

Function-on-function kriging, with applications to 3D printing of aortic tissues

Jialei Chen^{*†‡}, Simon Mak[§], V. Roshan Joseph^{*} and Chuck Zhang^{*† ¶}

March 2, 2020

Abstract

3D-printed medical prototypes, which use synthetic metamaterials to mimic biological tissue, are becoming increasingly important in urgent surgical applications. However, the mimicking of tissue mechanical properties via 3D-printed metamaterial can be difficult and time-consuming, due to the functional nature of both inputs (metamaterial structure) and outputs (mechanical response curve). To deal with this, we propose a novel function-on-function kriging model for efficient emulation and tissue-mimicking optimization. For functional inputs, a key novelty of our model is the spectral-distance (SpeD) correlation function, which captures important spectral differences between two functional inputs. Dependencies for functional outputs are then modeled via a co-kriging framework. We further adopt shrinkage priors on both the input spectra and the output co-kriging covariance matrix, which allows the emulator to learn and incorporate important physics (e.g., dominant input frequencies, output curve properties). Finally, we demonstrate the effectiveness of the proposed SpeD emulator in a real-world study on mimicking human aortic tissue, and show that it can provide quicker and more accurate tissue-mimicking performance compared to existing methods in the medical literature.

Keywords: Computer experiment; Gaussian process; Metamaterial; Sparsity; Tissue-mimicking; Translation-invariance.

^{*}H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology

[†]Georgia Tech Manufacturing Institute, Georgia Institute of Technology

[‡]Corresponding author

[§]Department of Statistical Science, Duke University

[¶]This work is supported by a U.S. National Science Foundation grant CMMI-1921646 and Piedmont Heart Institute.

1 Introduction

Three dimensional (3D) printing is an emerging technique of layer-by-layer additive manufacturing technology, with growing interest in medical applications (Rengier et al., 2010). This is because 3D-printed prototypes can provide precise mimicking of organ shape at an acceptable price and time cost (Ventola, 2014). Such prototypes can be extremely helpful for doctors to practice and be proficient in surgical procedures (Chen et al., 2018b) as well as personalized pre-surgical planning (Qian et al., 2017). One limitation is that the mechanical property (i.e., stress-strain curve) of printed prototypes is completely different from biological tissues (Raghavan et al., 1996). Currently, the state-of-the-art approach is to embed *metamaterial* structure to mimic the desired mechanical property of biological tissue (Wang et al., 2016a,b; see Figure 1). However, the optimization for this mimicking may take days or even weeks to perform, due to the *functional* nature of the metamaterial structure. This greatly limits the medical applicability of tissue-mimicking prototypes, since surgery timing is a critical factor for outcome success. In this paper, we propose a novel kriging model for emulating functional mechanical response over the design space of functional metamaterial structure, which can be used for efficient tissue-mimicking optimization in practical turnaround times.

There are two key reasons why state-of-the-art tissue-mimicking methods are impractical for *urgent* surgical needs. Firstly, such methods rely solely on both physical experiments and computer experiments, which are expensive and/or time-intensive to run. In particular, a single physical experiment (3D-printing and testing a prototype) takes hours to perform, and a single computer experiment (finite element analysis) requires at least 30 minutes for a reliable mechanical response simulation. Secondly, to optimize for a good structure which mimics the mechanical response of biological tissue, such methods require *many* experimental runs over the design space of functional metamaterial structures. This makes current tissue-mimicking methods prohibitively expensive for urgent surgical applications, where the tissue-mimicking prototype is needed within a day. One strategy (which we




Figure 1: (a) 3D-printed aortic valve (no metamaterial structure), (b) stress-strain curves of biological tissue and printable polymer, (c) a numerical (finite element) simulation case with sinusoidal metamaterial, (d) 3D-printed aortic valve with tissue-mimicking metamaterial.

adopt) is to train a *surrogate* model (or *emulator*, see Santner et al., 2013) which, given data over the design space, can efficiently *predict* the mechanical response of an untested metamaterial structure. However, due to the expensive nature and functional complexities, it is necessary to integrate the rich physics of the tissue-mimicking problem within the emulator model specification, in order to achieve accurate mimicking in a timely fashion.

The proposed emulator utilizes a technique called *kriging* (Matheron, 1963), which models the unknown simulation output via a Gaussian process (GP). Kriging is widely used in computer experiment modeling for its interpolating property, and the fact that both the predictor and its uncertainty have closed-form expressions (Santner et al., 2013). The literature on kriging for *functional outputs* typically involve some form of reduced-basis modeling, including principal component models (Higdon et al., 2008; Mak et al., 2018) and wavelet models (Bayarri et al., 2007), where only the weights of the bases need to be modeled. Other works utilize a co-kriging framework (Stein and Corsten, 1991) for modeling a discretization of functional output – such an approach typically works well when the output is highly correlated in the functional domain (Banerjee et al., 2014). For *functional inputs*, several modeling techniques have been proposed in the functional data analysis literature (see, e.g., Ramsay, 2005), including varying-coefficient models (Fan and Zhang, 2008) and historical functional linear models (Malfait and Ramsay, 2003). However, the literature on kriging with functional inputs is scarce. For the specific case of time-series

inputs, Morris (2012) proposed a kriging model with a covariance function depending on time order. Reduced-basis models were also proposed in Muehlenstaedt et al. (2017); Wang and Xu (2019). Such models, however, do not incorporate prior physical knowledge of the tissue-mimicking problem, and can therefore yield poor emulation and mimicking performance given the paucity and functional complexities of the experimental data.

To address this, we introduce in this work a new function-on-function kriging model which integrates an important source of physics: the spectral information of the functional metamaterial structure input. Previous studies (Wang et al., 2016a; Chen et al., 2018a) have shown that spectral differences in metamaterial structures result in variations to their mechanical properties. To incorporate this spectral information, we propose a new *spectral-distance* (or SpeD) correlation function, which uses the spectral-distance – the (weighted) Euclidean distance between two functional inputs in spectral domain – to model the process correlation of the GP. This new correlation function captures the appealing property of *translation-invariance*, where two input metamaterial structures which are the same except for a translation shift have the same mechanical properties. We then integrate this within a co-kriging framework for modeling the functional mechanical response output. A key novelty of our model is the incorporation of sparsity on the input spectral coefficients and the output co-kriging matrix. This sparsity serves two purposes for the tissue-mimicking problem. First, it allows us to extract important physics (e.g., dominant input frequencies, output curve properties), which furthers understanding of the complex problem. Second, by identifying key frequencies which most influence the mechanical response (thereby reducing the large functional input space), we can greatly speed up the tissue-mimicking optimization under the proposed emulator. This emulator-based mimicking approach allows for efficient and accurate mimicking of biological tissue, and broadens the applicability of printed prototypes for urgent surgical procedures.

The paper is structured as follows. Section 2 gives an overview of the tissue-mimicking problem and the underlying computer experiment. Section 3 presents the proposed SpeD emulation model and its shrinkage prior specification. Section 4 outlines the algorithm

for parameter estimation. Section 5 investigates the emulation accuracy, uncertainty quantification, physics extraction and computation time of the SpeD emulator in a real-world tissue-mimicking case study. Section 6 concludes the work.

2 Tissue-mimicking and finite element modeling

We first describe the tissue-mimicking problem (or the metamaterial design problem) and explain the physics of this problem. We then introduce the finite element (FE) analysis as a simulation tool, and provide a brief discussion on experimental design for the FE simulation.

2.1 Tissue-mimicking problem

As discussed in Section 1, 3D-printing technology can print patient-specific prototypes with precise geometry (Figure 1 (a)), but the mechanical properties of these printed prototypes can differ greatly from that for true organs (Figure 1 (b)). The considered mechanical property is the *stress-strain* curve (Malvern, 1969), defined as stress (the external load per area) as a function over strain (the overall deformation percentage)¹. The stress-strain curve of the biological tissue typically possesses the property of *strain-stiffening*, which means the curve is concave upward (see solid blue line in Figure 1 (b)), indicating it becomes stiffer as more load is introduced (Raghavan et al., 1996). However, for 3D-printable material, an opposite property of *strain-softening* is exhibited (see dotted red line in Figure 1 (b)) due to the plastic-slipping effect and energy dissipation (Hill, 1998).

To achieve the strain-stiffening property of the biological tissues, one approach is to introduce *metamaterial* structure (i.e., printed enhancement sub-structure) within the prototypes (Wang et al., 2016a). Figure 1 (c) shows an example of a metamaterial with sinusoidal structure. Here, the *stiffer* enhancement fiber is designed to have a sinusoidal shape, inside the cuboid matrix of a *soft* material. In this work, we treat the structure

¹Note that here, both stress and strain are not local values, rather averaged quantities over a sufficient volume to yield a continuum description, i.e., representative elementary volume.

(or shape) of the enhancement fiber (assumed to have uniform diameter) as the functional input for our SpeD model. Our goal is to mimic the target mechanical property of human tissues, by carefully choosing the shape of the enhancement fiber. Figure 1 (d) shows a printed “tissue-mimicking” aortic valve with the optimal metamaterial structure.

2.2 Finite element modeling and experimental design

In this work, FE modeling is used to simulate the output stress-strain curve of a given metamaterial structure. FE modeling is frequently used for stress analysis in solid mechanics; it transforms the partial differential equations to their integral form, so that a piece-wise linear formula can be used to approximate the true deformation profile (Zienkiewicz et al., 1977). The key advantage of FE simulations, compared to physical experiments, is that high accuracy can be achieved with no material cost or human error.

Here, FE simulations are performed using COMSOL Multiphysics. The overall size of the metamaterial cuboid (with one enhancement fiber inside) is $20mm$ by $4mm$ by $2mm$, with physics-based quadratic tetrahedral elements for meshing. To calculate the stress-strain curve of the metamaterial, one end of the cuboid is fixed while a series of load levels (up to 15% uniaxial deformation) is applied to the other end. The total computation time for one metamaterial design is around 30 minutes on 24 Intel Xeon E5-2650 2.20GHz processing cores. Figure 2 (a) shows the stress-strain curve obtained by FE simulation.

We use a sinusoidal wave structure for designing the training metamaterial structures, as such a form exhibits the best strain-stiffening property from a recent study (Wang et al., 2016b). The design space has four parameters: the diameter of the enhancement fiber d , and the amplitude A , frequency ω and initial phase ϕ of the sinusoidal wave:

$$I(t) = A \sin(\omega t + \phi). \quad (1)$$

The experimental design adopted for the sinusoidal coefficients is the maximum projection (MaxPro) design (Joseph et al., 2015). We note that the parametric sinusoidal form (1)




Figure 2: (a) The stress-strain output of the FE simulation. (b) The pairwise scatterplot of the training input (blue circle) and testing input (red cross) in the space of sinusoidal coefficients (1).

is used only to *generate* data for training the emulator; we will explore a bigger non-parametric input space for *prediction* and tissue-mimicking *optimization*. A total of $n = 58$ metamaterial structures are simulated over the design space as the training dataset, and an 18-run Sobol' sequence (Sobol', 1967) is used to generate the testing dataset, also on the same design space. Figure 2 (b) shows the pairwise scatterplots for the training and testing designs. Despite the relatively small training dataset ($n = 58$ samples), we show later that the functional stress-strain predictions from the proposed emulator are quite accurate, and provides noticeable improvements over a standard kriging model with four sinusoidal coefficients as inputs.

3 Emulation model

We now present the proposed emulation model in three parts. First, we introduce the proposed model for functional inputs, using the simplified setting of scalar outputs. We then extend this for functional outputs using a co-kriging structure, and provide closed-form expressions for prediction and uncertainty quantification (UQ). Finally, we discuss a prior specification for model parameters which encourages sparsity. A key theme of this section is

the incorporation of known physics within the emulation model. As mentioned in Section 1, this integration of known physical properties allows for improved predictive performance of the function-on-function kriging model for mechanical responses (see Sections 5.1 and 5.2), but also greatly reduces the computation time for tissue-mimicking optimization, which is crucial for urgent surgical scenarios. Table 1 summarizes these known physical properties and the corresponding emulator assumptions, which we discuss throughout the section.

3.1 Spectral-distance kriging model

We introduce first the proposed kriging model for functional inputs $I(\cdot) \in \mathcal{I}$, where $\mathcal{I}$ is the functional input space (to be defined later). For simplicity, assume first the case of scalar outputs (functional outputs are introduced next). For the map $y(\cdot) : \mathcal{I} \mapsto \mathbb{R}$ from functional inputs to scalar outputs, we propose the following GP model:

$$y(\cdot) \sim \text{GP}\{\mu, \sigma^2 \rho(\cdot, \cdot)\}, \quad (2)$$

where μ is the scalar process mean and σ^2 is the process variance. Here, $\rho(\cdot, \cdot) : \mathcal{I} \times \mathcal{I} \mapsto \mathbb{R}$ is the proposed spectral-distance (SpeD) correlation function, defined as:

$$\rho(I_1(\cdot), I_2(\cdot)) = \text{Corr}\{y(I_1(\cdot)), y(I_2(\cdot))\} = \exp\left(-D^2(|\mathcal{F}[I_1(\cdot)]|, |\mathcal{F}[I_2(\cdot)]|; \theta)\right). \quad (3)$$

In (3), $D(\cdot, \cdot; \theta)$ is a distance function (defined later), $|a\mathbf{i} + b|$ is the modulus of a complex number $a\mathbf{i} + b$ (where $\mathbf{i} = \sqrt{-1}$ is the unit imaginary number), and $\mathcal{F}[\cdot] : \mathcal{I} \rightarrow \hat{\mathcal{I}}$ is the Fourier transform from the input space of integrable functions, $\mathcal{I} = \{I(\cdot) : \int |I(t)|dt < \infty\}$, to its spectral space $\hat{\mathcal{I}}$. We will use the following definition of a Fourier transform for an input function $I(\cdot)$:

$$\hat{I}(\xi) = \mathcal{F}[I(t)] = \int I(t)e^{-2\pi\mathbf{i}t\xi}dt, \quad \xi \in \mathbb{R}. \quad (4)$$

Similar to the scale-parametrized distance function in the Gaussian correlation (which is widely used for GP emulation of computer experiments, see Santner et al., 2013), we will

Physical Properties	Model Assumptions
Translation-invariance property (Chen et al., 2018a)	Spectral-distance correlation function
Frequency information is important (Chen et al., 2018a)	
Only a few important frequencies (Chen et al., 2018a)	Sparsity in frequency coefficients
Known structure of mechanical response (Callister and Rethwisch, 2011)	Sparsity in co-kriging covariance matrix
Monotonicity of mechanical response (Malvern, 1969)	Basis representation of output mean

Table 1: *Known physical properties and corresponding assumptions for the emulation model.*

use the following scale-parametrized l_2 distance function in the *spectral* domain:

$$D(|\mathcal{F}[I_1(\cdot)]|, |\mathcal{F}[I_2(\cdot)]|; \theta) = \left[\int \theta(\xi) \left(|\hat{I}_1(\xi)| - |\hat{I}_2(\xi)| \right)^2 d\xi \right]^{1/2}. \quad (5)$$

Here, $\theta(\cdot)$ is a weight function in spectral space, with a larger value of $\theta(\xi)$ indicating greater importance of frequency ξ in the SpeD correlation function. In contrast to the standard Gaussian correlation, we assign importance to each *frequency component* of a functional input, rather than to each input *variable*. Plugging (5) into (3), the SpeD correlation function becomes:

$$\rho(I_1(\cdot), I_2(\cdot)) = \exp \left(- \int \theta(\xi) \left(|\hat{I}_1(\xi)| - |\hat{I}_2(\xi)| \right)^2 d\xi \right). \quad (6)$$

We will show later in Theorem 1 that $\rho(\cdot, \cdot)$ is indeed a positive semi-definite kernel.

One advantage of the SpeD correlation function is that it can capture known properties of the tissue-mimicking problem. First, recall that the *translation-shifting* property of Fourier transform (Bracewell and Bracewell, 1986): for any $t_0 > 0$, if $I_2(t) = I_1(t - t_0)$, then

$$\hat{I}_2(\xi) = e^{-2\pi i t_0 \xi} \hat{I}_1(\xi). \quad (7)$$




Figure 3: An illustration of the translation-invariance property: for the two input structures which are equivalent up to a translation shift of t_0 , their mechanical responses are the same.

For two metamaterial structures with a shift, i.e., $I_1(t) = I(t)$ and $I_2(t) = I(t - t_0)$, we can then show that their outputs are perfectly correlated, i.e.:

$$\rho(I_1(\cdot), I_2(\cdot)) = \exp \left(- \int \theta(\xi) \left(\left| \hat{I}_1(\xi) \right| - \left| e^{-2\pi i t_0 \xi} \hat{I}_1(\xi) \right| \right)^2 d\xi \right) = 1. \quad (8)$$

We call this the *translation-invariance* property of the SpeD correlation. As illustrated in Figure 3, this is a desirable property, since we know from physical knowledge that any translation of the metamaterial structure does not affect the output mechanical response. To contrast, the existing functional input models in Section 1 do not enjoy this property. Second, it is known that the stress-strain curve depends largely on frequency ω and amplitude A , but not on initial phase ϕ in the sinusoidal parametrization (1) (Wang et al., 2016a,b; Chen et al., 2018a). One can therefore expect that (i) the Fourier frequencies ξ are significant, and (ii) variations in mechanical response are largely due to differences in frequency intensities $|\hat{I}(\xi)|$. The proposed correlation function (6) nicely captures both of these properties.

For our tissue-mimicking problem, the specific choice of the Fourier transform with l_2 distance of the modulus gives an intuitive parametrization of known physical properties. For other applications, the SpeD correlation (6) can also be used with other spectral transforms (e.g., wavelet transforms) and other distance metrics (e.g., l_1 distance). The choice of spectral transform and distance should be made on a case-by-case basis, motivated by prior

information from the problem at hand.

The following theorem ensures that the SpeD correlation function $\rho(\cdot, \cdot)$ (6) is a valid positive semi-definite kernel.

Theorem 1. *The SpeD correlation function $\rho(\cdot, \cdot) : \mathcal{I} \times \mathcal{I} \mapsto \mathbb{R}$ in (6), is a positive semi-definite kernel, i.e.:*

$$\sum_{i=1}^n \sum_{j=1}^n c_i c_j \rho(I_i(\cdot), I_j(\cdot)) \geq 0, \quad (9)$$

holds for any $n \in \mathbb{N}$, $c_1, \dots, c_n \in \mathbb{R}$ and any distinct functions $I_1(\cdot), \dots, I_n(\cdot) \in \mathcal{I}$.

The proof of Theorem 1 is provided in Appendix A. This positive semi-definite property guarantees the validity of $\rho(\cdot, \cdot)$ as a proper correlation function to use for GP modeling. Note that $\rho(\cdot, \cdot)$ is not (strictly) *positive-definite*, in that an equality in (9) does not imply $c_i = 0$ for all $i = 1, \dots, n$. This can be seen by setting all input functions $(I_i(\cdot))_{i=1}^n$ to be the same modulo a translation shift; the resulting correlation matrix $[\rho(I_i(\cdot), I_j(\cdot))]_{i=1, j=1}^n$ then becomes a matrix of ones, which is clearly not positive definite. The fact that $\rho(\cdot, \cdot)$ is not positive-definite is not an issue, since for most space-filling designs (including the adopted MaxPro design, see Joseph et al. 2015), all training input functions are distinct even after translation shifts.

3.2 Spectral-distance co-kriging model

For the tissue-mimicking problem, the output – the stress-strain curve – is of functional form as well. Below, we generalize the scalar model in Section 3.1 to account for functional outputs. Denote the functional input as $I(\cdot) \in \mathcal{I}$ and functional output as $O(\cdot)$, where $O(s)$ is the output stress at strain level s . For our training dataset of $n = 58$ simulated structures, the functional outputs $O_i(\cdot), i = 1, \dots, n$ are discretized into m levels, yielding output vectors $\mathbf{y}_i \in \mathbb{R}^m, i = 1, \dots, n$. We assume the following SpeD co-kriging model on $\mathbf{y}(\cdot) : \mathcal{I} \mapsto \mathbb{R}^m$:

$$\mathbf{y}(\cdot) \sim \text{GP}\{\boldsymbol{\mu}, \mathbf{C}(\cdot, \cdot)\}, \quad (10)$$

where $\boldsymbol{\mu} \in \mathbb{R}^m$ is the process mean vector and $\mathbf{C}(\cdot, \cdot) : \mathcal{I} \times \mathcal{I} \mapsto \mathbb{R}^{m \times m}$ is the corresponding covariance matrix function.

Consider first the specification of the covariance matrix function $\mathbf{C}(\cdot, \cdot)$. Let

$$\mathbf{C}(I_1(\cdot), I_2(\cdot)) = \text{Cov}(\mathbf{y}(I_1(\cdot)), \mathbf{y}(I_2(\cdot))) = \rho(I_1(\cdot), I_2(\cdot))\boldsymbol{\Sigma}. \quad (11)$$

Here, $\rho(\cdot, \cdot)$ is the SpeD correlation kernel in (6), and $\boldsymbol{\Sigma} \in \mathbb{R}^{m \times m}$, $\boldsymbol{\Sigma} \succeq 0$ is a symmetric, positive definite co-kriging covariance matrix quantifying correlations between different output levels.

Equation (11) implicitly assumes separability in the co-kriging covariance structure. Here, separability means the covariance between output levels observed at different functional inputs can be decomposed as the product of the covariance between output levels and the covariance between functional inputs. This separability assumption is used extensively in the literature for reducing computational complexity (Banerjee et al., 2014). However, such an assumption can be justified by the physics of the tissue-mimicking problem as well. Here, the stress-strain curve output $O(\cdot)$ is an inherent property pertaining to the metamaterial structure itself, and the stress value at any specific strain level is not directly related to the metamaterial structure $I(\cdot)$ (Diest, 2013). Thus, the covariance between two strain levels s_1 and s_2 in the stress-strain curve should be similar between different metamaterial structures. Separability allows us to efficiently train our SpeD emulator for timely tissue-mimicking in real-world surgical problems.

Consider next the specification of the mean vector $\boldsymbol{\mu}$. We assume $\boldsymbol{\mu}$ follows the basis representation:

$$\boldsymbol{\mu} = \mathbf{P}\boldsymbol{\beta}, \quad (12)$$

where each column of $\mathbf{P} \in \mathbb{R}^{n \times q}$ represents a pre-specified basis function and $\boldsymbol{\beta} \in \mathbb{R}^q$ denotes its corresponding coefficients. The choice of basis functions in $\mathbf{P}$ should be guided by prior knowledge on the form of the output stress-strain curve. We will describe in Section 5.1 a specific parametrization of $\boldsymbol{\mu}$ which incorporates monotonicity information on the

stress-strain curve.

From this, we can derive the equations for prediction and UQ. Let $\mathbf{y}_{1:n} = [\mathbf{y}_1^T, \mathbf{y}_2^T, \dots, \mathbf{y}_n^T]^T$ denote the vector of functional outputs of the whole training set. Using the conditional distribution formula of the multivariate normal distribution, the discretized functional response $\mathbf{y}_{\text{new}}$ at a new functional input $I_{\text{new}}(\cdot) \in \mathcal{I}$ follows the multivariate normal distribution:

$$\mathbf{y}_{\text{new}}|\mathbf{y}_{1:n} \sim \mathcal{N}\left(\mathbf{P}\boldsymbol{\beta} + (\mathbf{r}_\theta \otimes \boldsymbol{\Sigma})^T (\mathbf{R}_\theta^{-1} \otimes \boldsymbol{\Sigma}^{-1}) (\mathbf{y}_{1:n} - \mathbf{1}_n \otimes \mathbf{P}\boldsymbol{\beta}), \right. \\ \left. \boldsymbol{\Sigma} - (\mathbf{r}_\theta \otimes \boldsymbol{\Sigma})^T (\mathbf{R}_\theta^{-1} \otimes \boldsymbol{\Sigma}^{-1}) (\mathbf{r}_\theta \otimes \boldsymbol{\Sigma})\right), \quad (13)$$

where $\otimes$ is the Kronecker product, $\mathbf{1}_n$ denotes 1-vector of n elements, $\boldsymbol{\beta}$ and $\boldsymbol{\Sigma}$ are regression coefficients and co-kriging covariance matrix, $\mathbf{r}_\theta = [\rho(I_{\text{new}}(\cdot), I_1(\cdot)), \dots, \rho(I_{\text{new}}(\cdot), I_n(\cdot))]^T$ and $\mathbf{R}_\theta = [\rho(I_i(\cdot), I_j(\cdot))]_{i=1, j=1}^n$. After algebraic manipulations, the posterior mean $\hat{\mathbf{y}}_{\text{new}} = \mathbb{E}\{\mathbf{y}_{\text{new}}|\mathbf{y}_{1:n}\}$ and posterior variance $\text{Var}\{\mathbf{y}_{\text{new}}|\mathbf{y}_{1:n}\}$ can be written in a more concise form:

$$\hat{\mathbf{y}}_{\text{new}} = \mathbb{E}\{\mathbf{y}_{\text{new}}|\mathbf{y}_{1:n}\} = \mathbf{P}\boldsymbol{\beta} + (\mathbf{r}_\theta^T \mathbf{R}_\theta^{-1} \otimes \mathbf{I}_m) (\mathbf{y}_{1:n} - \mathbf{1}_n \otimes \mathbf{P}\boldsymbol{\beta}), \quad (14)$$

$$\text{Var}\{\mathbf{y}_{\text{new}}|\mathbf{y}_{1:n}\} = (1 - \mathbf{r}_\theta^T \mathbf{R}_\theta^{-1} \mathbf{r}_\theta) \boldsymbol{\Sigma}, \quad (15)$$

where $\mathbf{I}_m$ denotes an $m \times m$ identity matrix. Equation (14) can be used to predict (or emulate) the stress-strain curve for a new metamaterial structure, while Equation (15) can be used to construct a confidence band for quantifying the uncertainty of this prediction.

3.3 Prior specification

Finally, we provide a prior specification for the model parameters $(\theta(\cdot), \boldsymbol{\Sigma}, \boldsymbol{\beta})$. Consider independent priors on each parameter in $(\theta(\cdot), \boldsymbol{\Sigma}, \boldsymbol{\beta})$. For the weight function $\theta(\cdot)$, we assign independent exponential priors at each frequency ξ , i.e.:

$$\theta(\xi) \stackrel{i.i.d.}{\sim} \text{Exp}(\lambda_I), \quad (16)$$

where λ_I is a rate parameter for the exponential priors. Similar to the Bayesian LASSO (Park and Casella, 2008), the intuition behind the shrinkage prior (16) is to encourage *sparsity* in $\theta(\cdot)$, by assigning large prior mass to regions where $\theta(\xi) = 0$. This sparsity is desired for two reasons. First, this allows us to identify dominant frequencies in metamaterial structure which influence mechanical response. Second, sparsity in $\theta(\cdot)$ greatly speeds up the tissue-mimicking procedure using the proposed emulator: it screens out inert frequencies from the optimization model, which then reduces the number of evaluations from the emulator. This speed-up is paramount for efficient tissue-mimicking in urgent surgical applications (see Section 5.4).

For the covariance matrix Σ , we assign the following prior:

$$\pi(\Sigma) \propto \exp(-\lambda_o \|\Sigma^{-1}\|_1). \quad (17)$$

Here, λ_o is a rate parameter, and $\|\cdot\|_1$ is the element-wise l_1 norm. The prior (17) on Σ can be viewed as a shrinkage prior which encourages sparsity on the elements of the inverse covariance matrix Σ^{-1} (Wang, 2012). This corresponds to the widely-used graphical LASSO (Friedman et al., 2008) method for sparse covariance estimation. For our problem, this sparsity can be used to identify important and interpretable physical couplings in the stress-strain relationship (see Section 5.3.1).

For the regression coefficients β , we assign a non-informative flat prior $\pi(\beta) \propto 1$, since little information is known on β prior to data in our problem. A more informative prior can be used on β , if additional domain knowledge is available on the mean trend of the stress-strain curve.

4 Parameter estimation

In implementation, the functional inputs $I(\cdot)$ are also discretized to p levels. Let $\mathbf{x}_1 = (x_1^k)_{k=0}^{p-1}$ and $\mathbf{x}_2 = (x_2^k)_{k=0}^{p-1}$ denote the discretized input vectors for both $I_1(\cdot)$ and $I_2(\cdot)$, respectively.

The proposed SpeD kernel $\rho(I_1(\cdot), I_2(\cdot))$ in (6) can be approximated as:

$$\rho(\mathbf{x}_1, \mathbf{x}_2) = \exp \left(- \sum_{k=0}^{(p-1)/2} \theta_k (|\hat{x}_1^k| - |\hat{x}_2^k|)^2 \right), \quad (18)$$

where $\boldsymbol{\theta} = (\theta_k)_{k=0}^{(p-1)/2}$ is the discretized weight vector, and $\hat{x}^k = \sum_{l=0}^{p-1} x^l e^{-\frac{2\pi i}{p}lk}$ is the k -th entry of the *discrete* Fourier transform $\hat{\mathbf{x}}$ for the discretized input $\mathbf{x}$. Note that $\hat{\mathbf{x}}$ is symmetric because $\mathbf{x}$ is real-valued (Sorensen et al., 1987); hence, only the first half of $\hat{\mathbf{x}}$ is used in (18).

With this input discretization, we adopt a maximum a posteriori (MAP) approach for estimating the parameters $(\boldsymbol{\beta}, \boldsymbol{\theta}, \boldsymbol{\Sigma})$. The main reason we prefer MAP over a fully Bayesian approach is computational efficiency, for both parameter estimation and tissue-mimicking optimization. For parameter estimation, a fully Bayesian approach typically requires Markov chain Monte Carlo sampling (MCMC; Gelman et al., 1995). Given the complexities of functional inputs and outputs, MCMC sampling can take several days, which is more time-consuming than a single computer experiment run! Furthermore, the primary application of the proposed emulator is for tissue-mimicking optimization, which typically requires *many* evaluations of the emulation predictor. Therefore, it can be *very* time-consuming in a fully Bayesian implementation, since *each* evaluation involves an average over all MCMC samples. Given the urgent nature of surgical planning, the MAP approach (described next) offers a quicker way to survey the metamaterial design space, which enables timely tissue-mimicking optimization.

From the GP model in (10) and (13), the MAP estimation of $(\boldsymbol{\beta}, \boldsymbol{\theta}, \boldsymbol{\Sigma})$ boils down to minimizing the following penalized negative log-posterior (Santner et al., 2013):

$$\begin{aligned} \min_{\boldsymbol{\beta}, \boldsymbol{\theta} \geq 0, \boldsymbol{\Sigma} \succeq 0} l_{\lambda}(\boldsymbol{\beta}, \boldsymbol{\theta}, \boldsymbol{\Sigma}) = & \min_{\boldsymbol{\beta}, \boldsymbol{\theta} \geq 0, \boldsymbol{\Sigma} \succeq 0} \left[n \log \det \boldsymbol{\Sigma} + m \log \det \mathbf{R}_{\boldsymbol{\theta}} + \lambda_I \|\boldsymbol{\theta}\|_1 + \lambda_o \|\boldsymbol{\Sigma}^{-1}\|_1 \right. \\ & \left. + (\mathbf{y}_{1:n} - \mathbf{1}_n \otimes \mathbf{P}\boldsymbol{\beta})^T (\mathbf{R}_{\boldsymbol{\theta}}^{-1} \otimes \boldsymbol{\Sigma}^{-1}) (\mathbf{y}_{1:n} - \mathbf{1}_n \otimes \mathbf{P}\boldsymbol{\beta}) \right]. \end{aligned} \quad (19)$$

Algorithm 1 BCD algorithm for minimizing the penalized negative log-likelihood (19)

```

1: • Set initial values  $\beta \leftarrow \mathbf{0}_q$ ,  $\Sigma \leftarrow \mathbf{I}_m$  and  $\theta \leftarrow \mathbf{1}_p$ , and set  $\mathbf{Y} \leftarrow [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n]^T$ 
2: repeat
3:   Optimizing  $\Sigma$ :
4:   • Set  $\mathbf{R}_\theta = \left[ \exp \left( - \sum_{k=0}^{(p-1)/2} \theta_q \left( |\hat{x}_i^k| - |\hat{x}_j^k| \right)^2 \right) \right]_{i=1, j=1}^{n \quad n}$  with  $\hat{x}^k = \sum_{l=0}^{p-1} x^l e^{-\frac{2\pi i}{p} l k}$ 
5:   • Set  $\mu = \mathbf{P}\beta$ 
6:   • Set  $\mathbf{W}_0 \leftarrow \frac{1}{n} (\mathbf{Y} - \mathbf{1}_n \otimes \mu^T)^T \mathbf{R}_\theta^{-1} (\mathbf{Y} - \mathbf{1}_n \otimes \mu^T) + \lambda_o \cdot \mathbf{I}_m$ 
7:   • Estimate  $\mathbf{W}$  by Graphical LASSO using  $\mathbf{W}_0$  as initialization
8:   • Update  $\Sigma \leftarrow \mathbf{W}^{-1}$ 
9:   Optimizing  $\beta$ :
10:  • Set  $\mathbf{S} = (\mathbf{P} \otimes \mathbf{1}_n)^T (\mathbf{R}_\theta^{-1} \otimes \mathbf{W}) (\mathbf{P} \otimes \mathbf{1}_n)$ 
11:  • Update  $\beta \leftarrow \mathbf{S}^{-1} (\mathbf{P} \otimes \mathbf{1}_n)^T (\mathbf{R}_\theta^{-1} \otimes \mathbf{W}) \mathbf{y}_{1:n}$ 
12:  Optimizing  $\theta$ :
13:  • Update  $\theta \leftarrow \arg\min_{\theta} l_{\lambda}(\beta, \Sigma, \theta)$  with L-BFGS
14: until  $\beta$ ,  $\Sigma$  and  $\theta$  converge
15: • return  $\beta$ ,  $\Sigma$  and  $\theta$ 

```

Here, $\mathbf{R}_\theta$ is the correlation matrix in (13) with scale parameters θ , and λ_I and λ_o are the rate parameters for the shrinkage priors in Section 3.3.

From a regularization perspective, the two prior terms $\lambda_I \|\theta\|_1$ and $\lambda_o \|\Sigma^{-1}\|_1$ in the negative log-posterior (19) can equivalently be viewed as penalty terms on θ and Σ^{-1} , with the rate parameters λ_I and λ_o corresponding to penalization parameters. In this sense, the parameters λ_I and λ_o control the degree of sparsity imposed on θ and Σ^{-1} , with a larger λ_I (or λ_o) resulting in a sparser estimate of θ (or Σ^{-1}), and vice versa. In practice, these penalization parameters can be estimated from the data itself, or specified from the problem at hand. For example, if predictive accuracy of the emulator is the end goal, then λ_I and λ_o can be estimated based on cross-validation techniques (Friedman et al., 2001). However, if the extraction of important physics is desired, then λ_I and λ_o can be set so that a desired number of physical features can be learned. We will return to this in Section 5.3.

Consider now the MAP optimization in (19) for fixed $\lambda_I > 0$ and $\lambda_o > 0$. We will use the following blockwise coordinate descent (BCD) optimization algorithm, described below. First, assign initial values for β , θ and Σ . Next, iterate the following three steps until

the convergence is achieved: (i) for fixed GP parameters $\boldsymbol{\theta}$ and regression coefficients $\boldsymbol{\beta}$, compute the correlation matrix $\mathbf{R}_{\boldsymbol{\theta}}$ and then optimize for covariance matrix $\boldsymbol{\Sigma}$ using the graphical LASSO algorithm (Friedman et al., 2008); (ii) for fixed $\boldsymbol{\theta}$ and $\boldsymbol{\Sigma}$, compute $\boldsymbol{\beta}$ using closed-form expressions (see Santner et al., 2013 for details); and (iii) for fixed $\boldsymbol{\beta}$ and $\boldsymbol{\Sigma}$, optimize for $\boldsymbol{\theta}$ using the L-BFGS algorithm (Liu and Nocedal, 1989). The full optimization procedure is provided in Algorithm 1. Since (19) is a non-convex optimization problem, the proposed BCD algorithm only converges to a stationary solution (Tseng, 2001). Because of this, we suggest performing multiple runs of Algorithm 1 with random initializations for each run, then taking the converged estimates for the run with smallest negative log-likelihood. These runs should be performed in parallel if possible, to take advantage of the parallel computing capabilities in many computing systems. ‘

5 Emulation results

In this section, we present the numerical performance of the proposed model for tissue-mimicking. This is presented in four parts. First, we compare the predictive performance of the proposed SpeD emulation model with two baseline emulation models. Second, we provide a comparison of the uncertainty quantification from these three emulation models. Third, we analyze the physical properties learned via shrinkage priors on $\boldsymbol{\theta}$ and $\boldsymbol{\Sigma}$. Finally, we demonstrate the usefulness of the fitted model for mimicking human aortic tissue.

5.1 Prediction accuracy

As mentioned in Section 2.2, the proposed SpeD model is fitted using the training data of $n = 58$ FE simulations. The input function $I(\cdot) \in \mathcal{I}$ is discretized to $p = 81$ parts, which we denote as a vector $\mathbf{x} \in \mathbb{R}^{81}$. In this specific tissue-mimicking problem, the diameter of the metamaterial enhancement $d \in \mathbb{R}$ (assumed to be uniform over the whole functional curve $I(\cdot)$) is also important. To account for this extra design variable, we use the following

separable correlation function $\rho_s(\cdot, \cdot) : \mathbb{R}^{82} \times \mathbb{R}^{82} \mapsto \mathbb{R}$:

$$\rho_s([d_1, \mathbf{x}_1], [d_2, \mathbf{x}_2]) = \rho(\mathbf{x}_1, \mathbf{x}_2) \exp(-\theta_d (d_1 - d_2)^2), \quad (20)$$

where $\rho(\cdot, \cdot)$ is the discretized SpeD kernel in (18) and θ_d is the scale parameter for diameter d . Let $\mathbf{s} \in \mathbb{R}^{m=41}$ denote the vector of strain levels equally spaced from $s = 0$ to $s = 15\%$, and let $\mathbf{y} = O(\mathbf{s}) \in \mathbb{R}^{41}$ be the discretized stress function $O(\cdot)$. Note that the discretization levels for input and output are selected due to heuristic observation – the output functions are quite smooth, and the input functions are slightly more rugged.

From the underlying physics of the stress-strain relationship, it is known that (i) the stress $O(s)$ is always positive, (ii) the stress is zero when the strain is zero (this is known as the free-standing state, see Malvern, 1969), and (iii) stress-strain curves are typically monotone and non-decreasing, since a larger force is needed to stretch further. To account for (i), a standard log-transformation of stress $O(s)$ is performed prior to modeling and parameter estimation, and the final results are transformed back to ensure the predicted stress is always positive. To account for (ii) and (iii), we choose the basis functions in (12) to be $\mathbf{P} = [\mathbf{1}_m, \log(\mathbf{s})]$, along with an additional constraint of $\beta_2 > 0$ to ensure the mean function is monotone and non-decreasing. This is equivalent to assuming the following form for the mean stress-strain curve $O(s) = as^b$, $a, b > 0$, which is a typical parametrization in biomedical literature (Rengier et al., 2010; Chen et al., 2018a). Another option is to adapt a monotone GP (Wang and Berger, 2016), but we choose not to use it here for the sake of simplicity.

For comparison, we also fit two different emulators as baseline methods, using the same dataset. The inputs of the first emulator are the parameters from the sinusoidal wave design $\mathbf{x}_p = [d, A, \omega, \phi]^T \in \mathbb{R}^4$, which represents the diameter of the metamaterial fiber, amplitude, period and initial phase of the sinusoidal wave (see Figure 1 and Equation (1)).




Figure 4: Predicted stress-strain curves for the l_2 -distance emulator (l_2), feature-based emulator (feature), and the proposed SpeD emulator (SpeD) on two test metamaterial structures. The corresponding MAREs are included in the legends.

This emulator uses a GP model with correlation function:

$$\rho_p(\mathbf{x}_1, \mathbf{x}_2) = \exp \left(- \sum_{k=1}^4 \theta_k (x_1^k - x_2^k)^2 \right). \quad (21)$$

The same correlation function (with scalar output) is used in Chen et al. (2018a). We refer this as the *feature-based* method. The second emulator also assumes a GP model with correlation function:

$$\rho_f(I_1(\cdot), I_2(\cdot)) = \exp \left(- \int \theta(t) (I_1(t) - I_2(t))^2 dt \right). \quad (22)$$

This correlation (22) is essentially the Gaussian correlation function, with distance taken to be the l_2 -distance between input functions. A similar correlation function is used in Morris (2012) for time-series inputs, with additional dependencies on time order. We refer this as the functional l_2 -distance method. Both baseline methods assume the same separable co-kriging structure for discretized outputs $\mathbf{y}_{1:n}$, along with MAP parameter estimation.

5.1.1 Predicting stress-strain curve

To test the performance of the proposed emulator, we compare the predictions of stress-strain curves (using Equation (14)) for the metamaterial designs from the test set (see Section 2.2).




Figure 5: Boxplots of the MARE ratio between the baseline emulators and SpeD emulator on the 18-run test set. The red line marks the MARE ratio of 1.0, indicating the same MARE of the baseline emulator and the SpeD emulator.

	Ave. MARE
SpeD	0.20
Feature-based	0.28
l_2 -distance	0.27

Table 2: The averaged MARE of the SpeD emulator and two baseline emulators over the 18-run test set.

Figure 4 shows the emulated stress-strain curves for two test metamaterial structures, along with the true stress-strain curve (ground truth) from FE simulations. To quantitatively measure the difference between the predicted and true curves, we use the following mean absolute relative error (MARE) metric:

$$\text{MARE} = \frac{\int_s |O(s) - \hat{O}(s)| ds}{\int_s |O(s)| ds}, \quad (23)$$

where s is the strain level, $O(s)$ is the stress at strain s from FE simulation (ground truth), and $\hat{O}(s)$ is the predicted stress from the emulators. The MARE values for the two test cases in Figure 4 are included in the legends. Table 2 reports the averaged MARE values for the three considered emulators, over the whole test set. The proposed SpeD emulator appears to perform very well, in that it achieves noticeably lower averaged MARE than the two existing emulators. Figure 5 shows the boxplots of the MARE ratio between the baseline emulators and the SpeD emulator for the 18 test cases (note that a ratio of 1.0 means the SpeD model yields similar MARE to a baseline model). We see that these ratios are mostly larger than one, which suggests that the proposed emulator is noticeably better in predicting the true stress-strain output curve. This is not surprising, since our model captures known physical properties of the tissue-mimicking problem (see Table 1).




Figure 6: (a) visualizes the three characteristics of mechanical performance: moduli E_1 and E_9 , and curvature κ . (b) and (c) show the pairwise absolute relative error for E_1 and E_9 between the two baseline emulators and the SpeD emulator. The red line marks a relative error of 1.0.

5.1.2 Predicting physical characteristics

In addition to predicting stress-strain curve $O(s)$ at untested metamaterial structures, engineers are also interested in predicting key physical characteristics. An accurate prediction of these physical characteristics can be as important as emulating the stress-strain curve itself, because it provides interpretability to the black-box emulation model. Two important physical characteristics to predict here are (i) the elastic modulus of the stress-strain curve, and (ii) the classification of material type as strain-stiffening or strain-softening. For (i), the *modulus*, i.e., the slope of the stress-strain curve at different strain levels, can be interpreted as the stiffness or hardness of the material (Raghavan et al., 1996). Here, we are interested in the elastic moduli E_1 and E_9 at strain levels 1% and 9%, respectively, where $E_k = \partial O(s) / \partial s|_{s=k\%}$; this allows us to evaluate the elastic moduli prediction over a wide range of strain levels. For (ii), we wish to classify the stress-strain curve as *strain-stiffening* or *strain-softening*; this is particularly important given the goal of mimicking biological tissues (see Section 2.1). One way to classify is to use the curvature of the stress-strain curve, which can be approximated by the slope of the two moduli, $\kappa = \partial^2 O / \partial s^2 \approx (E_9 - E_1) / (9\% - 1\%)$. Assuming no fluctuations in the strain region $s \in [1, 9]\%$ (Malvern, 1969), a positive curvature κ suggests a strain-stiffening property is present (due to increasing moduli), while a negative κ suggests a strain-softening property is present. Figure 6 (a) visualizes these physical characteristics from a stress-strain curve.

	SpeD	Feature-based	l_2 -distance
<i>True positive %</i>	12/12=100%	11/12=91.7%	7/12=58.3%
<i>True negative %</i>	6/6=100%	5/6=83.3%	5/6=83.3%
<i>Classification %</i>	18/18=100%	16/18=88.9%	13/18=72.2%

Table 3: *The true positive rate, true negative rate, and classification rate of strain-stiffening and strain-softening, for the three considered emulators.*

We now compare three emulators (SpeD and baselines) for predicting the moduli and material type. The moduli $\hat{E}_1$ and $\hat{E}_9$, computed from the emulated stress-strain curves, are compared with the moduli E_1 and E_9 from FE simulation. Figures 6 (b) and (c) show the pairwise absolute relative error $|\hat{E} - E|/E$, between the baselines and the SpeD emulator. We see that most of these ratios are larger than 1.0 in the test set, which shows that the proposed SpeD model outperforms both baseline emulators. For classification, the predicted curvature $\hat{\kappa}$, computed from the emulated curves, are compared with the true curvature κ from FE simulation. Table 3 shows the correct classification rates for the three emulators. The SpeD model has a perfect $18/18 = 100\%$ classification accuracy: it identified the correct strain-softening/-stiffening property for all 18 test structures. On the other hand, the feature-based model and the l_2 -distance model achieves only a $16/18 = 88.9\%$ and $13/18 = 72.2\%$ classification accuracy rate, respectively. One reason why the proposed SpeD model can better capture these physical characteristics (compared to existing emulators) is because it directly incorporates the underlying physics via the SpeD correlation function.

5.2 Uncertainty quantification

Particularly in healthcare applications, the quantification of predictive uncertainty can be as important as the prediction itself. For the proposed model, Equations (14) and (15) can be used to construct 90% pointwise highest posterior density predictive intervals (HPD-PIs) for the emulated stress-strain curves. Figure 7 shows the 90% HPD-PI for the three emulation models. Note that there is little predictive uncertainty at zero strain, with




Figure 7: A comparison of the 90% pointwise HPD-PIs for the three emulation models (left: SpeD, middle: feature-based, right: l_2 -distance). Different rows are for different test cases.

uncertainty increasing as strain levels increase. This is consistent with the physical intuition in Section 5.1: the stress always equals to zero when strain equals zero, i.e., no force at free-standing condition. The increasing uncertainty for higher strain levels may be due to the log-transformation of the functional output.

Comparing the predictive intervals for the three emulators, we see that the proposed SpeD model returns narrower HPD-PIs compared to both the l_2 -distance model and the feature-based model. This is particular evident for the test case in the top row of Figure 7. Moreover, the 90% HPD-PIs of the SpeD emulator covers the true stress-strain curves in 16/18 of the test cases, whereas the coverage for the feature-based and l_2 -distance emulators are only 12/18 and 14/18, respectively. For example, the bottom row of Figure 7 shows a test case where feature-based emulator fails to cover the true stress-strain curve. Over the whole test set, our SpeD emulator appears to give reliable coverage of the true stress-strain curve, with relatively low predictive uncertainty. The reasons for this may be two-fold: (i) the SpeD correlation directly captures the physics of the tissue-mimicking problem, which can be thought of as an additional source of data, (ii) the shrinkage priors on spectral




Figure 8: (a) The sparsity pattern visualization of the inverse covariance matrix Σ^{-1} by graphical LASSO with 40% of non-zero entries. The pattern indicates three different regions of the stress-strain curve, colored yellow, green and red. (b) The partition of strain-stress curves for soft materials into toe, elastic, and yield regions up to strain level of 15%.

coefficients screens out inert frequencies, which also helps reduce predictive uncertainty.

5.3 Learning physics via sparsity

In addition to prediction accuracy and uncertainty quantification, the SpeD emulator also provides a data-driven approach to learn important physics. This is achieved via the shrinkage priors on both the covariance matrix Σ and frequency coefficients θ .

5.3.1 Segmentation of stress-strain curve

We first analyze the important correlations selected by the shrinkage prior on the co-kriging covariance matrix Σ . Recall from Section 3.3 that the graphical LASSO is used to extract significant correlations in Σ^{-1} . Setting the penalty parameter λ_o such that 40% of the entries are non-zero, Figure 8 (a) visualizes the selected (important) covariances in Σ^{-1} . Recall that the (i, j) -th entry of Σ^{-1} represents the covariance between variables i and j , conditional on all other variables. We see that the stress-strain curve can be roughly segmented into three regions: small strain (from 0% to 3%) with high conditional correlation, medium strain (from 3% to 9%) with moderate conditional correlation and large strain

(from 9% to 15%) with high conditional correlation.

These three regions nicely corroborate known physical properties in material strength (Malvern, 1969; Martin et al., 1998), where the mechanical response of the soft bio-mimicking material can also be divided to three regions: the toe region, the elastic region and the yield region (see Figure 8 (b)). We see from Figure 8 (a) that there are fewer significant conditional correlations in the elastic region compared to the other two regions. One reason for this is that, within the elastic region, the stress-strain curve can be better approximated by the form $O(s) = as^b$ (which corresponds to the choice of basis functions in $\mathbf{P}$). Figure 8 (a) also suggests the presence of conditional correlations between the elastic and yield regions. Physically, this can be explained by the migration of both strain-stiffening and strain-softening properties from metamaterial straightening.

5.3.2 Learning dominant frequencies

The proposed approach can also learn important frequencies ξ which influence mechanical response, via the shrinkage priors on the weight function $\theta(\xi)$ (see Section 3.3). Figure 9 (c) shows the MAP estimate of $\theta(\cdot)$ in the spectral space, where the rate parameter λ_I is chosen via cross-validation. We see that $\theta(\cdot)$ shrinks to zero at low and high frequencies, with non-zero estimates only for medium frequencies between $50m^{-1}$ to $400m^{-1}$. For these two endpoint frequencies, Figures 9 (a) and (b) show the metamaterial structures with frequencies $\xi \approx 50m^{-1}$ and $\xi \approx 400m^{-1}$, respectively.

The selected frequencies in $\theta(\xi)$ are also in line with the physical understanding of the problem. For low frequencies (Figure 9 (a)), the fluctuation in metamaterial design is too weak to induce any effect on the stress-strain curve, whereas for high frequencies (Figure 9 (b)), the resulting strong fluctuation in metamaterial leads to mechanical properties similar to a straight fiber (given nonzero diameter d , see Chen et al., 2018a). While it is known that different frequencies affect mechanical response in different ways, a strict law is difficult to find for engineers. Here, our SpeD emulator sheds light on the influential frequencies, i.e., from $50m^{-1}$ to $400m^{-1}$; those frequencies should be carefully chosen for metamaterial




Figure 9: Examples of metamaterial structure with very low (a) or very high (b) frequency. (c) Sparse estimation of spectral parameters θ , where medium frequencies are non-zero.

design. We show next how this identification of important frequencies allows us to greatly speed up optimization for the target task of mimicking human tissue.

5.4 Mimicking aortic tissue via optimization

We now explore the motivating task of mimicking the mechanical properties of a target tissue with the proposed emulator. Recall that the key challenge in applying 3D-printing techniques to medical application is the discrepancy in mechanical properties between target tissues and printed prototypes. Here, the SpeD model can be used to find a good metamaterial design (both structure $I(\cdot)$ and diameter d) whose stress-strain curve matches the desired mechanical property $\mathbf{y}^*$. This is achieved via the following optimization problem:

$$(d^*, I^*(\cdot)) = \underset{d_{\text{new}}, I_{\text{new}}(\cdot) \in \mathcal{I}}{\operatorname{argmin}} \mathbb{E} \left\{ \|\mathbf{y}([d_{\text{new}}, I_{\text{new}}(\cdot)]) - \mathbf{y}^*\|_2^2 | \mathbf{y}_{1:n} \right\}. \quad (24)$$

Here, $I^*(\cdot)$ is the optimal metamaterial structure, d^* is the optimal fiber diameter, and $\mathbf{y}([d_{\text{new}}, I_{\text{new}}(\cdot)]) | \mathbf{y}_{1:n}$ is the conditional (discretized) stress-strain curve in (13) with diameter d_{new} and structure $I_{\text{new}}(\cdot)$. In words, equation (24) aims to find the optimal metamaterial design whose stress-strain curve from the proposed model (conditional on data) is closest to the target curve $\mathbf{y}^*$ in terms of mean-squared error (MSE).

This MSE criterion can be further decomposed as follows:

$$\|\hat{\mathbf{y}}([d_{\text{new}}, I_{\text{new}}(\cdot)]) - \mathbf{y}^*\|_2^2 + \text{tr}(\text{Var}\{\mathbf{y}([d_{\text{new}}, I_{\text{new}}(\cdot)])|\mathbf{y}_{1:n}\}). \quad (25)$$

Here, $\hat{\mathbf{y}}([d_{\text{new}}, I_{\text{new}}(\cdot)])$ and $\text{Var}\{\mathbf{y}([d_{\text{new}}, I_{\text{new}}(\cdot)])|\mathbf{y}_{1:n}\}$ are the conditional mean and variance of $\mathbf{y}([d_{\text{new}}, I_{\text{new}}(\cdot)])|\mathbf{y}_{1:n}$, respectively, and $\text{tr}(\mathbf{A}) = \sum_i A_{i,i}$ is the trace of the matrix $\mathbf{A}$. The first term can be interpreted as trying to minimize the l_2 -norm between the emulated stress-strain curve and the target curve. The second term can be viewed as trying to minimize the predictive variance of the emulated curve. Such a decomposition is quite intuitive, since we wish to find a metamaterial design whose emulated curve matches the desired curve, but also has low predictive uncertainty from the emulation model.

One difficulty in solving (24) is that the variable $I(\cdot)$ is *functional* in form, and its discretization $\mathbf{x} \in \mathbb{R}^p$, $p = 81$ can be too *high dimensional* to optimize numerically. Here is where the extracted important frequencies from Section 5.3.2 come into play. Let $\hat{\mathbf{x}}_{\text{new}} \in \mathbb{R}^7$ denote the seven non-zero Fourier coefficients (see Figure 9). Using these coefficients as inputs for optimization (and ignoring the other inert coefficients), we get the following lower-dimensional optimization problem:

$$(d^*, \hat{\mathbf{x}}^*) = \underset{d_{\text{new}} \in \mathbb{R}, \hat{\mathbf{x}}_{\text{new}} \in \mathbb{R}^7}{\text{argmin}} \mathbb{E}\{\|\mathbf{y}([d_{\text{new}}, \hat{\mathbf{x}}_{\text{new}}]) - \mathbf{y}^*\|_2^2 | \mathbf{y}_{1:n}\}, \quad (26)$$

where $\mathbf{y}([d_{\text{new}}, \hat{\mathbf{x}}_{\text{new}}])|\mathbf{y}_{1:n}$ is the conditional random vector taking the frequencies as input. While this problem is non-convex, it is much *lower-dimensional*, and can be effectively optimized using standard quasi-Newton methods (e.g., L-BFGS) and random initializations.

This framework (using the proposed SpeD model) offers significant speeds up for tissue-mimicking over the current state-of-the-art methods. To see why, consider first the optimization of (24) using only numerical FE simulations: this requires hundreds of evaluations of the optimization objective function, each of which requires around 30 minutes of computation time. This means tissue-mimicking with only FE simulations can require




Figure 10: A case study which uses the proposed SpeD model for mimicking human aortic tissue. (a) shows the stress-strain curves of the target aortic tissue (black), the mimicked curve from an existing method (blue) and the curve optimized from the SpeD method (red). (b) shows the optimal metamaterial design from our approach.

many days of computation, which is clearly unsuitable for urgent surgical planning (Chen et al., 2018a). To contrast, each evaluation of the proposed emulator requires only seconds of computation, which greatly speeds up the mimicking process. Furthermore, by exploiting sparsity in spectral coefficients, the dimension of the optimization problem reduces from 82 to 8 variables. This dimension reduction greatly cuts down on the number of predictions from the emulator, which yields significant reductions in computation time. Such speed-ups are paramount for performing tissue-mimicking in an accurate and timely manner. Section 5.5 provides a further comparison of timing.

Let us now explore the tissue-mimicking performance of the SpeD emulator in a real-world case study. Figure 10 (a) shows the stress-strain curve of a target aortic tissue (in black) from Hockaday et al. (2012), the stress-strain curve from the proposed mimicking procedure (in red), and the curve from an existing mimicking method (in blue) in Wang et al. (2016b). The latter method performs mimicking using only the four sinusoidal metamaterial features (see Section 2.2). Compared to the existing approach, which has an MARE (see Equation 23) of 0.528, the proposed SpeD approach achieves a much smaller MARE of 0.089. This improved tissue-mimicking performance can be seen in Figure 10 (a): the red curve (from the proposed method) closely mimics the desired black curve, whereas the blue

Modeling Step	Computation Time (in minutes)
Parameter estimation	21.72
Prediction & UQ at one input setting	0.01
Tissue-mimicking of a target material	2.69

Table 4: *Computation time for different modeling steps of the proposed emulator, parallelized over 24 processing cores.*

curve (from Wang et al., 2016b) overestimates stress at all strain levels. In particular, our method gives much better mimicking in small strain regions – this is important in medical applications due to the relatively small strain deformations in the human body.

There are two reasons for this improved performance. First, the existing mimicking approach in Wang et al. (2016b) is too restrictive, in that it uses only four sinusoidal features and not the full functional form of the input. Second, given a fixed timeframe, the proposed emulation-based approach permits a larger number of objective evaluations via the proposed SpeD model. Figure 10 (b) shows the optimal (discretized) metamaterial design $\mathbf{x}^*$ from our emulation-based approach, which is clearly not a sinusoidal function. By considering the broader class of functional inputs as well as allowing for more objective evaluations, the proposed method can identify better metamaterial designs for tissue-mimicking. While there is some discrepancy between the proposed stress-strain curve and the target curve, this may be because of the over-simplified optimization scheme and fact that the raw printable polymers are not strain-stiffening. Improving this will be the focus of future work.

5.5 Computation time

Another appeal of the SpeD emulator is its computational efficiency. Table 4 summarizes the computation time required for each step of the emulation process, with timing performed on a parallelized system of 24 Intel Xeon E5-2650 2.20GHz processing cores. The computation time required for parameter estimation (with cross-validation on λ_I) is 21.72 minutes, which is typically performed before the arrival of the patient. Once we fit the model, we can predict

for multiple settings very quickly (0.01 minutes per setting). To contrast, FE simulations requires 30 minutes for each structure, and the fully Bayesian counterpart, due to the large amount (e.g., 2000) of MCMC samples, requires $0.01 \times 2000 = 20$ minutes. Because of this, our SpeD emulator can effectively perform the tissue-mimicking procedure using only 3 minutes of computation; this greatly improves upon the standard tissue-mimicking approach with only FE simulations, which may require hours or even days to perform (Wang et al., 2016b) with much poorer mimicking performance (see Figure 10)!

Given the positive results reported in this section, the SpeD emulator can serve as an important tool for furthering the use of 3D printed organs in medical applications. Our model can provide effective and personalized pre-surgical practicing and planning (Qian et al., 2017; Chen et al., 2019) with dramatically lower costs, which then mitigates risk in complex surgical procedures.

6 Conclusion

We propose in this paper a novel function-on-function Gaussian process emulation model for tackling the challenging tissue-mimicking optimization, under urgent surgical demands. The key challenge is in building an efficient and accurate emulator, which incorporates the *functional* complexities of the input metamaterial structures and the *functional* output mechanical responses. To address this, the proposed co-kriging model uses a new spectral-distance (SpeD) correlation function, which integrates spectral information by directly modeling the effect of metamaterial frequencies on mechanical response. One appealing feature of this new correlation function is its *translation-invariance* property, which accounts for the fact that two metamaterial structures, which are equivalent modulo a translation shift, have the same mechanical properties. For parameter estimation, we use MAP with shrinkage priors, which identifies key frequencies and thereby reduces the large *functional* input space. This reduction greatly speeds up the tissue-mimicking optimization using the proposed emulator. Applied to a real-world tissue-mimicking study, the proposed

SpeD emulator outperforms existing models in (i) emulating and quantifying uncertainty on mechanical response, (ii) extracting meaningful physical insights, and (iii) providing efficient and accurate mimicking performance for human aortic tissue. With the development of multi-material 3D-printing technology, this new emulator can play an important role in furthering the impact of 3D-printing in important biomedical applications in surgery planning and healthcare.

References

- Banerjee, S., Carlin, B. P., and Gelfand, A. E. (2014). *Hierarchical Modeling and Analysis for Spatial Data*. CRC Press.
- Bayarri, M., Berger, J., Cafeo, J., Garcia-Donato, G., Liu, F., Palomo, J., Parthasarathy, R., Paulo, R., Sacks, J., and Walsh, D. (2007). Computer model validation with functional output. *The Annals of Statistics*, 35(5):1874–1906.
- Bracewell, R. N. and Bracewell, R. N. (1986). *The Fourier Transform and its Applications*, volume 31999. McGraw-Hill New York.
- Callister, W. D. and Rethwisch, D. G. (2011). *Materials Science and Engineering*, volume 5. John Wiley & Sons NY.
- Chen, J., Wang, K., Zhang, C., and Wang, B. (2018a). An efficient statistical approach to design 3D-printed metamaterials for mimicking mechanical properties of soft biological tissues. *Additive Manufacturing*, 24:341–352.
- Chen, J., Xie, Y., Wang, K., Wang, Z. H., Lahoti, G., Zhang, C., Vannan, M. A., Wang, B., and Qian, Z. (2018b). Generative invertible networks (GIN): Pathophysiology-interpretable feature mapping and virtual patient generation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 537–545. Springer.
- Chen, J., Xie, Y., Wang, K., Zhang, C., Vannan, M. A., Wang, B., and Qian, Z. (2019). Active image synthesis for efficient labeling. *arXiv preprint arXiv:1902.01522*.
- Diest, K. (2013). *Numerical Methods for Metamaterial Design*. Springer.
- Fan, J. and Zhang, W. (2008). Statistical methods with varying coefficient models. *Statistics and*

- its Interface*, 1(1):179.
- Friedman, J., Hastie, T., and Tibshirani, R. (2001). *The Elements of Statistical Learning*, volume 1. Springer Series in Statistics, NY.
- Friedman, J., Hastie, T., and Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441.
- Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (1995). *Bayesian Data Analysis*. Chapman and Hall/CRC.
- Higdon, D., Gattiker, J., Williams, B., and Rightley, M. (2008). Computer model calibration using high-dimensional output. *Journal of the American Statistical Association*, 103(482):570–583.
- Hill, R. (1998). *The Mathematical Theory of Plasticity*, volume 11. Oxford University Press.
- Hockaday, L., Kang, K., Colangelo, N., Cheung, P., Duan, B., Malone, E., Wu, J., Girardi, L., Bonassar, L., Lipson, H., Chu, C., and Butcher, J. (2012). Rapid 3D printing of anatomically accurate and mechanically heterogeneous aortic valve hydrogel scaffolds. *Biofabrication*, 4(3):035005.
- Joseph, V. R., Gul, E., and Ba, S. (2015). Maximum projection designs for computer experiments. *Biometrika*, 102(2):371–380.
- Liu, D. C. and Nocedal, J. (1989). On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45(1-3):503–528.
- Mak, S., Sung, C.-L., Wang, X., Yeh, S. T., Chang, Y. H., Joseph, V. R., Yang, V., and Wu, C. J. (2018). An efficient surrogate model for emulation and physics extraction of large eddy simulations. *Journal of the American Statistical Association*, pages 1–14.
- Malfait, N. and Ramsay, J. O. (2003). The historical functional linear model. *Canadian Journal of Statistics*, 31(2):115–128.
- Malvern, L. E. (1969). *Introduction to the Mechanics of a Continuous Medium*. Number Monograph.
- Martin, R. B., Burr, D. B., and Sharkey, N. A. (1998). *Skeletal Tissue Mechanics*, volume 190. Springer.
- Matheron, G. (1963). Principles of geostatistics. *Economic Geology*, 58(8):1246–1266.
- Morris, M. D. (2012). Gaussian surrogates for computer models with time-varying inputs and outputs. *Technometrics*, 54(1):42–50.
- Muehlenstaedt, T., Fruth, J., and Roustant, O. (2017). Computer experiments with functional

- inputs and scalar outputs by a norm-based approach. *Statistics and Computing*, 27(4):1083–1097.
- Park, T. and Casella, G. (2008). The Bayesian lasso. *Journal of the American Statistical Association*, 103(482):681–686.
- Qian, Z., Wang, K., Liu, S., Zhou, X., Rajagopal, V., Meduri, C., Kauten, J. R., Chang, Y. H., Wu, C., Zhang, C., Wang, B., and Vannan, M. A. (2017). Quantitative prediction of paravalvular leak in transcatheter aortic valve replacement based on tissue-mimicking 3d printing. *JACC: Cardiovascular Imaging*, 10(7):719–731.
- Raghavan, M. L., Webster, M. W., and Vorp, D. A. (1996). Ex vivo biomechanical behavior of abdominal aortic aneurysm: assessment using a new mathematical model. *Annals of Biomedical Engineering*, 24(5):573–582.
- Ramsay, J. (2005). *Functional Data Analysis*. Springer Series in Statistics, NY.
- Rengier, F., Mehndiratta, A., Von Tengg-Kobligh, H., Zechmann, C. M., Unterhinninghofen, R., Kauczor, H.-U., and Giesel, F. L. (2010). 3D printing based on imaging data: review of medical applications. *International Journal of Computer Assisted Radiology and Surgery*, 5(4):335–341.
- Santner, T. J., Williams, B. J., and Notz, W. I. (2013). *The Design and Analysis of Computer Experiments*. Springer Science & Business Media.
- Sobol’, I. M. (1967). On the distribution of points in a cube and the approximate evaluation of integrals. *Zhurnal Vychislitel’noi Matematiki i Matematicheskoi Fiziki*, 7(4):784–802.
- Sorensen, H. V., Jones, D., Heideman, M., and Burrus, C. (1987). Real-valued fast Fourier transform algorithms. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 35(6):849–863.
- Stein, A. and Corsten, L. (1991). Universal kriging and cokriging as a regression procedure. *Biometrics*, 47(2):575–587.
- Tseng, P. (2001). Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of Optimization Theory and Applications*, 109(3):475–494.
- Ventola, C. L. (2014). Medical applications for 3d printing: current and projected uses. *Pharmacy and Therapeutics*, 39(10):704.
- Wang, B. and Xu, A. (2019). Gaussian process methods for nonparametric functional regression with mixed predictors. *Computational Statistics & Data Analysis*, 131:80–90.
- Wang, H. (2012). Bayesian graphical lasso models and efficient posterior computation. *Bayesian Analysis*, 7(4):867–886.

- Wang, K., Wu, C., Qian, Z., Zhang, C., Wang, B., and Vannan, M. A. (2016a). Dual-material 3D printed metamaterials with tunable mechanical properties for patient-specific tissue-mimicking phantoms. *Additive Manufacturing*, 12:31–37.
- Wang, K., Zhao, Y., Chang, Y.-H., Qian, Z., Zhang, C., Wang, B., Vannan, M. A., and Wang, M.-J. (2016b). Controlling the mechanical behavior of dual-material 3D printed meta-materials for patient-specific tissue-mimicking phantoms. *Materials & Design*, 90:704–712.
- Wang, X. and Berger, J. O. (2016). Estimating shape constrained functions using gaussian processes. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1):1–25.
- Zienkiewicz, O. C., Taylor, R. L., Zienkiewicz, O. C., and Taylor, R. L. (1977). *The Finite Element Method*, volume 36. McGraw-Hill London.

Appendices

A Proof of Theorem 1

For any $n \in \mathbb{N}$, $c_1, \dots, c_n \in \mathbb{R}$ and function $I_1(\cdot), \dots, I_n(\cdot) \in \mathcal{I}$, we have

$$\sum_{i=1}^n \sum_{j=1}^n c_i c_j \rho(I_i(\cdot), I_j(\cdot)) \quad (27)$$

$$= \sum_{i=1}^n \sum_{j=1}^n c_i c_j \exp \left(- \int \theta(\xi) \left(\left| \hat{I}_i(\xi) \right| - \left| \hat{I}_j(\xi) \right| \right)^2 d\xi \right). \quad (28)$$

Note that the Fourier transform $\mathcal{F} : \mathcal{I} \mapsto \mathcal{I}$, where $\mathcal{I}$ is the space of integrable functions $I(\cdot) : \mathbb{R} \mapsto \mathbb{C}$, has a unique inverse $\mathcal{F}^{-1} : \mathcal{I} \mapsto \mathcal{I}$. Denote the standard Gaussian kernel as $K(\cdot, \cdot) : |\mathcal{I}| \times |\mathcal{I}| \mapsto \mathbb{R}$,

$$K(F_1(\cdot), F_2(\cdot)) = \exp \left(- \int \theta(t) (F_1(t) - F_2(t))^2 dt \right), \quad (29)$$

where $|\mathcal{I}|$ is the space of integrable functions $F(\cdot) : \mathbb{R} \mapsto \mathbb{R}$. Since $K(\cdot, \cdot)$ is a positive definite kernel, for the selected n and $c_1, \dots, c_n$ in Equation (27), and any function $F_1(\cdot), \dots, F_n(\cdot) \in \mathcal{I}$, we have,

$$\sum_{i=1}^n \sum_{j=1}^n c_i c_j K(F_i(\cdot), F_j(\cdot)) \geq 0. \quad (30)$$

Now let $F_k(\cdot) = \left| \hat{I}_k(\cdot) \right|$, where $k = 1, 2, \dots, n$. This is possible because $\hat{I}_k(\cdot) \in \mathcal{I}$ and therefore $\left| \hat{I}_k(\cdot) \right| \in \mathcal{I}$. Thus, we have

$$\sum_{i=1}^n \sum_{j=1}^n c_i c_j K \left(\left| \hat{I}_i(\cdot) \right|, \left| \hat{I}_j(\cdot) \right| \right) = \sum_{i=1}^n \sum_{j=1}^n c_i c_j \exp \left(- \int \theta(\xi) \left(\left| \hat{I}_i(\xi) \right| - \left| \hat{I}_j(\xi) \right| \right)^2 d\xi \right) \geq 0 \quad (31)$$

In other words,

$$\sum_{i=1}^n \sum_{j=1}^n c_i c_j \rho(I_i(\cdot), I_j(\cdot)) = \sum_{i=1}^n \sum_{j=1}^n c_i c_j K(F_i(\cdot), F_j(\cdot)) \geq 0, \quad (32)$$

i.e., the proposed SpeD correlation function is positive semi-definite.