recoverable three-dimensional ceramic nanolattices. EN. *Science*, September 2014. (Visited on 01/21/2022).
- [76] Xiaoyu Zheng, Howon Lee, Todd H. Weisgraber, Maxim Shusteff, Joshua DeOtte, Eric B. Duoss, Joshua D. Kuntz, Monika M. Biener, Qi Ge, Julie A. Jackson, Sergei O. Kucheyev, Nicholas X. Fang, and Christopher M. Spadaccini. Ultralight, ultrastiff mechanical metamaterials. EN. *Science*, June 2014. (Visited on 01/21/2022).
- [77] Haoran Fu, Kewang Nan, Wubin Bai, Wen Huang, Ke Bai, Luyao Lu, Chaoqun Zhou, Yunpeng Liu, Fei Liu, Juntong Wang, Mengdi Han, Zheng Yan, Haiwen Luan, Yijie Zhang, Yutong Zhang, Jianing Zhao, Xu Cheng, Moyang Li, Jung Woo Lee, Yuan Liu, Daining Fang, Xiuling Li, Yonggang Huang, Yihui Zhang, and John A. Rogers. Morphable 3D mesostructures and microelectronic devices by multistable buckling mechanics. en. *Nature Materials*, 17(3), March 2018. ISSN: 1476-4660. (Visited on 01/21/2022).
- [78] Huachen Cui, Ryan Hensleigh, Desheng Yao, Deepam Maurya, Prashant Kumar, Min Gyu Kang, Shashank Priya, and Xiaoyu (Rayne) Zheng. Three-dimensional printing of piezoelectric materials with designed anisotropy and directional response. en. *Nature Materials*, 18(3), March 2019. ISSN: 1476-4660. (Visited on 01/21/2022).
- [79] Xiaoxing Xia, Arman Afshar, Heng Yang, Carlos M. Portela, Dennis M. Kochmann, Claudio V. Di Leo, and Julia R. Greer. Electrochemically reconfigurable architected materials. en. *Nature*, 573(7773), September 2019. ISSN: 1476-4687. (Visited on 01/21/2022).
- [80] D Beli, JRF Arruda, and M Ruzzene. Wave propagation in elastic metamaterial beams and plates with interconnected resonators. *International Journal of Solids and Structures*, 139:105–120, 2018.

- [81] Jijie Huang, Tiening Jin, Shikhar Misra, Han Wang, Zhimin Qi, Yaomin Dai, Xing Sun, Leigang Li, Joseph Okkema, Hou-Tong Chen, et al. Tailorable optical response of au–linbo3 hybrid metamaterial thin films for optical waveguide applications. *Advanced Optical Materials*, 6(19):1800510, 2018.
- [82] Sha Yin, Weihua Guo, Huitian Wang, Yao Huang, Ruiheng Yang, Zihan Hu, Dianhao Chen, Jun Xu, and Robert O Ritchie. Strong and tough bioinspired additive-manufactured dual-phase mechanical metamaterial composites. *Journal of the Mechanics and Physics of Solids*, 149:104341, 2021.
- [83] Liang Wang, Ming-Ran An, and Hai-Tao Liu. Compression spin bio-inspired arm: a conceptual model based on compression–torsion cubic mechanical metamaterials with variable cross-section. *Extreme Mechanics Letters*, 41:101069, 2020.
- [84] Zian Jia, Matheus C Fernandes, Zhifei Deng, Ting Yang, Qiuting Zhang, Alfie Lethbridge, Jie Yin, Jae-Hwang Lee, Lin Han, James C Weaver, et al. Microstructural design for mechanical–optical multifunctionality in the exoskeleton of the flower beetle *torynorrhina flammea*. *Proceedings of the National Academy of Sciences*, 118(25), 2021.
- [85] Lingxi Huang, Yuping Duan, Jia Liu, Yuansong Zeng, Guojia Ma, Huifang Pang, Shao-hua Gao, and Weiping Zhang. Bioinspired gyrotropic metamaterials with multifarious wave adaptability and multifunctionality. *Advanced Optical Materials*, 8(12):2000012, 2020.
- [86] Matheus C Fernandes, Joanna Aizenberg, James C Weaver, and Katia Bertoldi. Mechanically robust lattices inspired by deep-sea glass sponges. *Nature Materials*, 20(2):237–241, 2021.
- [87] Andrew Gross, Panos Pantidis, Katia Bertoldi, and Simos Gerasimidis. Correlation between topology and elastic properties of imperfect truss-lattice materials. *Journal of the Mechanics and Physics of Solids*, 124:577–598, 2019.
- [88] Fani Derveni, Andrew J Gross, Kara D Peterman, and Simos Gerasimidis. Postbuckling behavior and imperfection sensitivity of elastic–plastic periodic plate-lattice materials. *Extreme Mechanics Letters*, 50:101510, 2022.
- [89] R. Schwaiger, L.R. Meza, and X. Li. The extreme mechanics of micro- and nanoarchitected materials. *MRS Bulletin*, 44(10):758–765, 2019. DOI: 10.1557/mrs.2019.230.
- [90] Oraib Al-Ketan, Rachid Rezgui, Reza Rowshan, Huifeng Du, Nicholas X. Fang, and Rashid K. Abu Al-Rub. Microarchitected Stretching-Dominated Mechanical Metamaterials with Minimal Surface Topologies. en. *Advanced Engineering Materials*, 20(9), 2018. ISSN: 1527-2648. (Visited on 01/21/2022).
- [91] Michael F Ashby. The properties of foams and lattices. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 364(1838):15–30, 2006.

- [92] Carlos M Portela, A Vidyasagar, Sebastian Krödel, Tamara Weissenbach, Daryl W Yee, Julia R Greer, and Dennis M Kochmann. Extreme mechanical resilience of self-assembled nanolabyrinthine materials. *Proceedings of the National Academy of Sciences*, 117(11):5686–5693, 2020.
- [93] Jan-Hendrik Bastek, Siddhant Kumar, Bastian Telgen, Raphaël N Glaesener, and Dennis M Kochmann. Inverting the structure–property map of truss metamaterials by deep learning. *Proceedings of the National Academy of Sciences*, 119(1), 2022.
- [94] Li Zheng, Siddhant Kumar, and Dennis M Kochmann. Data-driven topology optimization of spinodoid metamaterials with seamlessly tunable anisotropy. *Computer Methods in Applied Mechanics and Engineering*, 383:113894, 2021.
- [95] Qingping Ma, Lei Zhang, Junhao Ding, Shuo Qu, Jin Fu, Mingdong Zhou, Ming Wang Fu, Xu Song, and Michael Yu Wang. Elastically-isotropic open-cell minimal surface shell lattices with superior stiffness via variable thickness design. *Additive Manufacturing*, 47:102293, 2021.
- [96] Qingping Ma, Zhenjun Yan, Lei Zhang, and Michael Yu Wang. The family of elastically isotropic stretching-dominated cubic truss lattices. *International Journal of Solids and Structures*:111451, 2022.
- [97] Jiawei Feng, Bo Liu, Zhiwei Lin, and Jianzhong Fu. Isotropic porous structure design methods based on triply periodic minimal surfaces. *Materials & Design*, 210:110050, 2021.
- [98] Ming Lei, Wei Hong, Zeang Zhao, Craig Hamel, Mingji Chen, Haibao Lu, and H Jerry Qi. 3d printing of auxetic metamaterials with digitally reprogrammable shape. *ACS applied materials & interfaces*, 11(25):22768–22776, 2019.
- [99] Jianhu Shen, Shiwei Zhou, Xiaodong Huang, and Yi Min Xie. Simple cubic three-dimensional auxetic metamaterials. *physica status solidi (b)*, 251(8):1515–1522, 2014.
- [100] Romik Khajetourian and Dennis M Kochmann. A continuum description of substrate-free dissipative reconfigurable metamaterials. *Journal of the Mechanics and Physics of Solids*, 147:104217, 2021.
- [101] Mario Spagnuolo, Mustafa Yıldızdağ, Ugo Andreaus, and Antonio Cazzani. Are higher-gradient models also capable of predicting mechanical behavior in the case of wide-knit pantographic structures? *MATHEMATICS AND MECHANICS OF SOLIDS*, 2020.
- [102] Gourav Agrawal, Abhinav Gupta, Rajib Chowdhury, and Anupam Chakrabarti. Robust topology optimization of negative poisson’s ratio metamaterials under material uncertainty. *Finite Elements in Analysis and Design*, 198:103649, 2022.
- [103] Yafeng Chang, Hui Wang, and Qinxi Dong. Machine learning-based inverse design of auxetic metamaterial with zero poisson’s ratio. *Materials Today Communications*:103186, 2022.

- [104] Weifeng Jiang, Yangyang Zhu, Guofu Yin, Houhong Lu, Luofeng Xie, and Ming Yin. Dispersion relation prediction and structure inverse design of elastic metamaterials via deep learning. *Materials Today Physics*:100616, 2022.
- [105] Lingling Wu, Lei Liu, Yong Wang, Zirui Zhai, Houlong Zhuang, Deepakshyam Krishnaraju, Qianxuan Wang, and Hanqing Jiang. A machine learning-based method to design modular metamaterials. *Extreme Mechanics Letters*, 36:100657, 2020.
- [106] Adithya Challapalli, Dhruvil Patel, and Gouqiang Li. Inverse machine learning framework for optimizing lightweight metamaterials. *Materials & Design*, 208:109937, 2021.
- [107] Jackson K Wilt, Charles Yang, and Grace X Gu. Accelerating auxetic metamaterial design with deep learning. *Advanced Engineering Materials*:1901266, 2020.
- [108] Anthony P Garland, Benjamin C White, Scott C Jensen, and Brad L Boyce. Pragmatic generative optimization of novel structural lattice metamaterials with machine learning. *Materials & Design*, 203:109632, 2021.
- [109] Miguel A Bessa, Piotr Glowacki, and Michael Houlder. Bayesian machine learning in metamaterial design: fragile becomes supercompressible. *Advanced Materials*, 31(48):1904845, 2019.
- [110] Jordan Matthews, Timothy Klatt, Clinton Morris, Carolyn C Seepersad, Michael Haberman, and David Shahan. Hierarchical design of negative stiffness metamaterials using a bayesian network classifier. *Journal of Mechanical Design*, 138(4), 2016.
- [111] Conner Sharpe, Carolyn Conner Seepersad, Seth Watts, and Dan Tortorelli. Design of mechanical metamaterials via constrained bayesian optimization. In *ASME 2018 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. American Society of Mechanical Engineers Digital Collection, 2018.
- [112] Francisco Dos Reis and Nikolaos Karathanasopoulos. Inverse metamaterial design combining genetic algorithms with asymptotic homogenization schemes. *International Journal of Solids and Structures*, 250:111702, 2022. ISSN: 0020-7683. DOI: <https://doi.org/10.1016/j.ijsolstr.2022.111702>. URL: <https://www.sciencedirect.com/science/article/pii/S0020768322002104>.
- [113] Andrea Bacigalupo, Giorgio Gnecco, Marco Lepidi, and Luigi Gambarotta. Machine-learning techniques for the optimal design of acoustic metamaterials, 2019. DOI: 10.48550/ARXIV.1908.10645. URL: <https://arxiv.org/abs/1908.10645>.
- [114] Z Vangelatos, K Komvopoulos, and CP Grigoropoulos. Vacancies for controlling the behavior of microstructured three-dimensional mechanical metamaterials. *Mathematics and Mechanics of Solids*, 24(2):511–524, 2019.
- [115] Vikram S Deshpande, Norman A Fleck, and Michael F Ashby. Effective properties of the octet-truss lattice material. *Journal of the Mechanics and Physics of Solids*, 49(8):1747–1769, 2001.

- [116] Emanuele Viterbo and Ezio Biglieri. Computing the voronoi cell of a lattice: the diamond-cutting algorithm. *IEEE transactions on information theory*, 42(1):161–171, 1996.
- [117] John D Clayton. *Nonlinear mechanics of crystals*, volume 177. Springer Science & Business Media, 2010.
- [118] Pingbing Ming et al. Cauchy–born rule and the stability of crystalline solids: static problems. *Archive for rational mechanics and analysis*, 183(2):241–297, 2007.
- [119] JL2412005 Ericksen. On the cauchy—born rule. *Mathematics and mechanics of solids*, 13(3-4):199–220, 2008.
- [120] A Aghaei, MJ Abdolhosseini Qomi, MT Kazemi, and AR Khoei. Stability and size-dependency of cauchy–born hypothesis in three-dimensional applications. *International Journal of Solids and Structures*, 46(9):1925–1936, 2009.
- [121] J.L. Ericksen. On the cauchy—born rule. *Mathematics and Mechanics of Solids*, 13(3-4):199–220, 2008.
- [122] Wooju Lee, Da-Young Kang, Jihwan Song, Jun Hyuk Moon, and Dongchoul Kim. Controlled unusual stiffness of mechanical metamaterials. *Scientific reports*, 6(1):1–7, 2016.
- [123] Lorna J Gibson and Michael F Ashby. *Cellular solids: structure and properties*. Cambridge University Press, 1999.
- [124] Eric Brochu, Matthew W Hoffman, and Nando de Freitas. Portfolio allocation for bayesian optimization. *arXiv preprint arXiv:1009.5419*, 2010.
- [125] Z Vangelatos, K Komvopoulos, J Spanos, M Farsari, and C Grigoropoulos. Anisotropic and curved lattice members enhance the structural integrity and mechanical performance of architected metamaterials. *International Journal of Solids and Structures*, 193:287–301, 2020.
- [126] Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2 of number 3. MIT press Cambridge, MA, 2006.
- [127] Shanqing Xu, Jianhu Shen, Shiwei Zhou, Xiaodong Huang, and Yi Min Xie. Design of lattice structures with controlled anisotropy. *Materials & Design*, 93:443–447, 2016.
- [128] Zacharias Vangelatos, Chenyang Li, Costas P. Grigoropoulos, and Kyriakos Komvopoulos. Comparison of the mechanical performance of architected three-dimensional intertwined lattices at the macro/microscale. *Extreme Mechanics Letters*, 40:100930, 2020.
- [129] Andrew N Norris. Poisson’s ratio in cubic materials. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 462(2075):3385–3405, 2006.
- [130] Masaki Yamawaki, Masato Ohnishi, Shenghong Ju, and Junichiro Shiomi. Multifunctional structural design of graphene thermoelectrics by bayesian optimization. *Science advances*, 4(6):eaar4192, 2018.

- [131] Aldair E Gongora, Bowen Xu, Wyatt Perry, Chika Okoye, Patrick Riley, Kristofer G Reyes, Elise F Morgan, and Keith A Brown. A bayesian experimental autonomous researcher for mechanical design. *Science advances*, 6(15):eaaz1708, 2020.
- [132] Tomoki Yamashita, Nobuya Sato, Hiori Kino, Takashi Miyake, Koji Tsuda, and Tamio Oguchi. Crystal structure prediction accelerated by bayesian optimization. *Physical Review Materials*, 2(1):013803, 2018.
- [133] Giovanni Zanzotto. The cauchy–born hypothesis, nonlinear elasticity and mechanical twinning in crystals. *Acta Crystallographica Section A: Foundations of Crystallography*, 52(6):839–849, 1996.
- [134] Julio Andres Iglesias Martinez, Johnny Moughames, Gwenn Ulliac, Muamer Kadic, and Vincent Laude. Three-dimensional phononic crystal with ultra-wide bandgap at megahertz frequencies. *Applied Physics Letters*, 118(6):063507, 2021.
- [135] Mark Drela. *Pros & cons of airfoil optimization*. In *Frontiers of Computational Fluid Dynamics 1998*, pages 363–381. DOI: 10.1142/9789812815774_0019.
- [136] Eric Besnard, Adeline Schmitz, Edwan Boscher, Nicolas Garcia, and Tuncer Cebeci. Two-dimensional aircraft high lift system design and optimization. *36th AIAA Aerospace Sciences Meeting and Exhibit*, January 1998. DOI: 10.2514/6.1998-123.
- [137] Alessandro Vicini and Domenico Quagliarella. Airfoil and wing design through hybrid optimization strategies. *AIAA Journal*, 37(5):634–641, May 1999. DOI: 10.2514/2.764.
- [138] Ali Elham and Michel J.L. van Tooren. Winglet multi-objective shape optimization. *Aerospace Science and Technology*, 37:93–109, August 2014. DOI: 10.1016/j.ast.2014.05.011. URL: https://www.sciencedirect.com/science/article/pii/S1270963814001060?casa_token=EokdmYaC-ZQAAAAA:tzB75VWwq3AVLKRMypezV_KPufzSJFrh1iJBC25j9UCqGXFVpfVuo9n6fkPkj8qLyp57zUAZ4vM.
- [139] J-Y Li, R Li, Y Gao, and J Huang. Aerodynamic optimization of wind turbine airfoils using response surface techniques. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 224(6):827–838, June 2010. DOI: 10.1243/09576509jpe888. URL: <https://journals.sagepub.com/doi/10.1243/09576509JPE888>.
- [140] Y P Ju and C H Zhang. Multi-point robust design optimization of wind turbine airfoil under geometric uncertainty. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 226(2):245–261, October 2011. DOI: 10.1177/0957650911426540. URL: <https://journals.sagepub.com/doi/10.1177/0957650911426540>.
- [141] A.F.P. Ribeiro, A.M. Awruch, and H.M. Gomes. An airfoil optimization technique for wind turbines. *Applied Mathematical Modelling*, 36(10):4898–4907, October 2012. DOI: 10.1016/j.apm.2011.12.026. URL: <https://www.sciencedirect.com/science/article/pii/S0307904X11008146>.

- [142] Francesco Grasso. Hybrid optimization for wind turbine thick airfoils. *53rd AIAA/ASME/ASCE/AHS Structures, Structural Dynamics and Materials Conference; BR; 20th AIAA/ASME/AHS Adaptive Structures Conference; BR; 14th AIAA*, April 2012. DOI: 10.2514/6.2012-1354.
- [143] Adam Chehouri, Rafic Younes, Adrian Ilinca, and Jean Perron. Review of performance optimization techniques applied to wind turbines. *Applied Energy*, 142:361–388, March 2015. DOI: 10.1016/j.apenergy.2014.12.043. URL: <https://www.sciencedirect.com/science/article/abs/pii/S0306261914013002>.
- [144] Qazi Shahzad Ali and Man-Hoe Kim. Design and performance analysis of an airborne wind turbine for high-altitude energy harvesting. *Energy*, 230:120829, September 2021. DOI: 10.1016/j.energy.2021.120829. URL: <https://www.sciencedirect.com/science/article/abs/pii/S036054422101077X>.
- [145] Andras Sobester and Tom Barrett. Quest for a truly parsimonious airfoil parameterization scheme. *The 26th Congress of ICAS and 8th AIAA ATIO*, June 2008. DOI: 10.2514/6.2008-8879.
- [146] Vis Sripawadkul, Mattia Padulo, and Marin Guenov. A Comparison of Airfoil Shape Parameterization Techniques for Early Design Optimization. en. In *13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, Fort Worth, Texas. American Institute of Aeronautics and Astronautics, September 2010. ISBN: 9781600869549. DOI: 10.2514/6.2010-9050. URL: <https://arc.aiaa.org/doi/10.2514/6.2010-9050> (visited on 01/12/2022).
- [147] Dominic A. Masters, Nigel J. Taylor, T. Rendall, Christian B. Allen, and Daniel J. Poole. Review of Aerofoil Parameterisation Methods for Aerodynamic Shape Optimisation. en. In *53rd AIAA Aerospace Sciences Meeting*, Kissimmee, Florida. American Institute of Aeronautics and Astronautics, January 2015. ISBN: 9781624103438. DOI: 10.2514/6.2015-0761. URL: <https://arc.aiaa.org/doi/10.2514/6.2015-0761> (visited on 01/10/2022).
- [148] Haris Moazam Sheikh, Zeeshan Shabbir, Hassan Ahmed, Muhammad Hamza Waseem, and Muhammad Zubair Sheikh. Computational fluid dynamics analysis of a modified savonius rotor and optimization using response surface methodology. *Wind Engineering*, 41(5):285–296, 2017.
- [149] Antony Jameson. Aerodynamic design via control theory. *Journal of Scientific Computing*, 3(3):233–260, September 1988. DOI: 10.1007/bf01061285. URL: <https://link.springer.com/article/10.1007/BF01061285>.
- [150] Jamshid A. Samareh. Survey of shape parameterization techniques for high-fidelity multidisciplinary shape optimization. *AIAA Journal*, 39(5):877–884, May 2001. DOI: 10.2514/2.1391. URL: <https://ui.adsabs.harvard.edu/abs/2001AIAAJ..39..877S/abstract>.

- [151] Helmut Sobieczky. Parametric Airfoils and Wings. In Ernst Heinrich Hirschel, Kozo Fujii, Werner Haase, Bram van Leer, Michael A. Leschziner, Maurizio Pandolfi, Arthur Rizzi, Bernard Roux, Kozo Fujii, and George S. Dulikravich, editors, *Recent Development of Aerodynamic Design Methodologies*. Volume 65, pages 71–87. Vieweg+Teubner Verlag, Wiesbaden, 1999. ISBN: 9783322899545 9783322899521. DOI: 10.1007/978-3-322-89952-1_4. URL: http://link.springer.com/10.1007/978-3-322-89952-1_4 (visited on 01/12/2022).
- [152] Gerald Farin. Chapter 4 - the bernstein form of a bézier curve. In Gerald Farin, editor, *Curves and Surfaces for Computer-Aided Geometric Design (Third Edition)*, pages 41–63. Academic Press, Boston, third edition edition, 1993. ISBN: 978-0-12-249052-1. DOI: <https://doi.org/10.1016/B978-0-12-249052-1.50009-X>. URL: <https://www.sciencedirect.com/science/article/pii/B978012249052150009X>.
- [153] T. Rogalsky and R W Derksen. Bézier–parsec parameterization for airfoil optimization. *Canadian Aeronautics and Space Journal*, 55(3):163–174, 2009. DOI: 10.5589/q10-002. eprint: <https://doi.org/10.5589/q10-002>. URL: <https://doi.org/10.5589/q10-002>.
- [154] Sepehr Sanaye and Arash Hassanzadeh. Multi-objective optimization of airfoil shape for efficiency improvement and noise reduction in small wind turbines. en. *Journal of Renewable and Sustainable Energy*, 6(5):053105, September 2014. ISSN: 1941-7012. DOI: 10.1063/1.4895528. URL: <http://aip.scitation.org/doi/10.1063/1.4895528> (visited on 01/12/2022).
- [155] Xiacong Han and David W. Zingg. An adaptive geometry parametrization for aerodynamic shape optimization. en. *Optimization and Engineering*, 15(1):69–91, March 2014. ISSN: 1389-4420, 1573-2924. DOI: 10.1007/s11081-013-9213-y. URL: <http://link.springer.com/10.1007/s11081-013-9213-y> (visited on 01/12/2022).
- [156] Uwe Schramm, Walter D. Pilkey, Richard I. DeVries, and Mark P. Zebrowski. Shape design for thin-walled beam cross sections using rational b splines. *AIAA Journal*, 33(11):2205–2211, November 1995. DOI: 10.2514/3.12870.
- [157] Thomas W. Sederberg and Scott R. Parry. Free-form deformation of solid geometric models. *ACM SIGGRAPH Computer Graphics*, 20(4):151–160, August 1986. DOI: 10.1145/15886.15903. URL: <https://dl.acm.org/doi/epdf/10.1145/15886.15903>.
- [158] H.J. Lamousin and N.N. Waggenspack. Nurbs-based free-form deformations. *IEEE Computer Graphics and Applications*, 14(6):59–65, November 1994. DOI: 10.1109/38.329096. URL: <https://ieeexplore.ieee.org/document/329096>.
- [159] Holger Wendland. Scattered data approximation, December 2004. DOI: 10.1017/cbo9780511617539. URL: <https://www.cambridge.org/core/books/scattered-data-approximation/980EEC9DBC4CAA711D089187818135E3>.

- [160] Martin D. Buhmann. Radial basis functions, July 2003. DOI: 10.1017/cbo9780511543241. URL: <https://www.cambridge.org/core/books/radial-basis-functions/27D6586C6C128EABD473FDC08B07BD6D>.
- [161] David J. J. Toal, Neil W. Bressloff, Andy J. Keane, and Carren M. E. Holden. Geometric filtration using proper orthogonal decomposition for aerodynamic design optimization. *AIAA Journal*, 48(5):916–928, May 2010. DOI: 10.2514/1.41420.
- [162] Satyajit Ghoman, Zhicun Wang, Ping Chen, and Rakesh Kapania. A pod-based reduced order design scheme for shape optimization of air vehicles. *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference; BR; 20th AIAA/ASME/AHS Adaptive Structures Conference; BR; 14th AIAA*, April 2012. DOI: 10.2514/6.2012-1808.
- [163] Raymond M. Hicks and Preston A. Henne. Wing Design by Numerical Optimization. en. *Journal of Aircraft*, 15(7):407–412, July 1978. ISSN: 0021-8669, 1533-3868. DOI: 10.2514/3.58379. URL: <https://arc.aiaa.org/doi/10.2514/3.58379> (visited on 01/12/2022).
- [164] Brenda Kulfan and John Bussolletti. "Fundamental" Parametric Geometry Representations for Aircraft Component Shapes. en. In *11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Portsmouth, Virginia. American Institute of Aeronautics and Astronautics, September 2006. ISBN: 9781624100208. DOI: 10.2514/6.2006-6948. URL: <http://arc.aiaa.org/doi/10.2514/6.2006-6948> (visited on 01/12/2022).
- [165] Md Tausif Akram and Man-Hoe Kim. CFD Analysis and Shape Optimization of Airfoils Using Class Shape Transformation and Genetic Algorithm—Part I. en. *Applied Sciences*, 11(9):3791, April 2021. ISSN: 2076-3417. DOI: 10.3390/app11093791. URL: <https://www.mdpi.com/2076-3417/11/9/3791> (visited on 01/12/2022).
- [166] Asha Viswanath, A. I. J. Forrester, and A. J. Keane. Dimension reduction for aerodynamic design optimization. *AIAA Journal*, 49(6):1256–1266, June 2011. DOI: 10.2514/1.j050717.
- [167] Asha Viswanath, A. I. J. Forrester, and A. J. Keane. Constrained design optimization using generative topographic mapping. *AIAA Journal*, 52(5):1010–1023, May 2014. DOI: 10.2514/1.j052414.
- [168] Davide Cinquegrana and Emiliano Iuliano. Investigation of adaptive design variables bounds in dimensionality reduction for aerodynamic shape optimization. *Computers & Fluids*, 174:89–109, September 2018. DOI: 10.1016/j.compfluid.2018.07.012. URL: <https://www.sciencedirect.com/science/article/abs/pii/S0045793018304031>.
- [169] Wei Chen, Kevin Chiu, and Mark D. Fuge. Airfoil design parameterization and optimization using bézier generative adversarial networks. *AIAA Journal*, 58(11):4723–4735, November 2020. DOI: 10.2514/1.j059317.

- [170] Sahuck Oh, Chung-Hsiang Jiang, Chiyu Jiang, and Philip S. Marcus. Finding the optimal shape of the leading-and-trailing car of a high-speed train using design-by-morphing. *Computational Mechanics*, 62(1):23–45, July 2018. DOI: 10.1007/s00466-017-1482-4.
- [171] Haris Moazam Sheikh and Philip S. Marcus. Vertical Axis Wind Turbine Design Using Design-by-Morphing and Bayesian Optimization. In *APS Division of Fluid Dynamics Meeting Abstracts*, APS Meeting Abstracts, Q14.007, Q14.007, November 2019.
- [172] Haris Moazam Sheikh, Tess A Callan, Kealan J Hennessy, and Philip S Marcus. Optimization of the shape of a hydrokinetic turbine’s draft tube and hub assembly using design-by-morphing with bayesian optimization. *arXiv.org*, 2022. DOI: 10.48550/arXiv.2207.11451. URL: <https://arxiv.org/abs/2207.11451>.
- [173] Suleyman Murat Koroglu and Ibrahim Ozkol. Optimization of an airfoil characteristics to minimize the turn radius of a small unmanned aerial vehicle. *2019 IEEE 10th International Conference on Mechanical and Aerospace Engineering (ICMAE)*, July 2019. DOI: 10.1109/icmae.2019.8880954. URL: <https://ieeexplore.ieee.org/document/8880954>.
- [174] Michael S. Selig. UIUC airfoil data site, Department of Aeronautical and Astronautical Engineering University of Illinois at Urbana-Champaign, 1996. URL: <https://m-selig.ae.illinois.edu/ads.html>. Last access: Feb 2022.
- [175] Michael G. Piotrowski and David W. Zingg. Investigation of a smooth local correlation-based transition model in a discrete-adjoint aerodynamic shape optimization algorithm. *AIAA SCITECH 2022 Forum*, January 2022. DOI: 10.2514/6.2022-1865.
- [176] Xiaolong He, Jichao Li, Charles A. Mader, Anil Yildirim, and Joaquim R.R.A. Martins. Robust aerodynamic shape optimization—from a circle to an airfoil. *Aerospace Science and Technology*, 87:48–61, April 2019. DOI: 10.1016/j.ast.2019.01.051.
- [177] Gaetan K. W. Kenway and Joaquim R. R. A. Martins. Multipoint aerodynamic shape optimization investigations of the common research model wing. *AIAA Journal*, 54(1):113–128, January 2016. DOI: 10.2514/1.j054154.
- [178] Mark Drela. Xfoil: an analysis and design system for low reynolds number airfoils. *Lecture Notes in Engineering*:1–12, 1989. DOI: 10.1007/978-3-642-84010-4_1. URL: https://link.springer.com/chapter/10.1007/978-3-642-84010-4_1.
- [179] Göran Ronsten. Static pressure measurements on a rotating and a non-rotating 2.375 m wind turbine blade. comparison with 2d calculations. *Journal of Wind Engineering and Industrial Aerodynamics*, 39(1-3):105–118, January 1992. DOI: 10.1016/0167-6105(92)90537-k. URL: <https://www.sciencedirect.com/science/article/pii/016761059290537K?via%3Dihub>.

- [180] P. Gigue'ere and M. S. Selig. New airfoils for small horizontal axis wind turbines. *Journal of Solar Energy Engineering*, 120(2):108–114, May 1998. DOI: 10.1115/1.2888052. URL: <https://asmedigitalcollection.asme.org/solarenergyengineering/article-abstract/120/2/108/438974/New-Airfoils-for-Small-Horizontal-Axis-Wind?redirectedFrom=fulltext>.
- [181] Brian R. Jones, William A. Crossley, and Anastasios S. Lyrintzis. Aerodynamic and aeroacoustic optimization of rotorcraft airfoils via a parallel genetic algorithm. *Journal of Aircraft*, 37(6):1088–1096, November 2000. DOI: 10.2514/2.2717.
- [182] Thomas J. Mueller and James D. DeLaurier. Aerodynamics of small vehicles. *Annual Review of Fluid Mechanics*, 35(1):89–111, January 2003. DOI: 10.1146/annurev.fluid.35.101101.161102.
- [183] Forrester T. Johnson, Edward N. Tinoco, and N. Jong Yu. Thirty years of development and application of cfd at boeing commercial airplanes, seattle. *Computers & Fluids*, 34(10):1115–1151, December 2005. DOI: 10.1016/j.compfluid.2004.06.005. URL: <https://www.sciencedirect.com/science/article/pii/S0045793005000125?via%5C%3Dihub>.
- [184] W.M.J. Batten, A.S. Bahaj, A.F. Molland, and J.R. Chaplin. Hydrodynamics of marine current turbines. *Renewable Energy*, 31(2):249–256, February 2006. DOI: 10.1016/j.renene.2005.08.020. URL: <https://www.sciencedirect.com/science/article/pii/S0960148105002314?via%3Dihub>.
- [185] Cody Lafountain, Kelly Cohen, and Shaaban Abdallah. Use of xfoil in design of camber-controlled morphing uavs. *Computer Applications in Engineering Education*, 20(4):673–680, April 2010. DOI: 10.1002/cae.20437. URL: <https://onlinelibrary.wiley.com/doi/10.1002/cae.20437>.
- [186] Giridhar Ramanujam and Huseyin Ozdemir. Improving airfoil lift prediction. *35th Wind Energy Symposium*, January 2017. DOI: 10.2514/6.2017-1999.
- [187] Betty S. Walker Robert J. McGhee. Experimental results for the eppler 387 airfoil at low reynolds numbers in the langley low-turbulence pressure tunnel, 1988.
- [188] Michael S Selig. *Summary of low speed airfoil data. Vol. 1*. Virginia Beach, Va. Soartech Publications, 1995. ISBN: 9780964674714. URL: <https://experts.illinois.edu/en/publications/summary-of-low-speed-airfoil-data-vol-1>.
- [189] J. Morgado, R. Vizinho, M.A.R. Silvestre, and J.C. Páscoa. Xfoil vs cfd performance predictions for high lift low reynolds number airfoils. *Aerospace Science and Technology*, 52:207–214, May 2016. DOI: 10.1016/j.ast.2016.02.031. URL: <https://www.sciencedirect.com/science/article/pii/S1270963816300839?via%3Dihub>.
- [190] Miettinen Kaisa. *Nonlinear Multiobjective Optimization*, volume 12 of *International Series in Operations Research & Management Science*. Kluwer Academic Publishers, Boston, USA, 1999.

- [191] Nyoman Gunantara. A review of multi-objective optimization: methods and its applications. *Cogent Engineering*, 5(1):1502242, 2018. Qingsong Ai, editor. DOI: 10.1080/23311916.2018.1502242.
- [192] Barrett B. E. Barron F. H. Decision quality using ranked attribute weights. *Management Science*, 42:1515, 1996.
- [193] Dennis J. E. Das I. A closer look at drawbacks of minimizing weighted sums of objectives for pareto set generation in multi-criteria optimization problems. *Structural Optimization*, 14:63, 1997.
- [194] Chang C. T. Multi-choice goal programming. *Omega*, 35:389, 2007.
- [195] Ratiba F. Ghachi, Wael I. Alnahhal, Osama Abdeljaber, Jamil Renno, A. B. M. Tahidul Haque, Jongmin Shim, and Amjad Aref. Optimization of viscoelastic metamaterials for vibration attenuation properties. *International Journal of Applied Mechanics*, 12(10):2050116, 2020. DOI: 10.1142/S1758825120501161. eprint: <https://doi.org/10.1142/S1758825120501161>. URL: <https://doi.org/10.1142/S1758825120501161>.
- [196] Zacharias Vangelatos, Haris Moazam Sheikh, Philip S. Marcus, Costas P. Grigoropoulos, Victor Z. Lopez, George Flamourakis, and Maria Farsari. Strength through defects: a novel bayesian approach for the optimization of architected materials. *Science Advances*, 7(41):eabk2218, 2021. DOI: 10.1126/sciadv.abk2218. eprint: <https://www.science.org/doi/pdf/10.1126/sciadv.abk2218>. URL: <https://www.science.org/doi/abs/10.1126/sciadv.abk2218>.
- [197] Tomasz Schlieter and A. Długosz. Structural optimization of aerofoils for many criteria. In pages 448–451, January 2020. DOI: 10.21495/5896-3-448.
- [198] Adam Chehouri, Rafic Younes, Adrian Ilinca, and Jean Perron. Wind turbine design: multi-objective optimization. *Wind Turbines - Design, Control and Applications*, July 2016. DOI: 10.5772/63481. URL: <http://dx.doi.org/10.5772/63481>.
- [199] S. Rodrigues, P. Bauer, and Peter A.N. Bosman. Multi-objective optimization of wind farm layouts – complexity, constraint handling and scalability. *Renewable and Sustainable Energy Reviews*, 65:587–609, 2016. ISSN: 1364-0321. DOI: <https://doi.org/10.1016/j.rser.2016.07.021>. URL: <https://www.sciencedirect.com/science/article/pii/S1364032116303458>.
- [200] P Nguyen. A review of hybrid/combined methods for trajectory optimization of flight vehicles. *Journal of Physics: Conference Series*, 1958(1):012032, June 2021. DOI: 10.1088/1742-6596/1958/1/012032. URL: <https://doi.org/10.1088/1742-6596/1958/1/012032>.
- [201] Huang Gao, Yun Zhang, Xundao Zhou, and Dequn Li. Intelligent methods for the process parameter determination of plastic injection molding. *Frontiers of Mechanical Engineering*, 13(1):85, 85, 2018. DOI: 10.1007/s11465-018-0491-0.

- [202] Victor Franco Correia, José S. Moita, Filipa Moleiro, and Cristóvão M. Mota Soares. Optimization of metal–ceramic functionally graded plates using the simulated annealing algorithm. *Applied Sciences*, 11(2), 2021. ISSN: 2076-3417. DOI: 10.3390/app11020729. URL: <https://www.mdpi.com/2076-3417/11/2/729>.
- [203] Alan D. Fox, David W. Corne, C. Gabriela Mayorga Adame, Jeff A. Polton, Lea-Anne Henry, and J. Murray Roberts. An efficient multi-objective optimization method for use in the design of marine protected area networks. *Frontiers in Marine Science*, 6:17, 2019. ISSN: 2296-7745. DOI: 10.3389/fmars.2019.00017. URL: <https://www.frontiersin.org/article/10.3389/fmars.2019.00017>.
- [204] Hamid Afshari, Warren Hare, and Solomon Tesfamariam. Constrained multi-objective optimization algorithms: review and comparison with application in reinforced concrete structures. *Applied Soft Computing*, 83:105631, July 2019. DOI: 10.1016/j.asoc.2019.105631.
- [205] Di Wang, Ling Geng, Yu-Jun Zhao, Yang Yang, Yan Huang, Yang Zhang, and Hong-Bin Shen. Artificial intelligence-based multi-objective optimization protocol for protein structure refinement. *Bioinformatics*, 36(2):437–448, July 2019. ISSN: 1367-4803. DOI: 10.1093/bioinformatics/btz544. eprint: <https://academic.oup.com/bioinformatics/article-pdf/36/2/437/36759409/btz544.pdf>. URL: <https://doi.org/10.1093/bioinformatics/btz544>.
- [206] Ying Gao, Lei Shi, and Pingjing Yao. Study on multi-objective genetic algorithm. In *Proceedings of the 3rd World Congress on Intelligent Control and Automation (Cat. No.00EX393)*, volume 1, 646–650 vol.1, 2000. DOI: 10.1109/WCICA.2000.860052.
- [207] S.N. Skinner and H. Zare-Behtash. State-of-the-art in aerodynamic shape optimisation methods. en. *Applied Soft Computing*, 62:933–962, January 2018. ISSN: 15684946. DOI: 10.1016/j.asoc.2017.09.030. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1568494617305690> (visited on 01/12/2022).
- [208] Y. Rahmad, M. D. Robani, P. S. Palar, and L. R. Zuhail. Single- and multi-objective optimization of a low-speed airfoil using genetic algorithm. In page 020005, Tangerang Selatan, Indonesia, 2020. DOI: 10.1063/5.0002610. URL: <http://aip.scitation.org/doi/abs/10.1063/5.0002610> (visited on 01/12/2022).
- [209] Ke Zhao, Zhenghong Gao, and Jiang-tao Huang. Robust design of natural laminar flow supercritical airfoil by multi-objective evolution method. *Applied Mathematics and Mechanics*, 35:191–202, 2014.
- [210] Gavin K Ananda and Michael S Selig. Design of bird-like airfoils. In *2018 AIAA Aerospace Sciences Meeting*, page 0310, 2018.

- [211] Xingxing Li, Ke Yang, Jingyan Bai, and Jianzhong Xu. A method to evaluate the overall performance of the CAS-W1 airfoils for wind turbines. en. *Journal of Renewable and Sustainable Energy*, 5(6):063118, November 2013. ISSN: 1941-7012. DOI: 10.1063/1.4841056. URL: <http://aip.scitation.org/doi/10.1063/1.4841056> (visited on 01/31/2022).
- [212] Airfoil database search, 2022. URL: <http://airfoiltools.com/>. Last access: May 2022.
- [213] U.S. energy information administration monthly energy review, december 2021. <https://www.eia.gov/totalenergy/data/monthly/archive/00352112.pdf>, Report, 2022.
- [214] Taylor Curtis, Aaron Levine, and Kurt Johnson. State models to incentivize and streamline small hydropower development. In National Renewable Energy Laboratory (NREL), U.S. Department of Energy, October 2017.
- [215] United States Department of Energy. Hydropower Vision, 2001. Accessed 03-29-2022.
- [216] Tor Haakon Bakken, Håkon Sundt, Audun Ruud, and Atle Harby. Development of small versus large hydropower in norway– comparison of environmental impacts. *Energy Procedia*, 20:185–199, 2012. ISSN: 1876-6102. DOI: <https://doi.org/10.1016/j.egypro.2012.03.019>. URL: <https://www.sciencedirect.com/science/article/pii/S1876610212007497>. Technoport 2012 - Sharing Possibilities and 2nd Renewable Energy Research Conference (RERC2012).
- [217] Thiago Couto BA and Julian D. Olden. Global proliferation of small hydropower plants - science and policy. In *Front Ecol Environ*, volume 16, pages 91–100, 2018. DOI: <https://doi.org/10.1002/fee.1746>.
- [218] Douglas Hall G., Kelly S. Reeves, Julie Brizzee, Randy Lee D., Gregory Carroll R., and Garold Sommers L. Feasibility assessment of the water energy resources of the united states for new low power and small hydro classes of hydroelectric plants. In Idaho National Laboratory, 2006.
- [219] Martin Kaltschmitt, Wolfgang Streicher, and Andreas Wiese. *Renewable Energy: Technology, Environment and Economics*. Springer, Berlin, Heidelberg, 2007.
- [220] R. Mulley. *Flow of Industrial Fluids - Theory and Equations*. Taylor & Francis Group, LLC, 2004.
- [221] C. C. Warnick, Howard A. Mayo, James L. Carson, and Lee H. Sheldon. *Hydropower Engineering*. Prentice-Hall, Inc., 1984.
- [222] Lee H. Sheldon. Reviewing the approaches to hydro optimization. *Hydro Review*, 17:60–67, 1998.
- [223] Michel J. Cervantes. Effects of boundary conditions and unsteadiness on draft tube flow, ph.d. thesis. *Division of Fluid Mechanics, Luleå University of Technology*, 2003.
- [224] G. Krivchenko. *Hydraulic Machines: Turbines and Pumps*. CRC Press, Inc., 1994.

- [225] U. Andersson. An experimental study of the flow in a sharp-heel kaplan draft tube, ph.d. thesis. *Division of Fluid Mechanics, Luleå University of Technology*, 2009.
- [226] Håkan Nilsson. 3d numerical analysis of the unsteady turbulent swirling flow in a conical diffuser using fluent and openfoam. In *IAHR 2006*, 2006.
- [227] B. D. Marjavaara. Cfd driven optimization of hydraulic turbine draft tubes using surrogate models. *Division of Fluid Mechanics, Luleå University of Technology*, 2006.
- [228] Ole Gunnar Dahlhaug. A study of swirl flow in draft tubes. In Norwegian University of Science and Technology Trondheim, December 1997.
- [229] M.F. Gubin. *Draft Tubes of Hydro-electric Stations*. Energiya Press, Moscow, 1970.
- [230] Hölleforsen - vattenfall ab. <https://powerplants.vattenfall.com/holleforsen/>, 2022. Accessed: 2022-03-28.
- [231] N. Dahlbäck. Redesign of sharp heel draft tube - results from tests in model and prototype. In E. Cabrera, V. Espert, and F. Martínez, editors, *Proceedings of the XVIII IAHR Symposium on Hydraulic Machinery and Cavitation*, pages 985–993. Springer-Science+Business Media, B.V., 1996.
- [232] B.D. Marjavaara and T.S. Lundström. Automatic shape optimisation of a hydropower draft tube. In E. Cabrera, V. Espert, and F. Martínez, editors, *ASME/JSME 4th Joint Fluids Summer Engineering Conference*, volume 2, pages 1819–1824, 2003.
- [233] S. J. Daniels, A. A. M. Rahat, G. R. Tabor, J. E. Fieldsend, and R. M. Everson. Application of multi-objective bayesian shape optimisation to a sharp-heeled kaplan draft tube. *Optimization and Engineering*, 22, 4, 2021.
- [234] Steven J. Daniels, Alma A. M. Rahat, Richard M. Everson, Gavin R. Tabor, and Jonathan E. Fieldsend. A suite of computationally expensive shape optimisation problems using computational fluid dynamics. In Anne Auger, Carlos M. Fonseca, Nuno Lourenço, Penousal Machado, Luís Paquete, and Darrell Whitley, editors, *Parallel Problem Solving from Nature – PPSN XV*, pages 296–307, Cham. Springer International Publishing, 2018. ISBN: 978-3-319-99259-4.
- [235] R. Eisinger and A. Ruprecht. Automatic shape optimisation of hydro turbine components based on cfd. *TASK quarterly: Scietific Bulletin of Academic Computing Centre Gdańsk*, 6:101–111, 2002.
- [236] J. McNabb, C. Devals, S. Kyriacou, N. Murry, and B. Mullins. Cfd based draft-tube hydraulic design optimisation. In *27th IAHR symposium on hydraulic machinery and systems*, 2014.
- [237] Matthias Schramm, Bernhard Stoevesandt, and Joachim Peinke. Optimization of airfoils using the adjoint approach and the influence of adjoint turbulent viscosity. *Computation*, 6(1), 2018. ISSN: 2079-3197. DOI: 10.3390/computation6010005. URL: <https://www.mdpi.com/2079-3197/6/1/5>.

- [238] Jiaqin Chen, Vadim Shapiro, Krishnan Suresh, and Igor Tsukanov. Shape optimization with topological changes and parametric control. *International Journal for Numerical Methods in Engineering*, 71(3):313–346, July 2007. DOI: 10.1002/nme.1943.
- [239] Wei-Hong Zhang, Pierre Beckers, and Claude Fleury. A unified parametric design approach to structural shape optimization. *International Journal for Numerical Methods in Engineering*, 38(13):2283–2292, July 1995. DOI: 10.1002/nme.1620381309.
- [240] T. J. R. Hughes, J. A. Cottrell, and Y. Bazilevs. Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement. *Computer Methods in Applied Mechanics and Engineering*, 194(39-41):4135–4195, October 2005. DOI: 10.1016/j.cma.2004.10.008.
- [241] Wei Shyy, Nilay Papila, Rajkumar Vaidyanathan, and Kevin Tucker. Global design optimization for aerodynamics and rocket propulsion components. *Progress in Aerospace Sciences*, 37(1):59–118, January 2001. DOI: 10.1016/S0376-0421(01)00002-1.
- [242] X. D. Wang, C. Hirsch, Sh. Kang, and C. Lacor. Multi-objective optimization of turbomachinery using improved NSGA-II and approximation model. *Computer Methods in Applied Mechanics and Engineering*, 200(9-12):883–895, February 2011. DOI: 10.1016/j.cma.2010.11.014.
- [243] Xiaowei LI Liang FANG. Design optimization of unsteady airfoils with continuous adjoint method. *Applied Mathematics and Mechanics*, 36(10):1329, 1329, 2015. DOI: 10.1007/s10483-015-2010-9. URL: https://www.amm.shu.edu.cn/CN/abstract/article_16192.shtml.
- [244] M.J. Cervantes, T.F. Engström, and L.H. Gustavsson. In *Proceedings of the third IAHR/ERCOFTAC Workshop on draft tube flows: Turbine-99 III*. Division of Fluid Mechanics, Luleå University of Technology, 2005.
- [245] B. G. Mulu. An experimental and numerical investigation of a kaplan turbine model, ph.d. thesis. *Division of Fluid and Experimental Mechanics, Luleå University of Technology*, 2012.
- [246] Håkan Nilsson, Sebastian Muntean, and Romeo F. Susan-Resiga. Evaluation of open-foam for cfd of turbulent flow in water turbines. In *IAHR 2009*, 2009.
- [247] B. R. Gebart, L.H. Gustavsson, and R.I. Karlsson. In *Turbine-99: Workshop on Draft Tube Flow*. Division of Fluid Mechanics, Luleå University of Technology, 2000.
- [248] T.F. Engström, L.H. Gustavsson, and R.I. Karlsson. In *Proceedings of Turbine-99 - Workshop 2: The second ERCOFTAC Workshop on Draft Tube Flow*, Älvkarleby, Sweden. Division of Fluid Mechanics, Luleå University of Technology, 2001.
- [249] B.G. Mulu, P.P. Jonsson, and M.J. Cervantes. Experimental investigation of a kaplan draft tube – part i: best efficiency point. *Division of Fluid Mechanics, Luleå University of Technology*, 2012.
- [250] C. Geuzaine and J.-F. Remacle. Gmsh, version 4.6.0, 2022. URL: <http://gmsh.info/>.

- [251] Yulin Wu, Shuhong Liu, Hua-Shu Dou, Shangfeng Wu, and Tiejun Chen. Numerical prediction and similarity study of pressure fluctuation in a prototype kaplan turbine and the model turbine. *Computers & Fluids*, 56:128–142, March 2012. DOI: 10.1016/j.compfluid.2011.12.005.

Finally, a few last words...

Nitwit! Blubber! Oddment! Tweak!