

Accelerated, physics-inspired inference of skeletal muscle microstructure from diffusion-weighted MRI

Noel Naughton, Stacey Cahoon, Brad Sutton, *Senior Member, IEEE* and John G. Georgiadis

Abstract—Muscle health is a critical component of overall health and quality of life. However, current measures of skeletal muscle health take limited account of microstructural variations within muscle, which play a crucial role in mediating muscle function. To address this, we present a physics-inspired, machine learning-based framework for the non-invasive estimation of microstructural organization in skeletal muscle from diffusion-weighted MRI (dMRI) in an uncertainty-aware manner. To reduce the computational expense associated with direct numerical simulations of dMRI physics, a polynomial meta-model is developed that accurately represents the input/output relationships of a high-fidelity numerical model. This meta-model is used to develop a Gaussian process (GP) model that provides voxel-wise estimates and confidence intervals of microstructure organization in skeletal muscle. Given noise-free data, the GP model accurately estimates microstructural parameters. In the presence of noise, the diameter, intracellular diffusion coefficient, and membrane permeability are accurately estimated with narrow confidence intervals, while volume fraction and extracellular diffusion coefficient are poorly estimated and exhibit wide confidence intervals. A reduced-acquisition GP model, consisting of one-third the diffusion-encoding measurements, is shown to predict parameters with similar accuracy to the original model. The fiber diameter and volume fraction estimated by the reduced GP model is validated via histology, with both parameters accurately estimated, demonstrating the capability of the proposed framework as a promising non-invasive tool for assessing skeletal muscle health and function.

Index Terms—Diffusion-weighted MRI, Gaussian process, meta-model, microstructure, skeletal muscle

I. INTRODUCTION

Muscle health is strongly correlated to quality of life [1], [2], motivating a clinical need for interventional methods and quantitative diagnostics focused on its staging and improvement. This need is particularly acute for aging populations, as age-related loss of muscle mass is a primary determinant of

decreased muscle function and mobility [3], both of which are linked to increased mortality [4]. Skeletal muscle exhibits a hierarchical structure of elongated, tightly-packed muscle fibers that are surrounded by multiple levels of collagenous extracellular matrix, which plays an important role in force transmission [5]. However, traditional measures of muscle health take limited account for these structural features, restricting our understanding of muscle's structure-function relationship [6]. Non-invasive measurement of skeletal muscle structure is thus positioned to enable novel insights into the physiological changes of muscle during aging [7] or muscle pathology [8], aiding the development of effective, targeted treatments to increase muscle health.

Currently, biopsy and histology is the most common modality to quantitatively investigate skeletal muscle microstructure, but this measurement approach is invasive, labor-intensive, and highly local to the excised muscle region. A promising alternative with potential to address these limitations is diffusion-weighted magnetic resonance imaging (dMRI), which can provide *in-vivo*, non-invasive characterization of muscle microstructure organization over the entire muscle volume. dMRI is sensitive to the direction-dependent diffusion distance of water in tissue. In muscle, water diffuses faster in the fiber's axial direction than transverse direction, where barriers such as cell walls restrict diffusion. This results in anisotropic diffusion behavior that sensitizes the voxel's MR signal to these $\sim\mu\text{m}$ length-scale structures. Thus, although the resolution limit of clinical dMRI is $\sim\text{mm}$, the underlying tissue microstructure is encoded in the MR signal of each voxel.

Estimation of microstructural features from dMRI data has received extensive attention in the brain, with a particular focus on white matter organization [9], [10]. However, while both neural white matter and skeletal muscle share a similar gross morphology of long, axially aligned fibers, decoding the relationship between skeletal muscle microstructure and dMRI measurements is not as simple as applying neural tissue models to skeletal muscle. Skeletal muscle's larger fiber diameter ($\sim 50\ \mu\text{m}$ [11] vs axon's $< 2\ \mu\text{m}$ [12]) and increased membrane permeability [13] substantially complicates the underlying diffusion physics [14], necessitating use of high-fidelity numerical models to accurately capture the physical dynamics of diffusion MRI in muscle [15]–[21].

While these numerical models have successfully explored the forward problem of how microstructure variation influences

This work was supported in part by an NSF GRFP (N.N.), the R.A. Pritzker endowed chair (J.G.), NSF grants CMMI-1437113 & CMMI-1762774 and NIH grants HL090455 & EB018107 (J.G.).

Noel Naughton is with the Department of Mechanical Engineering at Virginia Tech, Blacksburg, VA and was previously with the Beckman Institute for Advanced Science and Technology, University of Illinois at Urbana-Champaign, Urbana, IL (e-mail: nnaughton@vt.edu).

Stacey Cahoon and John G. Georgiadis are with the Department of Biomedical Engineering, Illinois Institute of Technology, Chicago IL.

Brad Sutton is with the Beckman Institute for Advanced Science and Technology, University of Illinois at Urbana-Champaign, Urbana, IL.

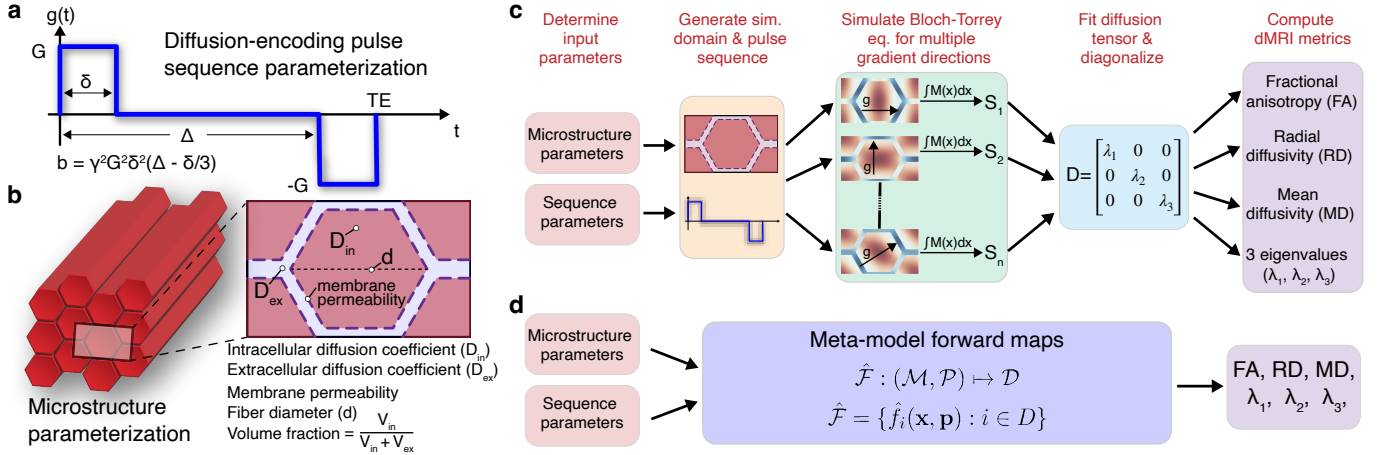


Fig. 1: In-silico experiments of skeletal muscle dMRI. (a) Schematic of the generalized diffusion sequence used to represent the diffusion-encoding sequence. (b) Schematic of numerical model of skeletal muscle as periodically packed hexagonal cylinders. (c) Overview of numerical simulation process with the generation of a pulse sequence and microstructural domain, the independent simulation of multiple diffusion directions, construction of diffusion tensors, and finally estimation of dMRI metrics. (d) Overview of meta-model approach where the microstructural and pulse sequence parameters are directly mapped to the dMRI metrics. Here standard DTI metrics are used, but more complex diffusion models can be straightforwardly included.

dMRI signal [20], [22]–[25], the inverse problem of estimating microstructural parameters from dMRI data remains unresolved. Numerical models are, in general, too computationally expensive to serve as the basis of an inverse solution while compartmental models, originally developed for neural tissue [26] but also at times applied to muscle [27]–[31], are based on assumptions regarding low permeability and small fiber size that do not hold in skeletal muscle, causing them to inaccurately model diffusion physics in muscle [32] and thus be inappropriate for skeletal muscle applications. More recently, the Random Permeable Barrier Model (RPBM) has been proposed, which abstracts muscle as a domain of random, permeable barriers [33], [34] and is sensitive to changes due to increased exercise, muscle atrophy, and muscular dystrophy, among others [34], [35]. However, the RPBM’s basis in a reduced-order physical model limits its ability to be extended to incorporate additional microstructural features. For example, its parameterization of muscle does not account for differences in the intracellular and extracellular domain, which is necessary if changes in the extracellular matrix, a critical mediator of skeletal muscle function [5], are to be considered. Thus, there remains a need for new approaches to characterizing the microstructural organization of skeletal muscle from dMRI.

To address this need, in this paper we combine the accuracy and realism of physics-based numerical models with the computational speed of analytical and data-driven models to develop a framework for estimating the microstructural organization of skeletal muscle from dMRI measurements. In Section II, we describe a numerical physics model of how microstructure and dMRI pulse features influence dMRI measurements (the ‘forward problem’). To accelerate time-to-solutions, we propose a meta-model framework for the forward problem that achieves orders of magnitude faster solutions while retaining a high degree of fidelity to the physics-based numerical model that drives it. In Section III, we describe

a method to estimate the microstructure of skeletal muscle within a voxel given a set of dMRI measurements (the ‘inverse problem’). To accomplish this, we adopt a data-driven approach that leverages the physics-inspired meta-model of Section II. In Section IV, we experimentally demonstrate the viability of the proposed model before proceeding in Sections V and VI with a discussion and conclusion.

Overall, this work provides a flexible framework for development of physics-inspired inversion models for the non-invasive estimation of tissue microstructure from dMRI measurements in an uncertainty-aware manner. While applied here to skeletal muscle modeled using a simplified periodic domain, this methodology is broadly applicable to many classes of biological tissues such as neural and cancer tissues.

II. FAST EVALUATION OF dMRI METRICS IN MUSCLE

To understand how a voxel’s diffusion MRI (dMRI) signal can be used to estimate the underlying microstructural organization of skeletal muscle, we first consider the forward problem of how microstructural variations influence the dMRI signal, the solution of which will form the basis of the inverse method to estimate microstructure from dMRI.

A. Forward problem parameterization

Bloch-Torrey equation: dMRI physics is governed by the Bloch-Torrey equation [36], which describes the time evolution of the dMRI signal in a voxel as

$$\frac{\partial M}{\partial t} = -i\gamma(\mathbf{x} \cdot \mathbf{g}(t))M + \nabla \cdot (D(\mathbf{x})\nabla M) \quad (1)$$

where $M(\mathbf{x}, t)$ is the complex-valued, transverse spin magnetization resulting from exciting longitudinal spins onto the transverse plane and which is manipulated by an externally applied magnetic field ($\mathbf{x} \cdot \mathbf{g}(t)$). Here γ is the gyromagnetic ratio of hydrogen, \mathbf{x} is the spin position vector, $\mathbf{g}(t)$ is the

time-varying magnetic field gradient vector used to encode diffusion, and $D(\mathbf{x})$ is the local diffusion coefficient.

Pulse sequence parameterization: The magnetic field gradient $\mathbf{g}(t)$ is operator-controlled and can be manipulated to probe different aspects of tissue structure [37]. We focus on two related diffusion-encoding pulse sequences: the Stejskal-Tanner pulsed-gradient spin echo (PGSE) [38] and the simulated echo acquisition mode (STEAM) diffusion sequence [39]. If T_1 and T_2 effects are set aside, both sequences can be minimally described by the generalized diffusion-weighted sequence described in Fig. 1a that consists of a bipolar magnetic gradient pulse of magnitude G [25], [40]. This generalized sequence is parameterized by four variables, echo time (TE), gradient duration (δ), gradient spacing (Δ), and b-value ($b = \gamma^2 G^2 \delta^2 (\Delta - \delta/3)$). Here, gradient duration is fixed at 5 ms and TE is defined as $\text{TE} = \Delta + \delta$, resulting in two free variables of gradient spacing and b-value. The diffusion time of the spins is $t_d = \Delta - \delta/3$. However, because our fixed gradient duration of 5 ms is short relative to the gradient spacing timings considered here, for simplicity we approximate the gradient spacing as the diffusion time (i.e. $t_d = \Delta$).

Muscle tissue model parameterization: At the microstructural level, skeletal muscle consists of parallel, elongated fibers, each surrounded by a semi-permeable membrane (sarcolemma) and embedded in an extracellular matrix. Informed by histologically based simulations [21], we represent muscle’s microstructural organization via a compact domain of infinitely long, parallel hexagonal cylinders arranged in a periodic array (Fig. 1b). We define a representative elemental volume (REV), which we parameterize to provide a parsimonious description of the muscle microstructure consisting of two morphological parameters (fiber diameter and muscle fiber volume fraction) and three tissue parameters (membrane permeability and intra/extracellular diffusion coefficients). Water diffusion in the intra- and extracellular domains is characterized by homogeneous (effective) diffusion coefficients that capture the cumulative effects of sub-cellular restrictions within each domain.

B. Physics-based numerical simulations

The governing Bloch-Torrey equation of Eq. 1 is integrated using the lattice Boltzmann method (LBM) on a D3Q7 stencil, supplemented with appropriate boundary conditions and the initial condition $\mathbf{M}(\mathbf{x}, t = 0) = 1$. Intra-domain, semi-permeable boundary conditions handle the effect of spins crossing the muscle’s sarcolemma membrane while appropriately-defined periodic boundary conditions are applied on the domain boundaries to represent a periodic REV geometry. Full details of both the boundary conditions and the numerical lattice Boltzmann scheme implementation are available in [21].

Solving the Bloch-Torrey equation over the prescribed domain results in a spatially localized distribution of the MR signal $\mathbf{M}(\mathbf{x}, t = \text{TE})$ (see Fig. 1c). In a physical MR experiment, image formation integrates this local signal distribution to provide the voxel’s dMRI signal. Integrating the numerical simulation result over the simulation domain

$$S = \left\| \int_V \mathbf{M}(\mathbf{x}, t = \text{TE}) d\mathbf{x} \right\| \quad (2)$$

TABLE I: Microstructure and pulse sequence parameter ranges spanned by meta-model.

Parameter	Range			
Muscle fiber diameter	10	–	80	μm
Volume fraction	0.7	–	0.95	
Membrane permeability	10	–	100	$\mu\text{m/s}$
Intracellular diffusion	0.5	–	2.5	$\mu\text{m}^2/\text{ms}$
Extracellular diffusion	0.5	–	2.5	$\mu\text{m}^2/\text{ms}$
Diffusion time	10	–	750	ms
b-value	300	–	1200	s/mm^2

where V is the voxel volume allows matching the MR signal that would be measured on a scanner. This equivalence enables *in silico* dMRI experiments, where a known microstructural domain can be defined and a dMRI experiment performed to computationally evaluate the signal.

A schematic overview of the simulation pipeline for these *in silico* experiments is given in Fig. 1c. For each experiment, six non-collinear gradient directions (S_i) and a non-diffusion-weighted acquisition (S_0) are simulated and used to fit a diffusion tensor using the fanDTasia ToolBox [41]. In muscle, the diffusion tensor is anisotropic and described by three eigenvalues (λ_1 , λ_2 , and λ_3), which correspond to the tensor’s principal directions. These eigenvalues are then used to compute the tensor invariants of fractional anisotropy $\text{FA} = \frac{\sqrt{(\lambda_1 - \lambda_2)^2 + (\lambda_2 - \lambda_3)^2 + (\lambda_3 - \lambda_1)^2}}{2(\lambda_1^2 + \lambda_2^2 + \lambda_3^2)}$, mean diffusivity $\text{MD} = (\lambda_1 + \lambda_2 + \lambda_3)/3$, and radial diffusivity $\text{RD} = (\lambda_2 + \lambda_3)/2$, which characterize the diffusion anisotropy and magnitude within the voxel. While in this work we use a standard diffusion tensor imaging (DTI) model, more complex post-processing and diffusion models can straightforwardly be incorporated.

C. Model acceleration via meta-modeling

While accurate, numerical simulation of the forward problem is computationally expensive, with a typical *in silico* dMRI experiment taking 1-2 minutes per voxel. Coupled with direct inverse solution approaches, which require many iterative forward solutions, use of this numerical model results in solution times on the order of 45 minutes per voxel to estimate underlying microstructural parameters [42]. Scaled to a typical 64×64 (or larger) resolution dMRI image with multiple slices, this approach quickly becomes computationally infeasible.

To increase solution speed, we exploit the insight that the intermediate steps of Fig. 1c—which are computationally expensive to obtain—are in fact only necessary to estimate a diffusion tensor for the tissue and can be bypassed (Fig. 1d) by the deployment of a physics-inspired data-driven forward mapping

$$\mathcal{F} : (\mathcal{M}, \mathcal{P}) \mapsto \mathcal{D} \quad (3)$$

that directly maps the microstructural ($\mathcal{M} \in \mathbb{R}^m$) and diffusion-encoding sequence ($\mathcal{P} \in \mathbb{R}^p$) domains to the dMRI metrics ($\mathcal{D} \in \mathbb{R}^d$). Beyond computational efficiency, this mapping entails a number of advantages. In particular, by directly estimating the diffusion tensor eigenvalues and other metrics, it is not sensitive to the number and direction of the diffusion-encoding gradient measurements while retaining the physical information of the numerical model it is based on.

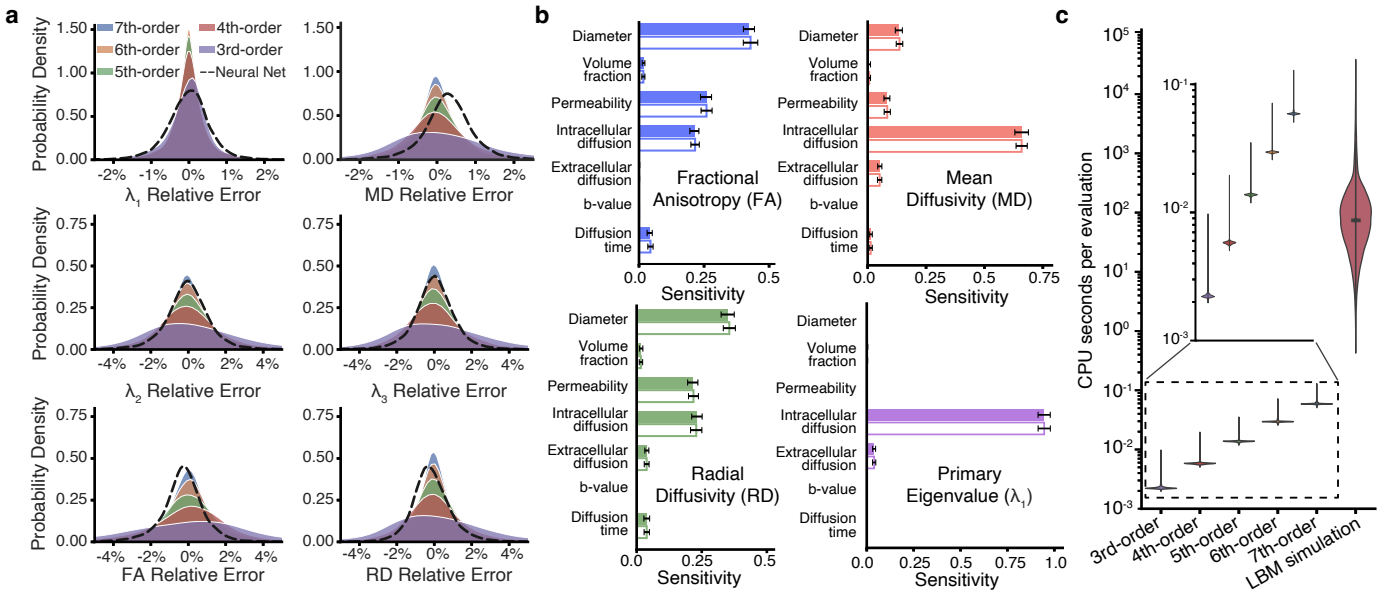


Fig. 2: Meta-model acceleration of dMRI simulations. (a) Density plots of meta-model error for dMRI metrics of FA, MD, RD, λ_1 , λ_2 , and λ_3 . Solid colors denote a polynomial meta-model while dashed lines are for a feed-forward neural network meta-model. As the order of the polynomial basis set increases, the meta-model accuracy improves. (b) Global sensitivity indices of the 7th-order meta-model (solid color) and LBM numerical model (white) demonstrating the meta-model captures the global behavior of parameter variation. Results for λ_2 and λ_3 and not shown as they are nearly identical to RD results. (c) The computational cost of meta-model evaluation increases with polynomial order but remains lower than the numerical model in all cases. The 7th-order meta-model has a mean evaluation time three orders of magnitude (1000x) faster than the numerical model and a maximum evaluation time five orders of magnitude faster. All models were evaluated on a machine with an Intel Xeon W-2265 processor.

Polynomial meta-model: We approximate the forward mapping \mathcal{F} using a set of meta-models to generate individual mappings for each dMRI metric

$$\hat{\mathcal{F}} = \{\hat{f}_i(\mathbf{x}, \mathbf{p}) : i \in D\} \quad (4)$$

for $D \in \{\text{FA, MD, RD, } \lambda_1, \lambda_2, \lambda_3\}$ and where $\mathbf{x} \in \mathcal{M}$ and $\mathbf{p} \in \mathcal{P}$. These meta-models (Fig. 1d) approximate the relationship between the $m = 5$ microstructural features (Fig. 1b) and $p = 2$ dMRI pulse sequence parameters (b-value and diffusion time; Fig. 1a) and the $d = 6$ dMRI metrics in a data-driven manner with no explicit dependence on the underlying physics.

While a number of meta-modeling frameworks [43] and machine learning techniques [44], [45] are available, here we adopt a polynomial expansion approach [46], [47]. Each meta-model is represented as an expansion of a properly selected polynomial basis truncated to a finite basis set

$$\hat{f}_i(\mathbf{x}, \mathbf{p}) = \sum_{j=0}^T \alpha_j \psi_j(\mathbf{z}) \quad (5)$$

where α_j are the basis weights, ψ_j are multivariate polynomials, $T = (N + C)! / (N! C!) - 1$ is the number of terms in the basis set for a maximum polynomial order of C and $N = m + p$ input parameters, and $\mathbf{z} = \Gamma(\mathbf{x}, \mathbf{p}) \in [-1, 1]^N$ is the linear transformation of the microstructure and pulse sequence parameter ranges given in Table I to the interval $[-1, 1]$ (e.g. for a single parameter $x \in [a, b]$, the transformation would be $2(x-a)/(b-a)-1$). The choice of basis polynomials ψ_j is

determined by the parameter distribution. To ensure even and unbiased coverage of the microstructural and pulse sequence input ranges in Table I, a uniform distribution is considered, for which Legendre polynomials are the appropriate basis [48].

Meta-model generation: To generate the meta-model, a set of 80,000 *in silico* dMRI experiments were performed using the LBM numerical model (Fig. 1c). For each simulation, a tissue domain parameterized by five microstructural features was defined (Fig. 1b) over which a dMRI measurement defined by the two pulse sequence features was numerically simulated and processed (Fig. 1a), resulting in a set of six dMRI metrics. Microstructure and pulse parameter combinations were generated via a Sobol sampling method, which is a low discrepancy sampling method, meaning the input data used to construct the model are evenly distributed across the input space [49]. The resulting dataset was used to fit the weights α_j of the PC expansion using a 70/30 train/test split. Data was fit using least-squares linear regression by the open-source *ChaosPy* Python package [50]. Each dataset entry consists of a microstructural parameter set \mathbf{x} , a pulse sequence parameter set \mathbf{p} , and dMRI metrics $f_i(\mathbf{x}, \mathbf{p})$ for $i \in D$ that resulted from an *in silico* dMRI experiment following the procedure described in Section II-A. Maximum polynomial orders $C \in \{3, \dots, 7\}$ of the Legendre polynomial basis were considered. To provide a baseline comparison, a multi-layer perceptron regression model was trained on the same data set. The regression model consisted of a feedforward neural network with three fully connected hidden layers of 200 neurons and

ReLU activation units. The input to the model is the normalized microstructural parameters while the output is the dMRI metric under consideration. The model was trained using the same 70/30 test/train data split as the polynomial meta-model using the Adam optimizer with a constant learning rate of 0.001 for 200 iterations using the scikit-learn package in Python. A separate neural network model was trained for each dMRI metric.

Accuracy and sensitivity: Evaluating the trained meta-models over the test data resulted in a vector of meta-model estimates $\hat{\mathbf{f}}_i$ and an associated vector of ground-truth, numerically-simulated dMRI metrics \mathbf{f}_i . To quantify the accuracy of the meta-model we consider the accuracy of each dMRI metric independently and define a relative error metric

$$e_i = \frac{\hat{f}_i(\mathbf{x}, \mathbf{p}) - f_i(\mathbf{x}, \mathbf{p})}{\bar{\mathbf{f}}_i} \quad (6)$$

where $\bar{\mathbf{f}}_i$ is the mean of the vector \mathbf{f}_i .

Figure 2a shows the distribution of relative error between the meta-models and the ground-truth numerical simulations in the test split of the data set for each dMRI metric. For increasing polynomial order, the accuracy of the meta-model improves, with the 7th-order polynomial model the most accurate overall, slightly outperforming the neural network meta-model implementation. While we use the polynomial meta-model throughout the rest of this paper due to its higher overall performance, we note that both approaches provide a level of performance sufficient to accurately represent the underlying numerical model, with errors on the order of <1% for all metrics. A limitation on polynomial order comes from the risk of overfitting. Here the training data set contained 56,000 entries while the number of coefficients for the maximum polynomial size of 7th-order was $T = 3431$ for $N = 7$ (input parameters) and $C = 7$ (polynomial order). Using a heuristic of 10 entries per free parameter to avoid overfitting suggests that 7th-order polynomials is the maximum order that should be used ($T = 6434$ for an 8th-order polynomial).

To further quantify the performance of the meta-model, the sensitivity of the meta-model to changes in its microstructure and pulse sequence parameter inputs was computed and compared to a previously reported sensitivity analysis of the Bloch-Torrey equation [25]. The results, shown in Fig. 2b, show the relative importance of each microstructural feature on the dMRI metrics. Further, the computed sensitivity indices of the meta-model are nearly identical to those of the full numerical model, demonstrating that the fitted meta-model accurately captures the global responses of the underlying numerical model for each microstructural input parameter.

Computational efficiency: As the polynomial order increases, the accuracy of the meta-model increases, but so too does the computational evaluation time (Fig. 2c). However, the trained meta-models are drastically faster to evaluate than the underlying LBM numerical model for all polynomial orders, with the mean evaluation time of the 7th-order meta-model three orders of magnitude faster than the mean evaluation time of the numerical model. This speed-up becomes even more pronounced when maximum evaluation time is considered, with the 7th-order meta-model five orders of magnitude faster than

the numerical model. The time-stepping nature of the LBM numerical model leads to its evaluation time being directly proportional to the diffusion time of the simulated sequence. The analytical nature of the meta-model, in contrast, removes this limitation, leading to more uniform evaluation times.

While the results of Fig. 2a suggest that for best accuracy the highest possible polynomial order should be used, incorporating consideration of the computational expense of increasing polynomial order suggests the optimal polynomial order will be case-specific based on the number of training samples available, the compute budget available for fitting and evaluation, and the desired accuracy of the meta-model. In the next section, both 5th- and 7th-order polynomial meta-models are used. This allows maximum accuracy (7th-order) with the possibility of a 5-7x speedup in evaluation speed (5th-order) when necessary (while maintaining a lower, yet still acceptable level of accuracy).

III. INVERSE PROBLEM SOLUTION

We next turn to the inverse problem of estimating an individual voxel's microstructural organization based on a set of dMRI measurements. That is, we ask if we can 'invert' the flow of data in the forward model to identify what microstructure yielded a set of dMRI measurements. This inverse model will take as its input the values of a dMRI experiment and provide as its output a quantitative estimate of microstructural parameters. Previous attempts have generally consisted of iteratively solving a forward model to converge on a set of microstructural parameters [27]–[31], [34], [42]; however, such approaches are computationally expensive due to the large number of function evaluations required. Our approach is to instead define a data-driven inverse map

$$\mathcal{G} : \mathcal{S} \mapsto \mathcal{M} \quad (7)$$

that directly maps the dMRI measurement domain $\mathcal{S} \in \mathbb{R}^{s \times d}$ to the microstructure parameter domain $\mathcal{M} \in \mathbb{R}^m$. Here, s is the number of dMRI measurements (distinct combinations of b-value and diffusion time) acquired, d is the number of dMRI metrics, and m is the number of microstructural features that parameterize the tissue.

While machine learning approaches have been widely used to generate cost-effective inverse solutions for a range of problems [44], [51], [52], many traditional approaches, such as random forest regression or deep neural networks, do not automatically account for uncertainty in their estimations, though substantial recent work has sought to address this limitation [53], [54]. Accounting for such uncertainty can be critically important to interpreting results as it provides knowledge of the confidence one should place in the estimate. To address this need, here we utilize Gaussian process (GP) regression, or kriging, to generate a data-driven inverse map that also provides confidence intervals of its predictions, substantially increasing the interpretability, and thus utility, of the estimates [55].

A. Gaussian process regression

Each microstructural parameter is modeled as its own Gaussian process, leading to the inverse map

$$G = \{g_i(\mathbf{s}) : i \in M\} \quad (8)$$

where $\mathbf{s} \in \mathcal{S}$ is a vector containing a list of d dMRI metrics for each of s diffusion-encoding sequences (defined in Section III-B) and M is the domain of all five microstructural parameters.

A Gaussian process is a generalization of a Gaussian probability distribution. It is a collection of random variables, any finite subset of which has a joint Gaussian distribution [56]. A Gaussian process

$$g_i(\mathbf{s}) \sim \mathcal{GP}(m_i(\mathbf{s}), k_i(\mathbf{s}, \mathbf{s}')). \quad (9)$$

is defined by a mean function $m_i(\mathbf{s})$ and covariance function $k_i(\mathbf{s}, \mathbf{s}')$ over the inputs \mathbf{s} . In practice, is it common to either subtract out the mean or directly set $m_i(\mathbf{s}) = 0$, allowing the Gaussian process to be written as

$$g_i(\mathbf{s}) \sim \mathcal{N}(0, k_i(\mathbf{s}, \mathbf{s}')) \quad (10)$$

where $\mathcal{N}(\cdot)$ denotes a normal distribution.

A key component of Gaussian process regression is the selection of the covariance kernel function $k_i(\mathbf{s}, \mathbf{s}')$, which can strongly influence the model's accuracy. Progress has been made towards automated kernel selection [57]; however empirical kernel selection is still generally required. Here, a radial basis function kernel and a linear kernel are combined with a Gaussian noise kernel to form the covariance function

$$k_i(\mathbf{s}, \mathbf{s}') = \sigma_r^2 \exp\left(-\frac{1}{2l^2}(\mathbf{s} - \mathbf{s}')^2\right) + \sigma_l^2(\mathbf{s} - c)(\mathbf{s}' - c) + \sigma_n^2 \delta_{ij} \quad (11)$$

where kernel variances σ and length-scales l , c are hyperparameters tuned to maximize the log marginal likelihood of the model over the training data. A unique set of hyperparameters is tuned for each GP $g_i(\mathbf{s})$. While it is possible to incorporate information about relationships between microstructural parameters (coregionalization), such methods substantially increase the difficulty of hyperparameter optimization and often lead to over-fitting [56]. As such, they were not considered here.

Here, training data are denoted as \mathbf{S} and $\mathbf{g}_i(\mathbf{S})$, where \mathbf{S} is the set of dMRI measurements (inputs) for which the associated microstructural parameters (outputs) $\mathbf{g}_i(\mathbf{S})$ are known. Similarly, \mathbf{S}^* and $\mathbf{g}_i^*(\mathbf{S}^*)$ are the testing or evaluation data, where $\mathbf{g}_i^*(\mathbf{S}^*)$ are the unknown microstructural parameters that need to be estimated based on some measured dMRI data \mathbf{S}^* . For the training data, $\mathbf{S} = \{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_n\}$ is a set of dMRI data for n voxels (or their computational equivalent from the forward model). Each voxel \mathbf{s}_i in turn consists of multiple independent dMRI measurements as described after Eq. 8. Similarly, for test data, $\mathbf{S}^* = \{\mathbf{s}_1^*, \mathbf{s}_2^*, \dots, \mathbf{s}_{n^*}^*\}$ are the dMRI data associated with n^* voxels. Note that the non-bolded term $g_i(\mathbf{s})$ of Eqs. 9 and 10 is the general Gaussian process while the bolded quantities $\mathbf{g}_i(\mathbf{S})$ and $\mathbf{g}_i^*(\mathbf{S}^*)$ can be viewed as drawn samples of $g_i(\mathbf{s})$.

The inverse problem is then formulated as follows: given a training set of n observations of microstructure parameters $\mathbf{g}_i(\mathbf{S})$ and their corresponding dMRI measurements \mathbf{S} , estimate the microstructural parameter distribution $\mathbf{g}_i^*(\mathbf{S}^*)$ associated with a set of n^* dMRI measurements \mathbf{S}^* for which microstructural parameters are unknown. According to the definition of Gaussian processes, the joint probability distribution of these

training and evaluation outputs is also Gaussian and can be written

$$\begin{bmatrix} \mathbf{g}_i \\ \mathbf{g}_i^* \end{bmatrix} \sim \mathcal{N}\left(0, \begin{bmatrix} K_i(\mathbf{S}, \mathbf{S}) & K_i(\mathbf{S}, \mathbf{S}^*) \\ K_i(\mathbf{S}^*, \mathbf{S}) & K_i(\mathbf{S}^*, \mathbf{S}^*) \end{bmatrix}\right) \quad (12)$$

where, $K_i(\mathbf{S}^*, \mathbf{S})$ is the $n^* \times n$ matrix of the covariances of \mathbf{S}^* and \mathbf{S} based on Eq. 11 (and similarly for all $K_i(\cdot, \cdot)$).

To allow the training data information $\mathbf{g}_i(\mathbf{S})$ to inform our estimation of \mathbf{g}_i^* , we condition the joint Gaussian prior distribution of Eq. 12 on the training observations \mathbf{g}_i , resulting in the posterior distribution

$$\mathbf{g}_i^* | \mathbf{S}^*, \mathbf{S}, \mathbf{g}_i \sim \mathcal{N}(\bar{\mathbf{g}}_i^*, \text{cov}(\mathbf{g}_i^*)) \quad (13)$$

where the mean $\bar{\mathbf{g}}_i^*$ and covariance matrix $\text{cov}(\mathbf{g}_i^*)$ are

$$\bar{\mathbf{g}}_i^* = K_i(\mathbf{S}^*, \mathbf{S})K_i(\mathbf{S}, \mathbf{S})^{-1}\mathbf{g}_i \quad (14)$$

$$\text{cov}(\mathbf{g}_i^*) = K_i(\mathbf{S}^*, \mathbf{S}^*) - K_i(\mathbf{S}^*, \mathbf{S})K_i(\mathbf{S}, \mathbf{S})^{-1}K_i(\mathbf{S}, \mathbf{S}^*) \quad (15)$$

Note that both the mean $\bar{\mathbf{g}}_i^*$ and covariance matrix $\text{cov}(\mathbf{g}_i^*)$ are solely functions of the training microstructural observations \mathbf{g}_i , the training dMRI data \mathbf{S} , and the observed dMRI data \mathbf{S}^* , all of which are known. Thus, given a set of dMRI measurements associated with n^* voxels $\mathbf{S}^* = \{\mathbf{s}_1^*, \mathbf{s}_2^*, \dots, \mathbf{s}_{n^*}^*\}$, we can evaluate the posterior distribution to compute voxel-wise mean estimates $\mu_i(\mathbf{s}_n^*)$ for each microstructural parameter (Eq. 14) along with a variance $\sigma_i(\mathbf{s}_n^*)$ (Eq. 15), which in turn provides a 95% confidence interval. Each voxel is considered independently and without regard to its neighboring voxels.

B. dMRI-inversion model formulation

To formulate an inverse mapping, a fixed set of diffusion-encoding parameter sets must first be defined. Because multiple microstructural parameters are being estimated, multiple diffusion-encoding measurements are required to constrain the inverse problem. Here, we focus on the effect of diffusion time and b-value. Multiple b-values provides sensitivity to non-Gaussianity of the diffusion behavior [58] while varying diffusion time sensitizes the diffusion behavior to microstructural features at different length scales [34], [59], [60].

Five b-values (400, 600, 800, 1000, and 1200 s/mm²) and six diffusion times (20, 50, 100, 200, 400, 700 ms) were selected to generate $s = 30$ unique diffusion-encoding parameter sets. Here, all sequences are modeled as STEAM sequences represented by the generalized diffusion sequence described in Section II-A. The b-value is achieved by adjusting the diffusion-encoding gradient strength. Both b-values and diffusion times were selected to span the ranges viable for a clinical scanner using a STEAM sequence for skeletal muscle dMRI.

Model training and evaluation: To generate training data for the GP model, the 7th-order meta-model was sampled using a Sobol sequence to generate 9000 microstructural parameter combinations. For each combination, the meta-model was evaluated for thirty diffusion-encoding parameter sets, resulting in 270,000 forward problem evaluations. For each diffusion-encoding, FA, λ_1 , and RD were computed and used as inputs to the model (λ_2 and λ_3 were not considered due to their similarity to RD and MD is a linear combination of λ_1 and

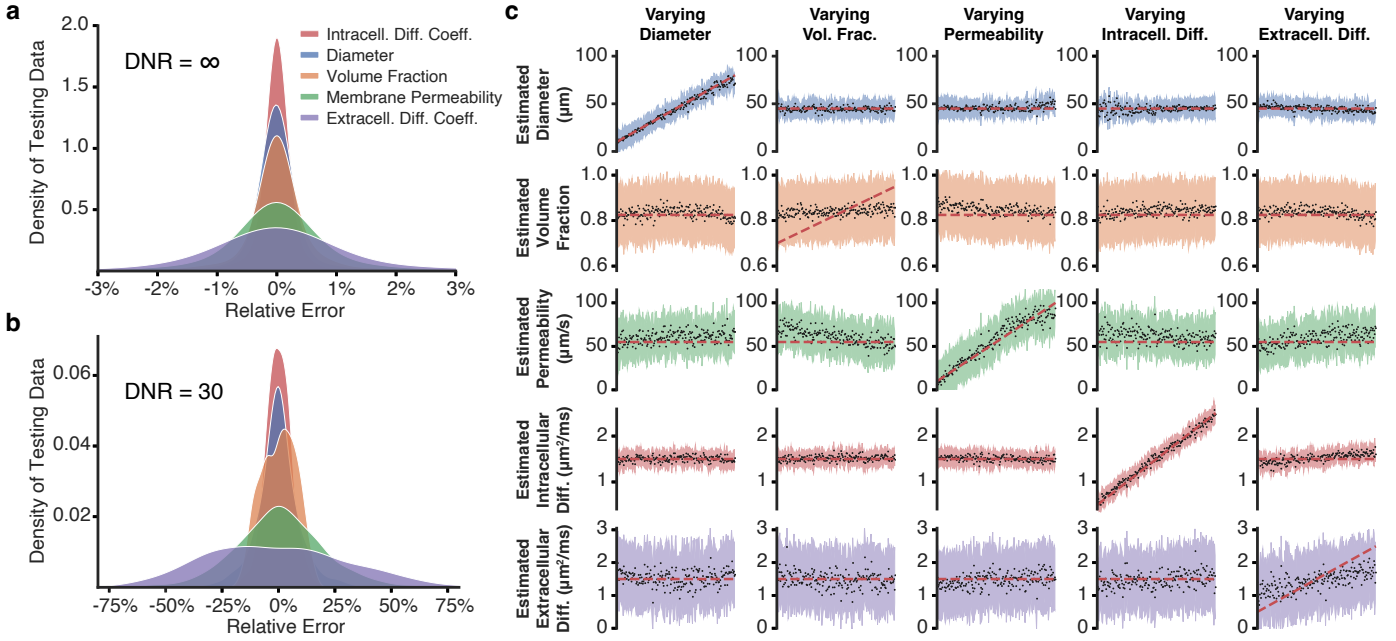


Fig. 3: Gaussian process inverse mapping. (a) Error density of GP model when evaluating synthetic, noise-free dMRI measurements. (b) Error density of GP model when evaluating noisy dMRI measurements with a diffusion-to-noise ratio of 30. (c) GP model results for dMRI measurements (DNR=30) when a single microstructural parameter is varied at a time. For each column, 200 meta-model evaluations were made as the varied microstructure parameter was linearly increased. Results are organized column-wise with vertical alignment relating to the same dMRI measurement and the dashed line denoting the true microstructural value.

RD). To reduce the dimensionality of the input data, these ninety inputs were combined into a single vector for each parameter set, and principal component analysis (PCA) was performed on the vectors. The 20 dimensions that best described the observed input variance were selected, scaled to zero-mean and unit-variance, and used as the inputs to the GP model.

The meta-model produces noise-free estimates of dMRI metrics; however, accounting for the influence of noise in dMRI measurements is a critical consideration of any inverse solution [61] and so synthetic noise was injected into the training data. The meta-model directly outputs dMRI metrics, bypassing the computationally expensive simulation of the raw signal within the tissue domain, which is the stage in a physical experiment where noise is introduced into the measurement. As such, it is not possible to incorporate noise directly into the raw signal. Instead, Gaussian noise is injected into the dMRI metrics based on a diffusion-to-noise (DNR) ratio. Five copies of the dMRI metrics were created. One remained noise-free with noise added to the others to achieve a DNR of 30.

The open-source *GP* Python package was used to fit the GP model and optimize the GP kernel hyperparameters using the L-BFGS-B algorithm [62]. For large data sets such as those considered here, this optimization can be slow. To accelerate the process, subsets of increasing size containing only noise-free data were iteratively used to optimize the hyperparameters over smaller data sets, allowing the GP model to quickly learn the coarse structure of the underlying data. In total, model training took an average of 2.2 hours for each microstructural parameter (13.2 hours total) using a machine with 2x Intel

Xeon E5-2698 processors and 1024 GB of DDR4 RAM.

GP model accuracy: The 7th-order meta-model was sampled using a testing set of 3000 microstructural parameter combinations determined by a different Sobol sequence. Evaluating the test data resulted in vectors of mean microstructural estimates μ_i for $i \in M$ and associated vectors of ground-truth microstructural values x_i . To quantify GP model accuracy, we define the relative error metric

$$e = \frac{\mu_i(\mathbf{s}) - x_i(\mathbf{s})}{\bar{x}_i} \quad (16)$$

where $x_i(\mathbf{s})$ is the true microstructural parameter. The accuracy of the model was quantified against both noise-free (Fig. 3a) and noisy (DNR = 30; Fig. 3b) versions of the test data. While the accuracy of microstructural estimates decreases in the presence of noise, the GP model's high accuracy when given noise-free data demonstrates both the general invertibility of the problem and that the GP model is learning the underlying data structure.

A third test is depicted in Fig. 3c wherein test data is generated by evaluating the forward meta-model and varying only a single microstructural parameter at a time (diagonal entries) while holding all others constant. As Fig. 3c shows, the GP model is able to accurately capture changes in fiber diameter, membrane permeability, and intracellular diffusion coefficient while struggling to identify variations in the volume fraction or extracellular diffusion coefficient. However, the GP model's confidence intervals in its mean estimate of the volume fraction and extracellular diffusion coefficient are wide, indicating the mean parameter estimate should be interpreted with caution. In contrast, its accurate estimates of fiber diameter,

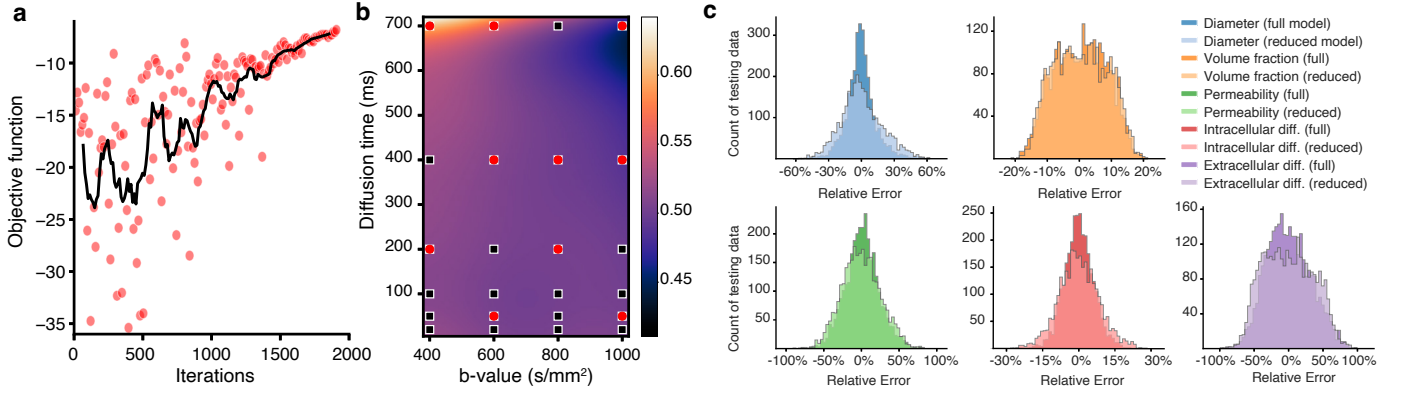


Fig. 4: Optimized diffusion-encoding sequence selection. (a) Convergence of CMA-ES algorithm's maximization of the objective function over 2000 iterations to select a compact set of ten diffusion-encoding sequences. (b) Map of the variance between the 1024 different microstructural parameters used for different diffusion time and b-value combinations. The locations of the sequence used for the full GP model (black boxes; $b=1200$ s/mm² not visualized) and ten sequences selected by the CMA-ES algorithm (red circles) are overlaid. (c) Comparisons of the relative error histograms of the full and reduced sequence GP models for all microstructural parameters shows limited loss in accuracy when the reduced set of sequences are used.

membrane permeability, and intracellular diffusion coefficient are accompanied by comparably narrower confidence intervals, indicating confidence in the mean estimate.

C. Reduced diffusion-encoding model

While it is broadly known how different diffusion-encoding pulse profiles affect the MR signal [58], which combinations of sequences encode the most microstructural information about the tissue is less clear. In the prior section, this uncertainty is addressed by densely sampling the possible combinations of diffusion time and b-value (Fig. 4b). However, because imaging time is proportional to the total number of diffusion-encoding parameter sets utilized, reducing the number of parameter sets used is paramount for clinical translatability.

To identify a compact set of diffusion-encoding sequences, the 5th-order meta-model (for computational efficiency) was sampled at four uniformly spaced points for all five microstructural parameters ($N = 1024$) and FA, MD, RD, and λ_1 metrics were precomputed for 80 pulse sequences defined on a grid of four evenly-spaced b-values between 400 and 1000 s/mm² and twenty evenly-spaced diffusion times between 10 and 700 ms. For a candidate subset of $s = 10$ sequences, a 3rd-order polynomial surface was fit to the dMRI metrics at each of the N microstructural parameter sets. The fitted surface was evaluated at all 80 pulse sequences resulting in a vector \mathbf{q}_{i_n} for $i \in \{\text{FA, MD, RD, } \lambda_1\}$ that was compared to the precomputed meta-model dMRI metrics \mathbf{Q}_{i_n} . An objective function was defined to quantify the ability of a subset of ten diffusion-encoding pulse sequences to capture the structural information encoded by this full set of sequences according to

$$\phi = - \sum_{n=1}^N \left\| \left(\|\mathbf{q}_{i_n} - \mathbf{Q}_{i_n}\| : i \in \{\text{FA, MD, RD, } \lambda_1\} \right) \right\| \quad (17)$$

We maximized Eq. 17 using the non-linear, derivative-free evolutionary search CMA-ES algorithm [63] (Fig. 4a). The ten sequences selected by the CMA-ES algorithm were matched

to the nearest corresponding pulse sequences used to define the previous GP model (Fig. 4b). The average r^2 between the surface \mathbf{q}_{i_n} of the selected sequences and \mathbf{Q}_{i_n} was 0.88.

Reduced GP model: A second GP model was fit using this reduced set of sequences following the same process as before. A comparison of the full and reduced GP model's accuracy is shown in Fig. 4c. To test the model's generalization performance, both models are tested against a dataset of 3000 solutions generated by the LBM numerical model. Results for both models are broadly similar, with only a slight decrease in accuracy for the reduced pulse sequence model. Using only one-third of the diffusion-encoding measurements, this set of $s = 10$ sequences defines a more experimentally plausible set of sequences for use in a physical experiment.

IV. EXPERIMENTS

A. Voxel-wise microstructure estimation

To test the accuracy of the proposed inversion scheme, a validation experiment was performed. Tissue from a bovine *biceps femoris* was procured and a 7 cm sample from a unipennate portion was excised and imaged on a Siemens Prisma 3T scanner using a prototype diffusion-weighted STEAM sequence with EPI readout [39]. Ten diffusion-weighted measurements based on the sequences selected by the CMA-ES algorithm were obtained, each with twelve gradient directions and 2 averages. Twenty slices with a 64×64 FOV matrix and 3.4 mm isotropic voxels were acquired. For all scans, TE = 49 ms. For scans with $\Delta > 500$ ms, TR = 1500 ms, otherwise TR = 1200 ms. The experiment took approximately one hour to complete. Image data was thresholded and post-processed with FSL [64] to compute diffusion metrics for each voxel.

This dataset (9607 voxels) was read into the reduced GP model and used to produce voxel-wise estimates of the microstructural parameters. Here, λ_3 was used instead of RD to better capture microstructural data based on the hypothesis that λ_3 is a stronger reflection of fiber-level transverse diffusion behavior [65], [66]. Two ROIs were also manually defined

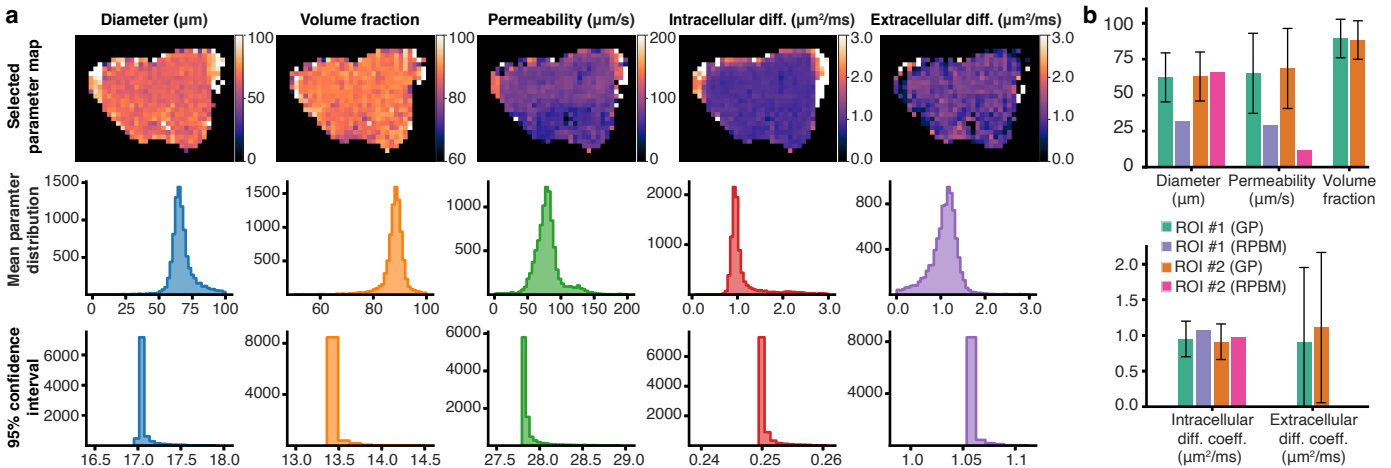


Fig. 5: Experimental estimates of microstructure organization. (a) Voxel-wise estimates of all five microstructural parameters for a single slice of the bovine *biceps femoris* dMRI data (top row); histograms of the distribution of the five microstructural parameters over the entire 3D domain after thresholding to exclude outliers related to edge voxels (middle row); and distribution of 95% confidence intervals from GP model over entire 3D domain (bottom row). (b) GP estimates and 95% confidence intervals of all five microstructural parameters and RPBMs estimates of three microstructural parameters for two ROIs.

within the 3D tissue volume, consisting of 12 and 32 voxels. Diffusion metrics from these voxels were averaged and used to estimate mean microstructural parameters and confidence intervals for the ROI. Additionally, the previously proposed Random Permeable Barrier Model (RPBM) [33], [67] was fit to the averaged ROI data for comparison. The RPBM treats muscle tissue as consisting of only an intracellular domain that is intersected by randomly oriented permeable membranes aligned in the cross-sectional plane. This parameterization yields three microstructural features, here reported as the diffusion coefficient, fiber diameter, and membrane permeability. Fitting the RPBM takes 1-2 seconds per ROI or voxel.

The top row of Figure 5a shows the resulting voxel-wise estimates of all five microstructural parameters for a representative slice of the 3D volume. For each of the five microstructural parameters, estimation for all 9607 voxels took 39.18 ± 0.05 seconds (4 ms/voxel). Parameter ranges appear consistent throughout the muscle beyond slight variation in voxels near the edge of the tissue, where partial volume effects likely distorted the signal. The middle row of Fig. 5a shows histograms of the distribution of the five microstructural parameters after thresholding to exclude outliers related to edge voxels while the bottom row of Fig. 5a shows the distribution of the associated 95% confidence intervals. Mean parameter estimates have unimodal distributions with relatively small tails while the confidence intervals exhibit a sharp minimum bound, indicative of experimental noise limiting the confidence of the inversion scheme. Results from the two ROIs for both GP and RPBM estimates are reported in Fig. 5b. Diffusion coefficient estimates are consistent between the GP and RPBM for both ROIs, while the RPBM estimates a substantially lower membrane permeability that is below the 95% confidence interval of the GP model. There is mixed agreement between the two models for fiber diameter, with estimates agreeing for ROI #2 but the RPBM estimating a fiber diameter below the 95% confidence interval of the GP model for ROI #1.

B. Histological validation

Validation of microstructure estimates is challenging due to the difficulty of independently measuring many of the estimated microstructural parameters, in particular intracellular and extracellular diffusion coefficients or sarcolemma membrane permeability. Morphological parameters, such as fiber diameter and volume fraction, are comparably easier to verify through histological examination. Tissue localized to the two ROIs was excised and fixed in 10% buffered formalin for one week. After fixation, tissue was dehydrated, embedded in paraffin wax, sectioned, and stained with hematoxylin and eosin. Microscopy images were acquired at 40x magnification. Images were processed using ImageJ [68] to binarize data, and fiber diameters for each ROI were approximated by the median Feret diameter using the ‘Analyze Particles’ tool in ImageJ. During tissue processing, distortion due to tissue shrinkage can skew measurements, possibly leading to underestimation of the fiber diameter. As such, we interpret the measured diameter as a lower bound of the true fiber diameter. To establish an upper bound, binarized images were processed in Matlab with a watershed transform to skeletonize the domain and eroded using ImageJ to reestablish a uniformly thin extracellular space. This processing effectively swells the fibers to fill the domain and the Feret diameter was again measured. Lower and upper bounds of the muscle fiber volume fraction were also computed using this approach.

The histology bounds of the muscle fiber diameters and volume fractions for both ROIs are reported in Table II along with the GP model’s estimated mean and confidence intervals for both muscle fiber diameter and volume fraction and the RPBM’s fiber diameter estimates (the RPBM does not consider volume fraction). The fiber diameter estimate of the RPBM agrees with the histology bounds for ROI #2 but substantially underestimated fiber diameter for ROI #1. The fiber diameter estimates of the GP model, in contrast, fall within the range

TABLE II: Diameter and volume fraction measurements of the two ROIs from histology, the GP model estimates with 95% confidence interval (CI), and the RPBM estimates.

ROI	Fiber diameter			Volume fraction	
	Histology (bounds)	GP \pm 95% CI	RPBM	Histology (bounds)	GP \pm 95% CI
#1	55.2 – 73.9 μm	62.3 \pm 17.0 μm	31.8 μm	0.674 – 0.894	0.894 \pm 0.134
#2	51.0 – 62.0 μm	62.9 \pm 17.0 μm	65.6 μm	0.514 – 0.902	0.883 \pm 0.134

of the bounds provided by histology for both ROIs though there is a relatively wide confidence interval. The GP model's confidence intervals of the ROIs are lower than the voxel-wise results due to the averaging of the dMRI metrics over the ROI increasing the SNR of the data. The upper bound of the histology measurements of volume fraction also matches the GP estimates. However, the wide confidence interval in both ROIs cautions against over-interpreting the agreement as the model is uncertain of its estimates. Overall, these results demonstrate the GP model can provide estimates of the fiber diameter and volume fraction that agree with histological measurements.

V. DISCUSSION

This paper presents a framework for computationally efficient estimation of skeletal muscle microstructural parameters from dMRI that accounts for model uncertainty. Use of Gaussian processes (GP) in the inverse mapping provides not only voxel-wise microstructure estimates but also uncertainty intervals, which increases the utility and interpretability of the model's estimates by identifying when a microstructure estimate can be strongly relied on or when it is likely an arbitrary guess. There are two approaches to interpreting these confidence intervals. The first is adopted here, where a single mean value represents the microstructure of the entire voxel and the confidence intervals are then a measure of uncertainty. The second interpretation treats the predicted Gaussian distribution as representing a distribution of parameters within the voxel (e.g. distribution of fiber diameters). Considering the distribution of microstructures that occur in skeletal muscle, such a perspective may further extend the insights available from GP-based inverse models. However, additional investigation, including *in vivo* experiments, remains necessary to establish ground truth parameter distributions and their relationship to GP model estimates and confidence intervals.

Comparison of the GP model with the RPBM approach over the two ROIs finds the GP model provides estimates of additional microstructural features than the RPBM (five vs three) and is substantially faster to evaluate per voxel (~ 20 ms vs ~ 1.5 s). While the two models generally agree in their estimates of the intracellular diffusion coefficient, they diverge in their estimates of fiber diameter and membrane permeability. For fiber diameter, both models provide similar estimates for ROI #1. Further, these estimates are both in line with, though slightly above, the upper bound of fiber diameter measured from histology. In ROI #2, however, the RPBM estimate of 32 μm is substantially below the histology lower bound of 55 μm , while the GP model estimate of 62 μm is squarely within the histology bounds. This is consistent with a recent simulation study that found the RPBM systematically underestimates muscle fiber

diameters [24]. There is little agreement between the two models in their estimate of membrane permeability, with the RPBM model providing a substantially lower estimate than the GP model. Due to the lack of experimental methods to directly measure sarcolemma membrane permeability, adjudicating the difference in these estimates is not currently possible, leaving the source of this discrepancy unclear and suggesting the need for additional investigation and validation of the two models.

Given noise-free data (Fig. 3a), the GP inverse model accurately inverts the problem for all five microstructural parameters, indicating the GP model learns the underlying data structure and suggesting model accuracy can be increased if higher SNR measurements are acquired. In the presence of noise (Fig. 3b,c), the GP inversion scheme accurately estimates three of the five microstructural features (diameter, intracellular diffusion coefficient, and membrane permeability). Notably, relative performance of different features matches the sensitivity analysis of the forward problem reported in Fig. 2b, with dMRI metrics exhibiting comparatively higher sensitivity to changes in diameter, intracellular diffusion coefficient, and permeability than the two lower-performing features (volume fraction and extracellular diffusion coefficient). This lack of sensitivity suggests experimental measurement noise may drown out the encoded microstructural information from these two features, requiring either higher SNR measurements or exploration of more advanced diffusion encoding sequences for muscle tissue that are more sensitive to these features [69]. We note, however, that our GP model can be straightforwardly extended to incorporate any such advancements due to its basis in the physics-based solution to the forward problem.

The framework developed here has a flexible, modular structure. It consists of multiple independent components, each of which can be refined or exchanged with alternative approaches to improve future iterations. For example, incorporating more sophisticated parameterizations of muscle microstructure will further improve the realism of the forward problem. The diffusion-encoding schemes considered here can also be advanced by increasing the fidelity of the PGSE and STEAM sequence simulations or by considering additional diffusion-encoding sequences such as OGSE [70]. Improving these components would require retraining the forward polynomial meta-model. This is a limitation of the current approach and motivates possible future incorporation of alternative meta-modeling techniques. For example, deep neural network architectures can provide capabilities such as fine-tuning to flexibly incorporate continual improvement of the forward model given improved data [71] or even transfer learning [72] for consideration of different tissue classes such as cancerous tumors [73]. Finally, combining diffusion-encoding sequence

optimization with compressed sensing frameworks may allow further reduction in the diffusion-encoding sequences needed, aiding in clinical feasibility and translatability efforts by further reducing imaging time. Changes in the imaging parameters or the forward model will necessitate retraining the inverse model. However, given the total training time of 12-14 hours, such a limitation is not overly prohibitive. Future improvements may additionally consider faster computers and introducing input transformation layers that can accommodate arbitrary imaging parameters.

VI. CONCLUSION

Overall, this work provides a flexible, modular framework for development of physics-inspired, data-driven, and uncertainty-aware inversion models for the estimation of tissue microstructure from dMRI. To reduce the computational expense associated with direct numerical simulations of dMRI physics, a polynomial meta-model is developed that accurately represents the numerical model and is used to develop a Gaussian process regression model to provide voxel-wise microstructural estimates and confidence intervals. The proposed methodology is broadly applicable to additional classes of biological tissues such as neural and cancer tissues, extending its potential impact. Applied here to skeletal muscle, its experimental implementation and validation demonstrate the capability of the framework as a promising non-invasive tool for *in vivo* assessment of skeletal muscle health and organization.

Data Availability: Code and model weights for the meta-model and Gaussian process regression models presented in this work are available online at <https://github.com/nmnaughton/dMRI-muscle-inverse>.

REFERENCES

- [1] R. T. Kell, G. Bell, and A. Quinney, "Musculoskeletal fitness, health outcomes and quality of life," *Sports Medicine*, vol. 31, no. 12, pp. 863–873, 2001.
- [2] D. Samuel, P. Rowe, V. Hood, and A. Nicol, "The relationships between muscle strength, biomechanical functional moments and health-related quality of life in non-elite older adults," *Age and Ageing*, vol. 41, no. 2, pp. 224–230, 2011.
- [3] R. Roubenoff, "Origins and clinical relevance of sarcopenia," *Canadian Journal of Applied Physiology*, vol. 26, no. 1, pp. 78–89, 2001.
- [4] D. T. Villareal, M. Banks, C. Siener, D. R. Sinacore, and S. Klein, "Physical frailty and body composition in obese elderly men and women," *Obesity Research*, vol. 12, no. 6, pp. 913–920, 2004.
- [5] P. P. Purslow, "The structure and functional significance of variations in the connective tissue within muscle," *Comparative Biochemistry and Physiology Part A: Molecular & Integrative Physiology*, vol. 133, no. 4, pp. 947–966, 2002.
- [6] R. Correa-de Araujo, M. O. Harris-Love, I. Miljkovic, M. S. Fragala, B. W. Anthony, and T. M. Manini, "The need for standardized assessment of muscle quality in skeletal muscle function deficit and other aging-related muscle dysfunctions: A symposium report," *Frontiers in Physiology*, vol. 8, p. 87, 2017.
- [7] S. B. Heymsfield, M. C. Gonzalez, J. Lu, G. Jia, and J. Zheng, "Skeletal muscle mass and quality: Evolution of modern measurement concepts in the context of sarcopenia," *Proceedings of the Nutrition Society*, vol. 74, no. 4, pp. 355–366, 2015.
- [8] B. H. Wokke *et al.*, "Quantitative MRI and strength measurements in the assessment of muscle quality in Duchenne muscular dystrophy," *Neuromuscular Disorders*, vol. 24, no. 5, pp. 409–416, 2014.
- [9] I. O. Jelescu and M. D. Budde, "Design and validation of diffusion mri models of white matter," *Frontiers in physics*, vol. 5, p. 61, 2017.
- [10] D. C. Alexander, T. B. Dyrby, M. Nilsson, and H. Zhang, "Imaging brain microstructure with diffusion mri: practicality and applications," *NMR in Biomedicine*, vol. 32, no. 4, p. e3841, 2019.
- [11] J. Polgar, M. Johnson, D. Weightman, and D. Appleton, "Data on fibre size in thirty-six human muscles: an autopsy study," *Journal of the neurological sciences*, vol. 19, no. 3, pp. 307–318, 1973.
- [12] J. M. Edgar and I. R. Griffiths, "White matter structure: a microscopist's view," in *Diffusion Mri*. Elsevier, 2009, pp. 74–103.
- [13] C. S. Landis *et al.*, "Equilibrium transcytolemmal water-exchange kinetics in skeletal muscle *in vivo*," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 42, no. 3, pp. 467–478, 1999.
- [14] D. S. Grebenkov, "Nmr survey of reflected brownian motion," *Reviews of Modern Physics*, vol. 79, no. 3, p. 1077, 2007.
- [15] S. N. Hwang, C.-L. Chin, F. W. Wehrli, and D. B. Hackney, "An image-based finite difference model for simulating restricted diffusion," *Magnetic Resonance in Medicine*, vol. 50, no. 2, pp. 373–382, 2003.
- [16] J. Xu, M. D. Does, and J. C. Gore, "Quantitative characterization of tissue microstructure with temporal diffusion spectroscopy," *Journal of Magnetic Resonance*, vol. 200, no. 2, pp. 189–197, 2009.
- [17] L. Beltrachini, Z. A. Taylor, and A. F. Frangi, "A parametric finite element solution of the generalised Bloch-Torrey equation for arbitrary domains," *Journal of Magnetic Resonance*, vol. 259, pp. 126–134, 2015.
- [18] V.-D. Nguyen, J. Jansson, J. Hoffman, and J.-R. Li, "A partition of unity finite element method for computational diffusion MRI," *Journal of Computational Physics*, vol. 375, pp. 271–290, 2018.
- [19] M. G. Hall and C. A. Clark, "Diffusion in hierarchical systems: A simulation study in models of healthy and diseased muscle tissue," *Magnetic Resonance in Medicine*, vol. 78, no. 3, pp. 1187–1198, 2017.
- [20] D. B. Berry, B. Regner, V. Galinsky, S. R. Ward, and L. R. Frank, "Relationships between tissue microstructure and the diffusion tensor in simulated skeletal muscle," *Magnetic Resonance in Medicine*, vol. 80, no. 1, pp. 317–329, 2018.
- [21] N. M. Naughton, C. G. Tennyson, and J. G. Georgiadis, "Lattice Boltzmann method for simulation of diffusion magnetic resonance imaging physics in multiphase tissue models," *Physical Review E*, vol. 102, no. 4, p. 043305, 2020.
- [22] M. G. Hall and C. A. Clark, "Diffusion in hierarchical systems: A simulation study in models of healthy and diseased muscle tissue," *Magnetic Resonance in Medicine*, vol. 78, no. 3, pp. 1187–1198, 2017.
- [23] J. Bates, I. Teh, D. McClymont, P. Kohl, J. E. Schneider, and V. Grau, "Monte Carlo simulations of diffusion weighted MRI in myocardium: Validation and sensitivity analysis," *IEEE Transactions on Medical Imaging*, vol. 36, no. 6, pp. 1316–1325, 2017.
- [24] D. B. Berry, E. K. Englund, V. Galinsky, L. R. Frank, and S. R. Ward, "Varying diffusion time to discriminate between simulated skeletal muscle injury models using stimulated echo diffusion tensor imaging," *Magnetic resonance in medicine*, vol. 85, no. 5, pp. 2524–2536, 2021.
- [25] N. M. Naughton and J. G. Georgiadis, "Global sensitivity analysis of skeletal muscle dmri metrics: Effects of microstructural and pulse parameters," *Magnetic Resonance in Medicine*, vol. 83, no. 4, pp. 1458–1470, 2020.
- [26] G. J. Stanisz, G. A. Wright, R. M. Henkelman, and A. Szafer, "An analytical model of restricted diffusion in bovine optic nerve," *Magnetic Resonance in Medicine*, vol. 37, no. 1, pp. 103–111, 1997.
- [27] T. Saotome, M. Sekino, F. Eto, and S. Ueno, "Evaluation of diffusional anisotropy and microscopic structure in skeletal muscles using magnetic resonance," *Magnetic Resonance Imaging*, vol. 24, no. 1, pp. 19–25, 2006.
- [28] L. Laghi, L. Venturi, N. Dellarosa, and M. Petracci, "Water diffusion to assess meat microstructure," *Food Chemistry*, vol. 236, pp. 15–20, 2017.
- [29] D. C. Karampinos, K. F. King, B. P. Sutton, and J. G. Georgiadis, "Myofiber ellipticity as an explanation for transverse asymmetry of skeletal muscle diffusion MRI *in vivo* signal," *Annals of Biomedical Engineering*, vol. 37, no. 12, pp. 2532–2546, 2009.
- [30] C. J. Galbán, S. Maderwald, K. Uffmann, and M. E. Ladd, "A diffusion tensor imaging analysis of gender differences in water diffusivity within human skeletal muscle," *NMR in Biomedicine*, vol. 18, no. 8, pp. 489–498, 2005.
- [31] S. Kim, G. Chi-Fishman, A. S. Barnett, and C. Pierpaoli, "Dependence on diffusion time of apparent diffusion tensor of *ex vivo* calf tongue and heart," *Magnetic Resonance in Medicine*, vol. 54, no. 6, pp. 1387–1396, 2005.
- [32] N. Naughton and J. Georgiadis, "Comparison of two-compartment exchange and continuum models of dMRI in skeletal muscle," *Physics in Medicine & Biology*, vol. 64, no. 15, p. 155004, 2019.
- [33] D. S. Novikov, E. Fieremans, J. H. Jensen, and J. A. Helpert, "Random walks with barriers," *Nature Physics*, vol. 7, no. 6, pp. 508–514, 2011.

- [34] E. Fieremans, G. Lemberskiy, J. Veraart, E. E. Sigmund, S. Gyftopoulos, and D. S. Novikov, "In vivo measurement of membrane permeability and myofiber size in human muscle using time-dependent diffusion tensor imaging and the random permeable barrier model," *NMR in Biomedicine*, vol. 30, no. 3, p. e3612, 2017.
- [35] K. V. Winters, O. Reynaud, D. S. Novikov, E. Fieremans, and S. G. Kim, "Quantifying myofiber integrity using diffusion MRI and random permeable barrier modeling in skeletal muscle growth and Duchenne muscular dystrophy model in mice," *Magnetic Resonance in Medicine*, 2018.
- [36] H. C. Torrey, "Bloch equations with diffusion terms," *Physical Review*, vol. 104, no. 3, p. 563, 1956.
- [37] I. Drobnjak, B. Siow, and D. C. Alexander, "Optimizing gradient waveforms for microstructure sensitivity in diffusion-weighted MR," *Journal of Magnetic Resonance*, vol. 206, no. 1, pp. 41–51, 2010.
- [38] E. O. Stejskal and J. E. Tanner, "Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient," *The Journal of Chemical Physics*, vol. 42, no. 1, pp. 288–292, 1965.
- [39] J. E. Tanner, "Use of the stimulated echo in NMR diffusion studies," *The Journal of Chemical Physics*, vol. 52, no. 5, pp. 2523–2526, 1970.
- [40] J. N. Rose, S. Nielles-Vallespin, P. F. Ferreira, D. N. Firmin, A. D. Scott, and D. J. Doorly, "Novel insights into in-vivo diffusion tensor cardiovascular magnetic resonance using computational modeling and a histology-based virtual microstructure," *Magnetic Resonance in Medicine*, vol. 81, no. 4, pp. 2759–2773, 2019.
- [41] A. Barmpoutis and B. C. Vemuri, "A unified framework for estimating diffusion tensors of any order with symmetric positive-definite constraints," in *Biomedical Imaging: From Nano to Macro, 2010 IEEE International Symposium on*. IEEE, 2010, pp. 1385–1388.
- [42] N. Naughton and J. Georgiadis, "Connecting diffusion MRI to skeletal muscle microstructure: Leveraging meta-models and GPU-acceleration," in *Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (learning)*. ACM, 2019, p. 7.
- [43] R. R. Barton and M. Meckesheimer, "Metamodel-based simulation optimization," *Handbooks in operations research and management science*, vol. 13, pp. 535–574, 2006.
- [44] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nature Reviews Physics*, vol. 3, no. 6, pp. 422–440, 2021.
- [45] J. Wang, T. Li, F. Cui, C.-Y. Hui, J. Yeo, and A. T. Zehnder, "Metamodeling of constitutive model using gaussian process machine learning," *Journal of the Mechanics and Physics of Solids*, vol. 154, p. 104532, 2021.
- [46] M. Eldred, "Recent advances in non-intrusive polynomial chaos and stochastic collocation methods for uncertainty analysis and design," in *50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2009.
- [47] A. Kaintura, T. Dhaene, and D. Spina, "Review of polynomial chaos-based methods for uncertainty quantification in modern integrated circuits," *Electronics*, vol. 7, no. 3, p. 30, 2018.
- [48] D. Xiu and G. E. Karniadakis, "The Wiener–Askey polynomial chaos for stochastic differential equations," *SIAM Journal on Scientific Computing*, vol. 24, no. 2, pp. 619–644, 2002.
- [49] I. M. Sobol, "On the distribution of points in a cube and the approximate evaluation of integrals," *Zhurnal Vychislitel'noi Matematiki i Matematicheskoi Fiziki*, vol. 7, no. 4, pp. 784–802, 1967.
- [50] J. Feinberg and H. P. Langtangen, "Chaospy: An open source tool for designing methods of uncertainty quantification," *Journal of Computational Science*, vol. 11, pp. 46–57, 2015.
- [51] S. Arridge, P. Maass, O. Öktem, and C.-B. Schönlieb, "Solving inverse problems using data-driven models," *Acta Numerica*, vol. 28, pp. 1–174, 2019.
- [52] G. Ongie, A. Jalal, C. A. Metzler, R. G. Baraniuk, A. G. Dimakis, and R. Willett, "Deep learning techniques for inverse problems in imaging," *IEEE Journal on Selected Areas in Information Theory*, vol. 1, no. 1, pp. 39–56, 2020.
- [53] J. Gawlikowski *et al.*, "A survey of uncertainty in deep neural networks," *Artificial Intelligence Review*, vol. 56, no. Suppl 1, pp. 1513–1589, 2023.
- [54] A. F. Psaros, X. Meng, Z. Zou, L. Guo, and G. E. Karniadakis, "Uncertainty quantification in scientific machine learning: Methods, metrics, and comparisons," *Journal of Computational Physics*, vol. 477, p. 111902, 2023.
- [55] O. Bousquet, U. von Luxburg, and G. Rätsch, *Advanced Lectures on Machine Learning: ML Summer Schools 2003, Canberra, Australia, February 2-14, 2003, Tübingen, Germany, August 4-16, 2003, Revised Lectures*. Springer, 2011, vol. 3176.
- [56] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006.
- [57] D. Duvenaud, J. Lloyd, R. Grosse, J. Tenenbaum, and G. Zoubin, "Structure discovery in nonparametric regression through compositional kernel search," in *International Conference on Machine Learning*. PMLR, 2013, pp. 1166–1174.
- [58] D. S. Grebenkov, "Use, misuse, and abuse of apparent diffusion coefficients," *Concepts in Magnetic Resonance Part A: An Educational Journal*, vol. 36, no. 1, pp. 24–35, 2010.
- [59] P. Porcari, M. G. Hall, C. A. Clark, E. Greally, V. Straub, and A. M. Blamire, "The effects of ageing on mouse muscle microstructure: a comparative study of time-dependent diffusion MRI and histological assessment," *NMR in Biomedicine*, vol. 31, no. 3, p. e3881, 2018.
- [60] —, "Time-dependent diffusion mri as a probe of microstructural changes in a mouse model of duchenne muscular dystrophy," *NMR in Biomedicine*, vol. 33, no. 5, p. e4276, 2020.
- [61] B. M. Damon, "Effects of image noise in muscle diffusion tensor (DT)-MRI assessed using numerical simulations," *Magnetic Resonance in Medicine*, vol. 60, no. 4, pp. 934–944, 2008.
- [62] GPy, "GPy: A Gaussian process framework in python," <http://github.com/SheffieldML/GPy>, since 2012.
- [63] N. Hansen, S. Muller, and P. Koumoutsakos, "Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES)," *Evolutionary Computation*, vol. 11, no. 1, pp. 1–18, 2003.
- [64] M. Jenkinson, C. F. Beckmann, T. E. Behrens, M. W. Woolrich, and S. M. Smith, "FSL," *Neuroimage*, vol. 62, no. 2, pp. 782–790, 2012.
- [65] C. J. Galbán, S. Maderwald, K. Uffmann, A. de Greiff, and M. E. Ladd, "Diffusive sensitivity to muscle architecture: a magnetic resonance diffusion tensor imaging study of the human calf," *European Journal of Applied Physiology*, vol. 93, no. 3, pp. 253–262, 2004.
- [66] N. Naughton, A. Wang, and J. Georgiadis, "Fascicle ellipticity as an explanation of transverse anisotropy in diffusion MRI measurements of skeletal muscle," in *Proceedings of the 27th Annual Meeting of the International Society for Magnetic Resonance in Medicine*, Montreal, Canada, 2019, abstract 1276.
- [67] E. E. Sigmund *et al.*, "Time-dependent diffusion in skeletal muscle with the random permeable barrier model (RPBM): Application to normal controls and chronic exertional compartment syndrome patients," *NMR in Biomedicine*, vol. 27, no. 5, pp. 519–528, 2014.
- [68] C. A. Schneider, W. S. Rasband, and K. W. Eliceiri, "NIH Image to ImageJ: 25 years of image analysis," *Nature Methods*, vol. 9, no. 7, p. 671, 2012.
- [69] D. Berry, V. Galinsky, E. Hutchinson, J. Galons, S. Ward, and L. Frank, "Double pulsed field gradient diffusion mri to assess skeletal muscle microstructure," *Magnetic Resonance in Medicine*, vol. 90, no. 4, pp. 1582–1593, 2023.
- [70] M. D. Does, E. C. Parsons, and J. C. Gore, "Oscillating gradient measurements of water diffusion in normal and globally ischemic rat brain," *Magnetic Resonance in Medicine*, vol. 49, no. 2, pp. 206–215, 2003.
- [71] C. Käding, E. Rodner, A. Freytag, and J. Denzler, "Fine-tuning deep neural networks in continuous learning scenarios," in *Computer Vision—ACCV 2016 Workshops: ACCV 2016 International Workshops, Taipei, Taiwan, November 20-24, 2016, Revised Selected Papers, Part III 13*. Springer, 2017, pp. 588–605.
- [72] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on knowledge and data engineering*, vol. 22, no. 10, pp. 1345–1359, 2009.
- [73] X. Jiang, H. Li, J. Xie, P. Zhao, J. C. Gore, and J. Xu, "Quantification of cell size using temporal diffusion spectroscopy," *Magnetic resonance in medicine*, vol. 75, no. 3, pp. 1076–1085, 2016.