2026 IEEE/SICE International Symposium on System Integration (SII)
January 11-14, 2026. Cancún, México

Co-Optimization of Design and Manufacturing Parameters for Low-Cost Robotic Actuation

Gregory M. Campbell, Yi Cao, Hannah Escritor, Zihao Zhou, and Mark Yim

***Abstract*—Additive and low-cost manufacturing techniques promise increased access to robotic actuation at the cost of mechanical precision. In this work, we employ principled Design of Experiments (DoE), including Taguchi orthogonal arrays, in parallel to sequential experimentation enabled by Bayesian Optimization (BO) for co-optimization of design and manufacturing parameters across two design case studies. We optimize for a combination of gear ratio and backdrivability in a 3D-printed compound Wolfrom bilateral gearbox. We also optimize for crack pressure and steady-state pressure differential of an injection-molded silicone check valve. Using BO, we find a 3D-printing compatible gear design with a gear ratio of 63.6 that backdrives without ever needing more than 0.35 Nm of input torque. This represents a 49% increase in 'score' over the Taguchi method. Similarly, we find a BO valve with lower combined crack and steady-state pressure errors than the Taguchi trials, decreasing cumulative error by 55%. Tracking model uncertainty throughout training, we conclude that further model training is necessary to reach optimal results in both cases. We further conclude that BO via the Ax platform is not yet as "plug-and-play" as Taguchi arrays.**

Keywords: hardware design, machine learning, Wolfrom gear, Bayesian Optimization, Taguchi orthogonal array

I. INTRODUCTION

Rapid and low-cost manufactured parts are both accessible and time-efficient for product development. However, they suffer from a variety of inaccuracies compared to carefully machined counterparts. While theoretical models give insight into the mechanism and broad effects of design parameters, manufacturing imprecisions and ductile materials negatively affect theoretical precision [1]. Prototypes and experimentation therefore become instrumental for optimization. We investigate efficient experimentation to both relieve associated costs and improve performance of these parts.

Researchers have employed a principled Design of Experiments (DoE) to quantify parametric design spaces - adding efficiency when capturing data is difficult or expensive. These approaches aim to span the design space with few experiments, sometimes leveraging statistics [2]. Similar methods are used in Robust Parameter Design (RPD) [3], which is more robust to noise factors that are outside the researcher's control. One RPD tool, Taguchi orthogonal arrays, is easily applied for varying numbers of parameters and levels, and notably discretizes all continuous parameters, allowing for the mixing of continuous, integer, and categorical elements.

Parallel to DoE, the growth of machine learning (ML) has allowed for precise surrogate models to approximate the effects of design choices in even the most non-linear of interactions [4], [5]. These surrogate models allow for interpolation between sampled design combinations and fast speeds at inference time. ML has been applied to data collected through DoE primarily to increase prediction accuracy, with many common ML algorithms using neural network (NN) [6].

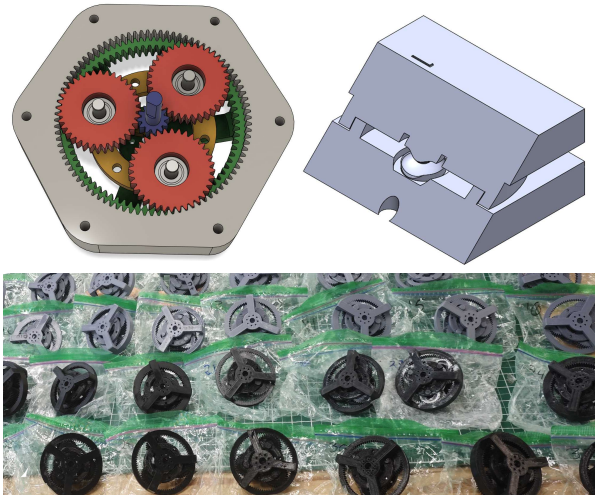

Fig. 1. (left) Case study 1: 3D-printed Wolfrom gear. CAD includes design parameters in planet gears (red), sun gear (blue), stage 2 ring (green), and carriers (gold). (right) Case study 2: injection-molded soft valve. CAD of valve mold includes design parameters in domed valve and thickness. (bottom) Assembled PLA (grey) and ABS (black) 3D-printed gears after sequential experimentation. Gear outer diameter of 84 mm.

Sequential experimentation (adaptive sampling) and active learning provide means of collecting new data to minimize overall error of a ML model [7]. Beyond model improvement, Bayesian Optimization (BO) uses targeted acquisition functions with a Gaussian Process Regression surrogate to maximize the likelihood of improving a specific output metric during each experiment. BO has become widely accessible via open-source software [8], notably BoTorch [9]. BO approaches substantially reduce the number of experiments required for optimization [10], including within the areas of mechanical design of rapid-prototyped elements [11] and manufacturing hyperparameters using continuous parameterizations [12]. Using BoTorch, Meta's Adaptive Experimentation Platform (Ax) has provided an accessible platform for BO [13] that inherently combines continuous elements with categorical parameters. Ax's efficacy in mechanical design applications has not been thoroughly explored.

The goal of this work is to investigate the potential for improved functionality in low-cost robotic systems using

978-1-6654-5784-2/26/$31.00 ©2026 IEEE

mechanical optimization. Specifically, whether BO can offer improved outputs or increase experimental efficiency compared to DoE. Such a pipeline for optimizing the functionality of accessible or on-hand hardware represents a step away from high-precision and high-cost requirements in robotic actuation. We conduct two case studies relevant to this topic, one for a classical, geared drivetrain and one for a soft, pneumatic valve.

In our first case study, we co-design categorical, integer, and continuous design and manufacturing parameters to maximize a single objective combination of gear ratio and backdrivability for a compound Wolfrom bilateral gearbox [14]. The property of high gear ratio is crucial for getting larger torques from DC motors that run at high speeds but low torques. This, in tandem with bidirectional force transmission through backdrivability, allows for relatively inexpensive robotic systems that are relevant in haptics and exoskeletons [15]. However, there is a natural tradeoff between gear ratio and backdrivability, as higher gear ratio negatively influences reflected inertia, damping and friction [16]. In this work, we develop an accessible and low cost 3D-printing compatible gear based on [14]. This material choice presents tolerancing, stiffness, and surface roughness challenges which increase backdrivability limitations and are also difficult to model with standard gear design. Backdrivable gears with ratios up to 30:1 have been achieved with 3D-printing before [15], but in this work we focus on optimizing for higher gear ratios with quantified backdrivability.

In the second case study, we vary categorical and continuous design parameters to maximize multiple objectives relevant to the crack pressure and steady-state pressure difference across an injection-molded silicone check valve. Precise control over these metrics during the design of pneumatic systems can allow for the timely activation of different pneumatic actuators when a specific internal pressure is reached. Eventually, we aim to co-design these valves with the soft pneumatic actuators themselves to enable multi-actuator lift trajectories [4]. For this study, we again impose system-match specifications (tubing, pump, pressure chambers, release valve), limiting the effectiveness of existing theoretical models [17], [18]. Specifically, the use of a pneumatic pump couples the control of pressure and flow, while set system components require that all design changes be applied at the valve specifically, unlike in [17].

We apply surrogate model uncertainty quantification to analyze the effect of input data on model training progression for both case studies. We also apply parameter sensitivity analysis to highlight important features. We provide context for the relative advantages of BO and Taguchi methods in rapid prototyping for mechanical design, and highlight questions that are important to future applications.

II. Bayesian Optimization

This brief primer on BO offers only broad context. Consider [8] for a more rigorous introduction to the topic.

BO aims to maximize the output of a given function with minimal experimental data. The framework requires that we describe all of our parameterized design choices in a multi-dimensional state vector (x) and our output as the value (y). BO is a 'sequential' approach that updates the next design setup based on a comprehensive 'surrogate' model of all possibilities.

We use a Gaussian Process (GP) as the surrogate to predict all potential design scores based on previously collected data. The GP models' score for any location in the design space (x) is a mean prediction (μ) and an epistemic uncertainty (σ) of the output (y). A high μ represents a promising score while a high σ represents uncertainty. We retrain the model with each new data to reduce the error of these predictions for all known data.

Prior to each new experiment, we apply an acquisition function to this surrogate model to choose the next experimental design set (x_i). The acquisition function assumes a 'similarity' between x and y that is the basis of BO: that prior results inform future results. It thereby predicts design areas where the GP represents potential of a high output score. Differing choices of acquisition function will emphasize breadth of search ('exploring') or targeting of already high-scoring areas ('exploiting').

We iterate through the design space, changing design inputs (x) with each experiment to move closer to a maximized output (y) while monitoring the change in GP uncertainty.

III. Modeling & Parameter choice

A. Modeling

Gears: We wish to design a compound Wolfrom bilateral gearbox with high back-driving efficiency at high reduction ratio for low-cost versatile robot actuation. This compound gearbox is composed of two coaxial planetary gearboxes. The sun gear is present only on the forward-drive input side. We refer to the forward-drive input side of the compound gearbox as gear 1 and the output side as gear 2. The geometry of each individual gear in the compound gearbox is fully specified by its module M, tooth number Z, and profile shift coefficient X. Assigning fixed module values a priori, there are then a total of 10 free design parameters, namely:

$$X_{p1}, X_{p2}, X_{r1}, X_{r2}, X_s, Z_{p1}, Z_{p2}, Z_{r1}, Z_{r2}, Z_s,$$

where the subscript r denotes the ring gear, the subscript p denotes the planet gear, and the subscript s denotes the sun gear. However, these 10 parameters are not pairwise independent. The coaxial geometry design imposes geometry constraints that allow us to render 4 of the 10 parameters dependent on the remaining 6, as follows. The original 10 design parameters are subjected to the coaxial geometry constraint:

$$r_a = r_b = r_c \tag{1}$$

where r_a, r_b, r_c are respectively the center distances between s and p_1, p_1 and r_1, and between p_2 and r_2. Let subscript t denote the tip diameter and subscript r denote the root diameter, and we introduce a variable clearance C between the meshing gears. Then

$$r_a = r_{s_t} + r_{p_r} + C_{s \to p1} \tag{2}$$

$$r_b = r_{r1_r} - r_{p1_t} - C_{p1\to r1} \tag{3}$$

where $C_{s\to p1}$ denotes the clearance between sun and planet one and $C_{p1\to r1}$ denotes the clearance between planet one and ring one.

For profile shifted spur gears, radius is related to the profile shift coefficient [14], with $+1$ for tip and -1 for root:

$$r_i = \frac{m_i}{2}(Z_i + 2x_i \pm 1) \tag{4}$$

Substituting Eqn. 2 3 4 into Eqn. 1 and simplifying yields:

$$Z_{p1} + 2X_{p1} = \frac{1}{2}[(Z_{r1} + 2X_{r1}) - (Z_s + 2X_s)] - \frac{C_{s\to p1} + C_{p1\to r1}}{m1} \tag{5}$$

$$Z_{p2} + 2X_{p2} = \frac{m_1}{m_2}[(Z_{p1} + 2X_{p1}) - (Z_{r1} + 2X_{r1})] + (Z_{r2} + 2X_{r2}) + 2(C_{p1\to r1} - C_{p2\to r2}) \tag{6}$$

where m_1 is the module of the sun, ring one, and planet one gears. As shown in Eqn. 5 and 6, the coaxial constraint Eqn.1 renders $Z_{p1}, Z_{p2}, X_{p1}, X_{p2}$, dependent on the remaining 6 design parameters and the variable clearances. In addition, we fix the values of Z_{r1} and X_{r1} to ease experimentation, which leaves our free design parameters:

$$X_s, X_{r2}, Z_s, Z_{r2}$$

These four parameters are the primary design variables in our Taguchi array for the compound Wolfrom gearbox and fully determine the right-hand side Eqn 6. We can then solve for the values of the remaining independent variables $X_{p1}, X_{p2}, Z_{p1}, Z_{p2}$. The number of gear teeth must be an integer value. We choose X to range between -1 and 1 such that there are then only four valid combinations of Z_p and X_p. Among them, we define the standard solution to be $Z_{p_s} = floor(Y)$ and $X_{p_s} = \frac{1}{2}(Z_{p_s} - floor(Y))$ where Y denotes the right-hand side of Eqn 6. The remaining 3 valid solutions contain Z_p values that are integer offsets from the standard solution Z_{p_s}, quantified by offset parameters $p_{1\ offset}$ and $p_{2\ offset}$ that take the integer range $[-1, 2]$. We include these offset parameters as secondary design variables in our Taguchi array, along with the variable clearances. We choose to enforce all clearance values as equal and define a single clearance parameter Cl.

In addition to the above design parameters, we study the effects of manufacturing choices. We define the parameters material and lubricant with 2 and 4 levels respectively, along with the gear thickness parameter g_{t_h} that we assign integer values between 3 and 6 mm. These variables primarily affect friction between meshing gears, which is a critical factor influencing the driving efficiencies of the gearbox. We assume an average coulomb friction constant [19], which directly affects drive efficiency, as calculated in [14].

To bound our parameter space, we define the basic driving efficiency η_0 of a pair of involute gears as the ratio of power output to power input, and simplify the expressions using no-slip contact condition on the pitch circle.

$$\eta_0 = \frac{\omega_{out}\tau_{out}}{\omega_{in}\tau_{in}} = \frac{f_{out}}{f_{in}}; \tag{7}$$

where f is the tangential force, ω is the rotational velocity, and τ is the gear torque from the driving (in) and driven (out) gears.

From kinematics, we have the following gear ratio equation, which also helps define efficiency:

$$G_r = \frac{1 - I_2}{1 + I_1}; \tag{8}$$

where

$$I_1 = \frac{r_{r1}r_{p2}}{r_{r2}r_{p12}} = \frac{Z_{r1}}{Z_s}; \quad I_2 = \frac{r_{r1}r_{p11}}{r_s r_{p12}} = \frac{Z_{r1}Z_{p2}}{Z_{r2}Z_{p1}}; \tag{9}$$

r_{p11} and r_{p12} are respectively the pitch circles of planet one when engaged to sun and ring one.

In steady-state operation, the net torque acting on the gears is zero, allowing us to solve for all forces at the contact points as a function of radii, basic driving efficiencies, and the input torque. Reducing these equations, see [14], yields the back-drive efficiencies η' of the Wolfrom gearbox as follows:

$$\eta' = \frac{\tau_{\text{out}}\omega_s}{\tau_{\text{in}}\omega_{r2}} = \frac{(1 + I_1)\eta_a(\eta_b\eta_c - I_2)}{\eta_c(\eta_a\eta_b + I_1)(1 - I_2)} \tag{10}$$

for $I_2 < 1$, and

$$\eta' = \frac{(1 + I_1)\eta_a(1 - \eta_b\eta_c I_2)}{(\eta_a\eta_b I_1)(1 - I_2)} \tag{11}$$

for $I_2 > 1$. Where:

$$\eta_i = 1 - \mu\pi\left(\frac{1}{z_{i1}} + \text{sgn}\frac{1}{z_{i2}}\right)\varepsilon_i, \quad i \in \{a, b, c\} \tag{12}$$

and a, b, c are defined as:

$$a_1 = s,\ a_2 = p_1,\ b_1 = p_1,\ b_2 = r_1,\ c_1 = p_2,\ c_2 = r_2$$

The final Taguchi array uses 10 variables $Z_s, Z_{r2}, X_s, X_{r2}, p1_{offset}, p2_{offset}, g_{t_h}, Cl$, material, lubricant, with ranges shown in in Table I.

We use one-factor-at-a-time analysis of Eqn. 10 and 11 to predict driving efficiencies and inform the parameter ranges designated in Table I.

Valves: We wish to design a soft check valve that achieves a crack pressure of 10 kPa and a steady-state pressure of 5 kPa for the pneumatic capacitance and resistance of our specific two-chamber test system. These values are assigned to enforce a delayed activation of the second pneumatic chamber and a targeted offset in relative inflation trajectories of soft actuators.

The valve consists of a single silicone body with dimensions determined by injection molding. The functional properties of the valve depend on mechanical and geometric properties of the silicone body, and also on fluidic properties of the pneumatic system [18]. We hold constant the fluidic properties of the pump, piping, pneumatic capacitance, and pneumatic resistance.

The mechanics of the dome are modeled as the superposition of a linear spring with stiffness k_s and a rotational spring k_R. The apex of the dome, node P, is defined by the position x, which is considered 1-dimensionally and radially towards the center of the dome. The equilibrium position of x is determined from a force balance between the pneumatic pressure (Δp) applied over the area of the dome (A_V) and the mechanical spring force F_s, which is defined as the derivative of the silicone dome's elastic energy, E, with respect to x [18]:

$$F_S = \frac{dE}{dx} = -k_S \left(\frac{\sqrt{h_{dome}^2 + b_0^2} - \sqrt{x^2 + b_0^2}}{\sqrt{x^2 + b_0^2}} \right) x + \frac{k_R}{b_0} \left[\frac{\arctan\left(\frac{h_{dome}}{b_0}\right) + \arctan\left(\frac{x}{b_0}\right)}{\frac{x^2}{b_0^2} + 1} \right] \tag{13}$$

where h_{dome} is the no strain dome height and b_0 is the constant dome radius. Crack pressure is then defined by the pressure that supplies sufficient force to drive x to 0: $\Delta p_{crack} * A_V - F_S(x_*) = 0$, where x_* is the valve crack position of $x = 0$ and Δp_{crack} is the crack pressure.

We begin with design parameters h_{dome} and b_0 based on the printable files provided with [17]. Fixing the valve radius to allow for consistent testing equipment, the curvature of the elliptical dome is then described completely by the dome height. The material constants k_s and k_R parameterize based on the silicone material and the thickness of the valve, which we set to be constant for the entire piece.

Geometry and length of the cuts also affect the material constants. Geometry denotes the number of cuts from the center, while the length of each of those cuts is measured in the tangential plane of the dome from the point P.

k_s and k_R are solved experimentally in [18]. Using their reported values of $k_s = 2.9 \times 10^7 \frac{N}{mm}$ and $k_R = 3.6 \times 10^8 \frac{Nmm}{rad}$ for Dragon Skin 20, we solve for a cracking pressure of $\Delta p = 1.5 \times 10^9$ kPa for a 3 mm h_{dome} and 3.75 mm b_0. This result is orders of magnitude higher than the experimental value, or the pressure capabilities of our system (< 100 kPa). Rather than re-solving these constants, we focus our experiments on matching the pressure characteristics relevant to our use without further bounding of parameters.

B. Parameters

The 10 gear parameters and 5 valve parameters are described in Table I and II respectively. Note that 'h' in z_{sh} and g_{t_h} denotes a half-value, restricting the actual values (sun gear teeth and gear thickness) to even integers. Gear thickness is restricted this way to interface correctly with purchased rods and bearings, which are themselves discrete sizes. Gear parameter values are chosen based on peak efficiency regions from Eqn. 10 and 11 , while valve parameters are chosen to span a design space fitting the fixed 3.75 mm b_0.

PLA and ABS plastics both use 1.75 mm filament. Lithium grease (LI), mineral oil (oil) , and Polytetrafluoroethylene lubricant (PTFE) are spread evenly over the gears by manual spinning prior to testing. Dragon Skin (DS) and Ecoflex (E) silicone rubbers are both mixed in 1:1 ratios.

TABLE I
GEAR PARAMETERS

Parameter	Range/Options
Integer Parameters	
z_{sh}	[6, 9]
z_{r2}	[76, 85]
$g_{\text{t\_h}}$ $[mm]$	[3, 6]
$p1_{\text{offset}}$	[-1, 2]
$p2_{\text{offset}}$	[-1, 2]
Choice Parameters	
material	PLA, ABS
lubricant	None, LI, Oil, PTFE
Float Parameters	
x_{s}	[-0.8, 0.8]
x_{r2}	[-0.25, 1.0]
$Cl[mm]$	[0, 3]

TABLE II
VALVE PARAMETERS

Parameter	Range/Options
Choice Parameters	
material	DS10, DS20, E0030, E0050
cut	triple, x, slit
Float Parameters	
t [mm]	[0.2, 1.5]
h_{dome} [mm]	[1.0, 3.8]
l_{cut} [mm]	[0.75, 2.0]

IV. EXPERIMENTATION

A. Fabrication

Gears: The 3D-printed (Bambu X1C) components of the Wolfrom gearbox include the planetary carrier, sun gear, two planet gears, and two ring gears (see Fig. 1). We print the full gear on a single build plate, streamlining the workflow. Because the stage 1 ring gear was reused across multiple design iterations, only the unique components needed to be reprinted, which significantly reduced print time and material consumption. We sliced all parts using a consistent 15% infill with the Gyroid pattern to maintain mechanical consistency.

Material selection, ABS [Overture] or PLA [Matter Hackers], was defined in the design parameters for each iteration. After printing, we post-process parts to remove supports and clean contact surfaces. We press-fit steel dowel pins [McMaster] through the ball bearings [Yamaso], which we then press-fit into the bores of the planet gears. We mount these gear-bearing-pin assemblies onto the stage 2 planetary carrier and mesh them with the rest of the gearbox. We fasten the complete gearbox onto the 3D-printed fixture, which secures it coaxially with the Dynamixel MX-28 motor.

Valves: The fabrication of each valve directly follows [17] and involves three main steps: mold creation (FormLabs3 - Tough1500), silicone injection-molding, and laser cutting. Molds are designed using 'equation editor' in Solidworks, parameterized for thicknesses and dome height. The valve molds consist of one male and one female part, as dictated by the dome protrusion and the dome cavity, respectively.

One of four materials is used to create the valves - Dragon Skin 10, Dragon Skin 20, Ecoflex 00-30, and Ecoflex 00-50. We use equal parts A and B by weight. We add a thinning agent (NOVOCS) to all Dragon Skin to reduce viscocity and allow degassing. Before and after casting, solutions are degassed in a vacuum chamber. Before injecting, we place the valve molds in a FormCure for 60 minutes at 70° C, wash them with soap and water, and dust with baby powder to ensure that the molds are dry and the valves cure properly.

Valves are set to cure for a minimum of four hours. We place the completed valve in a vessel [PLA] similar to expose the dome of the valve through a circular cutout ($radius = 3\ mm$), allowing us to then laser-cut.

B. Experimentation

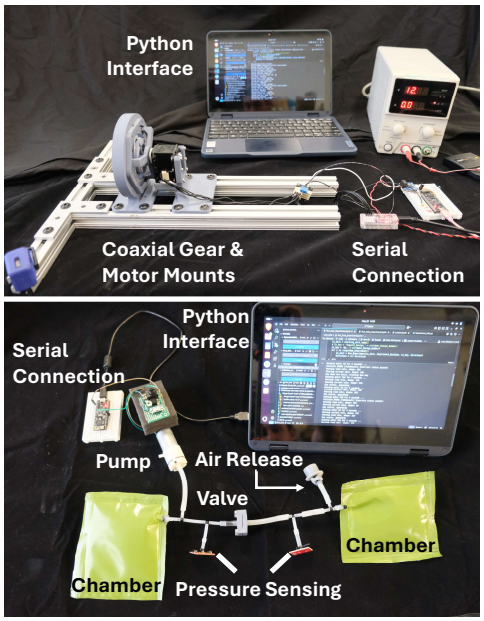

Fig. 2. Experimental test setups for (top) gear spin and (bottom) valve pressurization.

Gears: We mount the gears to a fixed gear mount and coaxially to the DC motor (Dynamixel MX-28). 12 V DC is applied through a current sensor (INA260), monitored via a microcontroller interface (ESP32) and reported serially to the Python interface. The Python interface communicates with the motor through serial communication to drive the motor and record torque load. The motor functions in 'wheel' mode, with a goal speed of 11 RPM for five seconds. Six trials are performed for each gear, three with clockwise rotation and three with counterclockwise rotation.

For each gear, two types of data are measured during testing: torque output[1] as "DXT Torque %" over time (ms), and current draw (mA), measured via the INA260 power sensor. The torque readings represent a percentage of the maximum rated torque as inferred by the built-in sensor of the Dynamixel MX-28. For performance evaluation, the highest torque value from the best five of the six trials is selected and used to compute the reward score.

Valves: We mount the valves in a 3D-printed holder chamber, which is positioned between two air pressure sensors (MPRLS0025), each connected to an ESP32 microcontroller to transmit pressure data. Air pressure is supplied by a 4.5 V DC air pump (ZR370-02PM) on the entry side of the valve. TPU-Coated nylon pressure chambers provide pneumatic capacitance on either side of the valve and compressed open-cell foam acts as a pneumatic resistor on the exit side of the valve. We activate the pump for 50 seconds and record pressure at entry to and exit from the valve.

When pressure increases on the output side of the valve, the corresponding pressure on the inlet side is noted as Δp_{crack}. The average difference between pressures on the inlet and outlet side for all data-points after Δp_{crack} is taken as the steady-state pressure differential Δp_{ss}. These values are verified manually for each valve by looking at the time-dependent pressure data. If a valve does not crack during experimentation, the Δp_{crack} and Δp_{ss} are both recorded as 100 kPa. Δp values are scaled between 0 and 1 by dividing by 100 kPa before they are input into the learning client.

C. Active Learning

To carry out active learning (AL) [7], we use the Ax platform with modeling from BoTorch. We initialize active learning with a set of n=10 and n=5 tests chosen randomly from the Taguchi trials for gears and valves respectively. This is meant to proxy for a space-filling method without requiring additional testing beyond what we had already performed. The number of initialization trials is chosen as equal to the number of design parameters in each case.

Gears: For the single task, single objective gear optimization, we specify a 'AdditiveMapSaasSingleTaskGP' model, which requires minimal hyperparameter tuning [20]. We specify a 'warp' input transform class, SingleTaskGP model class, and Matern-5/2 kernel. Mean-squared-error is chosen as the model evaluation criterion. The acquisition function is detailed as a 'qLogNoisyExpectedImprovement'.

The active learning procedure is then carried out by selecting the gear $M \in \mathbb{R}^{10}$ which best balances exploration and exploitation as determined by the acquisition function. Each iteration, $q = 2$ sets of gear designs are then obtained simultaneously. We continue this procedure iteratively, until enough gears are acquired to match the Taguchi trials.

We use a singular score function for the gear optimization that takes into account both the gear ratio (higher is better) and the torque required to backdrive (lower is better). The score to maximize is:

$$\mathrm{R}_y = \frac{\text{Gear Ratio}}{\max^{n}_{M_i \neq \max_j M_j} M_i} \tag{14}$$

where R_y is the reward for gear design y, $M_i = max(\tau(t))$ is the maximum motor torque over trial i as a percentage of the motor's stall torque, and n is the total number of trials for y. The denominator is simply the maximum torque percentage, ignoring the worst trial. A higher reward is better.

Valves: For the single task multi-objective valve optimization, we use Ax's default BoTorch model class 'SingleTaskGP' and default acquisition function 'qLogNoisyExpectedHypervolumeImprovement'.

The active learning procedure is then carried out by selecting the valve $M \in \mathbb{R}^5$ which best balances exploration and exploitation, as determined by the acquisition function.

[1]These torque values represent inferred values from the MX28 motor driving the gears. While they are relevant predictions for scoring, there can be significant discrepancy in actual torque output from the motor.

We carry out active learning in the parallel setting, where at each iteration $q = 3$ valves are obtained simultaneously for fabrication and testing. We continue this procedure iteratively, until enough valves are acquired to match the Taguchi trials.

Errors to minimize each objective are designated as:

$$\mathrm{L}_y = \sqrt{(p_{\text{target}} - p_{\text{actual}})^2} \tag{15}$$

Where L_y is the loss for valve y and p is the pressure. The equation is the same for both crack pressure and steady-state pressure difference. A lower loss is better.

V. Results

A. Performance

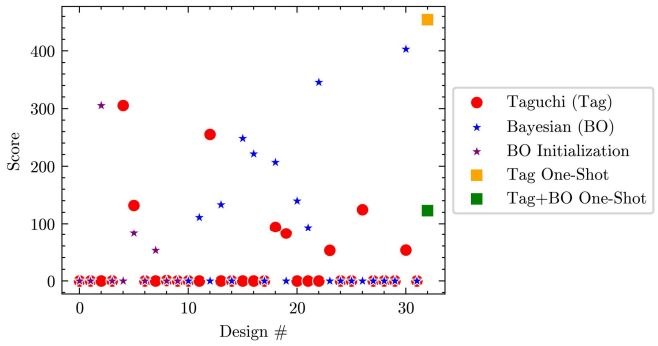

Fig. 3. Reward function score (Eqn. 14) for gear design iterations among Taguchi, Bayesian Optimization, and one-shot BO design strategies.

Gears: The Taguchi trials yielded a maximum score of 305 on the fifth design iteration. The seventh iteration of Bayesian Optimization (BO) outperformed the Taguchi methods with a score of 346. The best overall BO design was on the eleventh iteration, achieving a score of 403. A single 'one-shot' iteration of BO performed with a model that trained on all Taguchi data outperformed all other methods with a score of 454. A similar 'one-shot' iteration of BO with a model that trained on all Taguchi and BO trials did significantly worse with a score of 123.

The best scoring 'one-shot' gearbox uses a gear ratio of 63.6 and backdrives without ever needing more than 0.35 *Nm* of torque. The next best gearbox from BO uses a gear ratio of 56.9 and reaches a maximum backdrive torque of 0.35 *Nm*. The other BO gear of note uses a gear ratio of 107.7 and reaches a maximum backdrive torque of 0.78 *Nm*. The parameters for each of these three gears are as follows:

Material	Lubricant	z_{sh}	z_{r2}	x_s	x_{r2}	Gear Ratio
PLA	LI	6	76	-0.5872	-0.0764	63.6
PLA	None	6	80	-0.80	0.544	56.9
PLA	LI	6	76	0.7513	0.1686	107.7

Cl [mm]	$g_{\text{thickness\_h}}$ [mm]	$p1_{\text{offset}}$	$p2_{\text{offset}}$	Max Torque [Nm]
.033	6	2	2	0.35
3	3	2	-1	0.35
2.71	3	-1	2	0.78

Valves: Because no valve is optimal for both crack pressure and steady state pressure, there is ambiguity among which valve is best. Any output that is not outperformed by some other valve simultaneously in both metrics lies on the so-called Pareto front. Taguchi trials yielded a four-valve Pareto front with crack pressure errors as low as 0.4 kPa and steady-state pressure errors as low as 0.5 kPa. Only one BO trial offered sufficient performance to near this front. A 'one-shot' iteration of BO performed with a model trained on all Taguchi data did not enter the Pareto front, but the one-shot BO trained on all Taguchi and BO trial data outperformed the previous front with crack and steady-state pressure errors of 1 kPa and 0.5 kPa respectively. This cumulative 1.5 kPa error represents a 55% decrease from the next best summed error of 3.3 kPa (Taguchi design 7).

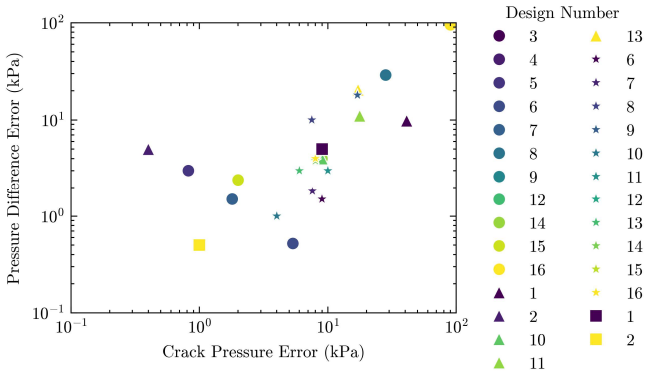

Fig. 4. Valve pressure error values (Eqn. 15) for Taguchi (circle and triangle), Bayesian Optimization (triangle initialization, star generated), and one-shot BO (square) designs.

The parameters for each of the valves in the final, combined Pareto front are as follows:

Design	Material	t [mm]	h_{dome} [mm]	l_{cut} [mm]	cut
2	DS20	0.20	1.0	0.75	X
5	E0030	0.75	1.0	1.00	triple
OS 2	E0050	0.31	1.95	0.75	triple

Δp_{crack} [kPa]	Δp_{ss} [kPa]	Crack Error [kPa]	SS Error [kPa]
10.4	10.0	0.4	5.0
9.2	8.0	0.8	3.0
11.0	4.5	1.0	0.5

B. Modeling

The surrogate model uncertainty is determined sequentially with the addition of each trial for both Taguchi methods and BO methods beyond the 5 or 10 trial initialization. Standard deviation is averaged across 10_000 sets of design parameters from a Latin hypercube [21] spanning the design space as detailed in Table I. We used the trained model to predict the outcome score or crack pressure at each set of parameters, and the average standard deviation from these predictions is used as our measure of model uncertainty.

VI. Discussion

Gears: Looking at the uncertainty metrics in Fig. 5, we are likely functioning on too-few data points for successful

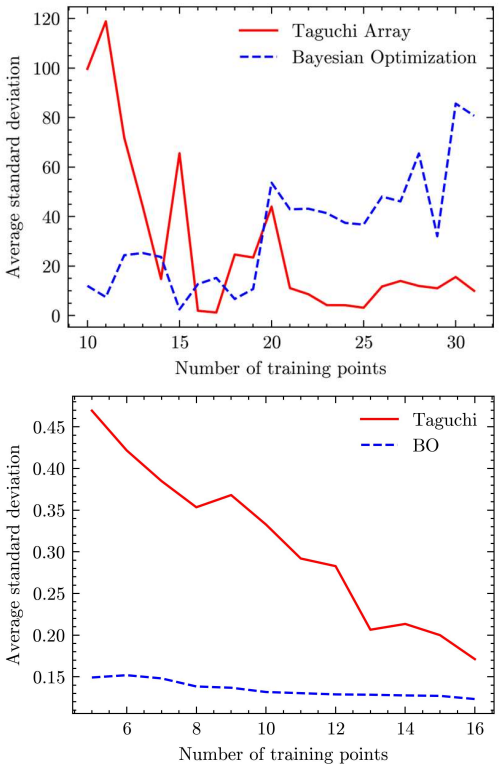

Fig. 5. Prediction uncertainty for gear (top) and valve (bottom) models. Uncertainty (standard deviation) versus number of iterations included in the training set for a Gaussian Process Regression (GPR) model trained on BO (blue) or Taguchi (red) data.

training of the Gaussian Process Regression (GPR) model. Our standard deviations are still on the rise for the BO trials. The Ax platform acknowledges that the model is not sufficiently trained upon request for best output with the warning: "Metric [score] was unable to be reliably fit." The initially low standard deviations after the BO initialization (n=10) imply an overconfidence with insufficient training data. Standard deviations for the Taguchi set are starting to reduce near the completion of the Taguchi trials (n>25).

Despite the high model uncertainty, the expected improvement (EI) acquisition function is still successful in outperforming the space-filling Taguchi method. The highest score comes from a mixture of the two methods, with the one-shot BO using a model trained on Taguchi data only.

Looking across all trials, results qualitatively indicated that PLA with a lithium grease lubricant performed best overall, along with a mix of low z_{r2} with high z_s. Relatively low clearance and high $p_{2\ offset}$ also seemed to contribute to positive results. Quantitatively, Ax provides second-order parameter sensitivity analysis for each surrogate model. The metrics provided by the surrogate model trained on Taguchi data (the model that returned the highest performing design) predict that z_{r2} is the best predictor of score, followed closely by material. These predictions are substantially different from those of the model trained during BO trials, which predicts a linear combination of z_s and $p_{1\ offset}$ as the best predictor of score, followed by $p_{1\ offset}$.

Valves: Looking at the uncertainty metrics in Fig. 5, the Taguchi trials provide excellent model training data for reducing uncertainty. Again, the initially low standard deviations after the BO initialization (n=5) imply an overconfidence, which remains throughout all executed tests.

In this case study, the model overconfidence and limited initialization data are correlated to over-investigation of a local minimum, with the model suggesting designs similar to Taguchi design 2: using DS20 material and 'X' cut geometry. This is likely because these were the materials used in the only member of the Taguchi Pareto front in initialization.

Looking across all trials, the lowest combined error (by both average and RMSE) across both crack and steady-state pressure came from the one-shot BO trained on all Taguchi and BO trials, which used Ecoflex 00-50, and a 'triple' cut geometry. The metrics provided by the corresponding surrogate model predict that valve thickness has the largest effect on both Δp_{crack} and Δp_{ss}.

VII. Conclusions & Future Work

In this work, we use Bayesian Optimization for design optimization across two case studies:3D-printed gearboxes for large gear ratios and backdrivability, and pressure response of a silicone check-valve. We show that BO improves score functions for desired outputs relative to a Taguchi baseline approach in both case studies, and evaluate metrics of surrogate model uncertainty as it is trained on the two test sets. Despite performance improvements in both studies, we predict further model training is required to reach optimal results based on lingering epistemic uncertainty.

Based on these two case studies, it seems that sequential testing with BO is not yet in the 'plug-and-play' regime to match that of Taguchi arrays. Based on 'one-shot' methods reaching the highest scores in both instances, Taguchi data seems more valuable than the data points chosen by Bayesian Optimization (Expected Improvement (EI)) for model training. The relative performance of the 'one-shot' BO models also implies that sufficiently spanning training data is the most valuable factor for improving overall BO performance. Looking at the standard deviation metric for both cases, training on Taguchi data allows for a more consistent decrease in model standard deviation. The valve results from sequential testing further emphasize this issue, as the minimal initialization set correlated to model overconfidence and overemphasis on exploitation over exploration. Further work is needed to find the correct point at which to transition from a model-improving acquisition function (i.e., minimizing total uncertainty) to greedier optimization function (i.e., EI).

The outputs of our machine learning approach to mechanical design provide helpful sensitivity analysis for understanding relative importance of model parameters. However, the volatility of this analysis (i.e., as we alter training data across trials) indicates that we may still be far from removing epistemic uncertainty from our models.

Further work will be required to investigate the efficacy of alternate surrogate and acquisition functions in this type of

sequential experimentation. Due to the nature of our physical design experiments, it is expensive to run sequential testing for multiple model strategies. Therefore, open access to the data from individual case studies will be instrumental in enabling a review across many design problems.

We were successful in using BO to identify precise valve parameters that match the desired pressure differentials for our use case, but struggled to efficiently reach these results during sequential experimentation. The surrogate model seemed to require a spanning dataset beyond just five initialization trials before it could accurately predict the relation between input and output parameters.

We were also successful in identifying multiple options for 3D-printed gears with sufficiently high gear ratio and low back-drive torque. The minimum $0.35Nm$ torque estimation is an order of magnitude greater than torques achieved in [14]. This is the cost for a far more accessible and adaptable 3D-printed assembly. In this case study, sequential experimentation also quickly improved scores beyond the Taguchi baseline. While a number of parameters seem to be optimized, perfecting these results and scaling the gears for individual mechanical systems is left to future works. This can also include expanding the parameter ranges for future BO, which would no longer need to match the baseline ranges.

ACKNOWLEDGMENTS

Thanks to Jack Granite for assistance with polymer curing.

DATA AVAILABILITY

The majority of the work above was coded in Python, including in Jupyter Notebooks. This code, select CAD files, and .pkl files containing dictionaries of all the test data used herein, is available on Github: https://github.com/gmcampbell/Mod_DOE.

REFERENCES

[1] S. Yang and Y. F. Zhao, "Additive manufacturing-enabled design theory and methodology: A critical review," *The International Journal of Advanced Manufacturing Technology*, vol. 80, no. 1, pp. 327–342, 2015.

[2] D. C. Montgomery, *Design and Analysis of Experiments, 8th Edition*. John Wiley & Sons, Incorporated, 2013, ISBN: 9781118146927.

[3] T. J. Robinson, C. M. Borror, and R. H. Myers, "Robust parameter design: A review," *Quality and reliability engineering international*, vol. 20, no. 1, pp. 81–101, 2004.

[4] G. M. Campbell, G. Muhaxheri, L. F. Guilhoto, *et al.*, "Active learning design: Modeling force output for axisymmetric soft pneumatic actuators," *arXiv preprint arXiv:2504.01156*, 2025.

[5] N. Feizi, S. F. Atashzar, M. R. Kermani, and R. V. Patel, "Modeling and high-definition control of a smart electroadhesive actuator: Toward application in rehabilitation," *IEEE Transactions on Medical Robotics and Bionics*, vol. 4, no. 4, pp. 1057–1067, 2022.

[6] R. Arboretti, R. Ceccato, L. Pegoraro, and L. Salmaso, "Design of experiments and machine learning for product innovation: A systematic literature review," *Quality and Reliability Engineering International*, vol. 38, no. 2, pp. 1131–1156, 2022.

[7] P. Ren, Y. Xiao, X. Chang, *et al.*, "A survey of deep active learning," *ACM computing surveys (CSUR)*, vol. 54, no. 9, pp. 1–40, 2021.

[8] S. Greenhill, S. Rana, S. Gupta, P. Vellanki, and S. Venkatesh, "Bayesian optimization for adaptive experimental design: A review," *IEEE access*, vol. 8, pp. 13 937–13 948, 2020.

[9] M. Balandat, B. Karrer, D. Jiang, *et al.*, "Botorch: A framework for efficient monte-carlo bayesian optimization," *Advances in neural information processing systems*, vol. 33, pp. 21 524–21 538, 2020.

[10] K. Wang and A. W. Dowling, "Bayesian optimization for chemical products and functional materials," *Current Opinion in Chemical Engineering*, vol. 36, p. 100 728, 2022.

[11] A. E. Gongora, B. Xu, W. Perry, *et al.*, "A bayesian experimental autonomous researcher for mechanical design," *Science advances*, vol. 6, no. 15, eaaz1708, 2020.

[12] J. R. Deneault, J. Chang, J. Myung, *et al.*, "Toward autonomous additive manufacturing: Bayesian optimization on a 3d printer," *MRS Bulletin*, vol. 46, pp. 566–575, 2021.

[13] E. Bakshy, L. Dworkin, B. Karrer, *et al.*, "Ae: A domain-agnostic platform for adaptive experimentation," in *Conference on neural information processing systems*, 2018, pp. 1–8.

[14] H. Matsuki, K. Nagano, and Y. Fujimoto, "Bilateral drive gear—a highly backdrivable reduction gearbox for robotic actuators," *IEEE/ASME Transactions On Mechatronics*, vol. 24, no. 6, pp. 2661–2673, 2019.

[15] G. Gambirasio, M. Pesenti, M. Panzenbeck, M. Gandolla, L. Roveda, M. Covarrubias, *et al.*, "Design and integration of a low-back exoskeleton: A 3d-printed cycloidal drive actuator for flexible human-robot interaction.," *Computer-aided Design and Applications*, vol. 21, pp. 791–806, 2024.

[16] A. Wang and S. Kim, "Directional efficiency in geared transmissions: Characterization of backdrivability towards improved proprioceptive control," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2015, pp. 1055–1062.

[17] L. C. van Laake, J. de Vries, S. M. Kani, and J. T. Overvelde, "A fluidic relaxation oscillator for reprogrammable sequential actuation in soft robots," *Matter*, vol. 5, no. 9, pp. 2898–2917, 2022.

[18] L. C. van Laake, A. Comoretto, and J. T. Overvelde, "On the coexistence of pressure regulation and oscillation modes in soft hysteretic valves," *Journal of Fluids and Structures*, vol. 126, p. 104 090, 2024.

[19] T. Yada, "Review of gear efficiency equation and force treatment," *JSME international journal. Ser. C, Dynamics, control, robotics, design and manufacturing*, vol. 40, no. 1, pp. 1–8, 1997.

[20] D. Eriksson and M. Jankowiak, "High-dimensional bayesian optimization with sparse axis-aligned subspaces," in *Uncertainty in Artificial Intelligence*, PMLR, 2021, pp. 493–503.

[21] J. C. Helton and F. J. Davis, "Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems," *Reliability Engineering & System Safety*, vol. 81, no. 1, pp. 23–69, 2003.