Connecting Diffusion MRI to Skeletal Muscle Microstructure: Leveraging Meta-Models and GPU-acceleration

Noel Naughton

nnaught2@illinois.edu

Mechanical Science and Engineering
University of Illinois at Urbana-Champaign
Urbana, Illinois

John Georgiadis

jgeorgia@iit.edu

Biomedical Engineering
Illinois Institute of Technology
Chicago, Illinois

ABSTRACT

Due to its non-invasive nature, diffusion-weighted MRI (dMRI) has shown promise as a method to quantify skeletal muscle's microstructure; however, connecting the dMRI signal of muscle to the underlying microstructure is difficult. Numerical models of dMRI can parameterize this relationship, but the associated computational expense has prohibited extensive use. Here, efficient numerical methods are presented to address this problem. In particular, a meta-model representation of a lattice Boltzmann model of dMRI is formulated and shown to be both accurate and several orders of magnitude faster to evaluate. It is also demonstrated how such a meta-model can help inform dMRI pulse profile selection in encoding microstructural information into the dMRI signal. Additionally, histologically-informed simulations are performed, allowing comparison of the numerical model's simplified parameterization with the more complex topology of skeletal muscle. Finally, an efficient inversion method is proposed to infer microstructural parameters of muscle from dMRI signal using a GPU-accelerated numerical model. The inversion method is able to infer microstructural parameters from dMRI signal when the underlying geometry matches the numerical model's, however, the simplified numerical model does not agree with simulations of more complex muscle tissue.

CCS CONCEPTS

• **Computing methodologies** → **Modeling methodologies**; *Parallel algorithms*; • **Applied computing** → **Imaging**; *Engineering*.

KEYWORDS

lattice Boltzmann method, skeletal muscle, diffusion MRI, meta-model, GPU-acceleration

ACM Reference Format:

Noel Naughton and John Georgiadis. 2019. Connecting Diffusion MRI to Skeletal Muscle Microstructure: Leveraging Meta-Models and GPU-acceleration. In *PEARC '19: Practice and Experience in Advanced Research Computing, July 28– August 01, 2019, Chicago, IL, USA*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3332186.3333054>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PEARC '19, July 28– August 01, 2019, Chicago, IL, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-7227-5/19/07...\$15.00

<https://doi.org/10.1145/3332186.3333054>

1 INTRODUCTION

Healthy muscle is linked to increased quality of life measures [19, 23], particularly among the aging population [35]. Muscle's functional ability is intimately linked to its structural organization. Muscle has a hierarchical structure leading to a hierarchical multi-scale organization, which ranges from the macro-scale (whole muscle tissue) through the micro-scale (muscle cell organization) down to the nano-scale (actin-myosin cross-linked chains). At the micro-scale, individual muscle cells (myocytes) are elongated fibers that are tightly packed together. Muscle fibers are surrounded by, and connected to, a collagenous extracellular matrix (ECM), which plays an important role in force transmission [32]. Understanding changes in this microstructure is a necessary first step towards developing diagnostic tools, interventional techniques and rehabilitation methods to improve muscle function.

A better measurement of muscle structure will improve our understanding of how muscle functions, especially in instances when autonomous movement is impaired, making measures of functional ability difficult. Such conditions are found during fetal development and for bed-ridden, comatose, or injured individuals. Better measurements will also provide insight into the physiological changes of muscle during aging [13] and muscular dystrophy [41]. Measuring the muscle's microstructure allows combining of functional measures of muscle ability, such as force measurements, with structural measurements, enabling increased understanding of the different phenomena affecting muscle function. Functional measurements of muscle convolve many different effects with the effects of muscle structure [41], such as the motivation of the subject [25], neural-muscular coordinations [16, 37], lipid infiltration [4], metabolic changes [5], and protein expression [20].

By independently measuring muscle's microstructure, functional measures can be normalized by these measurements, enabling deeper understanding of the influence that both microstructure and physiology have on muscle's functional ability. Currently, biopsy and histology are the best methods available to investigate this microstructure, but these ex-vivo methods are invasive, labor-intensive, and only examine a small region of the muscle. A promising alternative to these techniques, however, is diffusion-weighted magnetic resonance imaging (dMRI). In contrast to biopsy and histology, dMRI provides quantitative measures of muscle microstructure in a non-invasive and in-vivo way. It can also measure microstructure throughout the entire muscle as opposed to biopsy's inherently local nature. These advantages enable the possibility of measuring muscle microstructure in situations that heretofore have not been considered feasible.

Diffusion-weighted MRI measures the cumulative effect of water molecule movement due to diffusion. Microstructural barriers within and surrounding individual muscle fibers restrict the free-diffusion of water, leading to anisotropic diffusion in the tissue. Water diffuses faster in the axial direction of the muscle fiber than in the transverse direction where barriers, such as cell membranes, restrict the water molecules movement. The root-mean-squared displacement of the water molecule in each voxel is used to calculate an apparent diffusion tensor, which reflects the average effect of these microstructural restrictions on the self-diffusion of water. Although the resolution limit of clinical dMRI is often $\sim 1 - 2$ mm, the diffusion distance of the water molecules during the measured diffusion time is $\sim 10 \mu\text{m}$. As such, the effects of diffusion barriers measured by the dMRI signal are only related to structures at this ten micron length scale, thus allowing the signal in each voxel to be related to the muscle cell microstructure.

Although the measured dMRI signal is reflective of the underlying microstructure, the relationship with morphological parameters of muscle is not straightforward. Many analytical models that have been successful in explaining diffusion in the brain are not appropriate to use in muscle due to the dense packing and larger sizes of muscle cells compared with axons [26, 27]. The models that do exist ignore the effects of the ECM [7], however, due to the ECM's effect on force transmission, models capable of estimating ECM's properties are desired. This leaves numerical models of dMRI as the best candidates. Numerical models have been developed using an array of methods such as finite difference [15, 34, 43], finite elements [3, 31, 40], Monte Carlo [10], and lattice Boltzmann [29]. Compared with analytical models, numerical models are orders of magnitude more expensive to evaluate, and with five to nine independent parameters, there is limited ability to sample all possible parameter combinations. These constraints motivate the development of methods that maximize the usefulness of every model evaluation in order to allow efficient use of computational resources.

In this paper we present techniques to leverage precomputed datasets and HPC resources, such as multi-core processing and GPU-acceleration, to both gain insight into how muscle microstructure impacts dMRI signal and also to infer microstructural parameters of the muscle from dMRI measurements. These techniques allow efficient utilization of computational resources, while enabling the use of fitting techniques to infer microstructural parameters from numerical model's of dMRI signal that had only been possible with analytical models before. In section 2 we present a numerical model of dMRI in skeletal muscle microstructure. In section 3 we show how the model can be efficiently approximated with a polynomial meta-model, and how this meta-model allows investigations that were computationally prohibitive with the previous model. In section 4 we show how our simplified numerical model compares with more complex and realistic simulations of skeletal muscle. Finally in section 5 we present an inversion scheme to infer microstructural parameters of skeletal muscle from dMRI measurements.

2 DESCRIPTION OF MUSCLE dMRI MODEL

Skeletal muscle has a relatively uniform cross-section geometry (Figure 1). Our model is premised on abstracting this uniform structure as a periodic unit cell, or representative elementary volume

Figure 1: Left panel: Example of skeletal muscle cross-section taken from chicken leg (by permission of E. Perreault lab). Right panel: representative elementary volume (REV) used as the periodic muscle fiber model showing the simplified representation of the extracellular domain (blue), the individual muscle cells (green), and the semi-permeable sarcolemma membrane surrounding them (red).

(REV), with periodic boundary conditions, thus allowing a parsimonious parameterization of the muscle microstructure. The REV consists of hexagonal cylinders packed together and surrounded by a semi-permeable membrane (sarcolemma). The hexagonal cylinders are assumed to be infinitely long, thus reducing the problem to 2-D [29]. Parameterizing the REV yields seven microstructural parameters that can be independently varied. They are the two geometrical parameters of cell diameter and cell volume fraction, as well as the five tissue parameters of membrane permeability, intracellular and extracellular diffusion coefficients, and intracellular and extracellular spin-spin relaxation (T_2) time.

With this parameterized REV, we use the lattice Boltzmann method (LBM) to solve the Bloch-Torrey equation over the REV, which is the governing partial differential equation of diffusion MRI [29]. Our model is written in Fortran 90 with two versions developed. One was written with Message Passing Interface (MPI) to utilize multiple CPUs. This MPI-enabled model has been shown to scale efficiently up to at least 168 cores [29]. A second version of the model was written using the same Fortran code but adding OpenACC directives to allow GPU-acceleration. This GPU-accelerated version of the model was found to be approximately 8x faster than a single core (Xeon E5-2600v4 2.20 GHz) implementation of the code when run on Nvidia Pascal Titan X GPUs.

To simulate the dMRI signal, we use the pulsed-gradient spin echo (PGSE) pulse, which consists of two equal and opposite pulses separated by a diffusion time. The PGSE pulse is a traditional diffusion-weighting pulse [38] that can be parameterized by gradient strength (g), diffusion time (Δ) and gradient duration, which was held constant at 5 ms for all simulations. dMRI measures the diffusion of water in the direction of the applied gradient. Multiple gradient directions are used to measure the effect of microstructure on the diffusion of water in each direction. These measurements are used to calculate an anisotropic apparent diffusion tensor using algorithms from the fanDTasia Matlab toolbox [1] rewritten in Python 3. This apparent diffusion tensor represents the cumulative impact of microstructural restrictions on the free diffusion of water;

Figure 2: Relative error of meta-model compared to underlying LBM model for different order polynomials.

its deviation from water’s free-diffusion coefficient in a particular direction is indicative of the microstructural restrictions in that direction [2, 11]. From this tensor, six dMRI metrics are calculated, which describe the principle directions of the tensor ($\lambda_1, \lambda_2, \lambda_3$), the mean diffusivity ($MD = (\lambda_1 + \lambda_2 + \lambda_3)/3$), the radial diffusivity ($RD = (\lambda_2 + \lambda_3)/2$) and the fractional anisotropy (FA), which is a measure of the total anisotropy of the tensor.

3 POLYNOMIAL META-MODEL

3.1 Meta-Model Creation and Accuracy

A previous parameter sensitivity study of the computational model generated a dataset of 100,000 model evaluations. The study examined seven microstructural parameters and two pulse parameters. These nine parameters were sampled using a Sobol low-discrepancy sampling method [12] and the apparent diffusion tensor was calculated from the dMRI signal. Six different tensor invariants were calculated from the apparent diffusion tensor (FA, MD, RD, λ_1 , λ_2 and λ_3). The simulations were performed on SDSC’s *Comet* cluster [39]. 40 nodes were used to run three simulations simultaneously on each node, with each simulation using eight cores. The entire study took approximately 7000 core hours with a median wall time of 10 seconds per simulation; however, the longest simulation took 6.4 hours. This discrepancy in simulation times is due to the resolution requirements for different cell volume fractions in the REV. The LBM model involves space discretization with a Cartesian grid. As the volume fraction increases, the extracellular space decreases. It was found that to reach a minimum accuracy threshold, there needed to be at least four nodes across the extracellular domain regardless of size. For small cell diameters and large volume fractions this condition requires smaller spatial discretizations leading to a larger number of nodes and longer solution times. Additionally, in the LBM model, the space and time discretization is coupled via $dx^2 \sim dt$. This relationship means that for a 2-D domain, halving the spatial discretization leads to a 16x increase in run time.

These long simulation times motivate the pursuit of more efficient solution techniques, such as a meta-model. A meta-model, or surrogate model, is a model of a model. To construct the meta-model, the relationship between the inputs and outputs of the model is examined with no regard for the underlying physics. The advantage

Table 1: Run time of models in CPU seconds. LBM-REV model was run in parallel with 8 CPUs leading to a median wall time of 10.13 seconds.

Model	Median Run Time
3rd order meta-model	0.046 CPU seconds
4th order meta-model	0.091 CPU seconds
5th order meta-model	0.233 CPU seconds
6th order meta-model	0.587 CPU seconds
7th order meta-model	1.343 CPU seconds
LBM-REV model	81.03 CPU seconds

of meta-models is that they are often much faster to evaluate, thus allowing fast approximate solutions. The dataset from the previous sensitivity study was used to fit a polynomial meta-model to relate the input microstructural and pulse parameters with the calculated diffusion metrics using the open-source *ChaosPy* Python package [6]. Because the input parameters were sampled from a uniform distribution, the relationship between the inputs and outputs can be modeled with a basis set of Legendre polynomials [42]. This basis set used to interpolate the data using a least-squares method to form a 9-dimensional, n^{th} order polynomial for each of the six diffusion metrics. To investigate the effect of different degree polynomials, n was varied from 3 to 7. 75% of the dataset was used to fit the polynomial, with 25% reserved to test the accuracy of the fit. Because the dataset was generated via a Sobol sampling method, which is a low discrepancy sampling method, the samples used to construct the model are evenly distributed across the input space. Additionally, the testing dataset is extracted from the same sequence as the fitting dataset, so it is not only uniformly distributed, it also has low discrepancy with the training dataset. This means the testing dataset will be sampled at points far away from the training dataset, allowing good confidence that these results are an upper limit of error.

Before using the meta-model, it is important to determine its accuracy. The meta-model was compared with the 25,000 LBM evaluations that were not used to create it for four dMRI metrics (FA, MD, RD, and λ_1). The LBM-REV model produces transversely isotropic diffusion leading to minimal difference between RD, λ_2 and λ_3 , so only RD was considered. The relative error between the meta-model and the underlying LBM model is shown in Figure 2. For increasing polynomial order, the accuracy of the meta-model improves, though the accuracy does not substantially increase beyond 4th order. The relative error is normalized by the average value of the LBM dataset to avoid over-emphasizing differences between small values. These results suggest that, as a whole, the meta-model is an accurate representation of the underlying LBM model and relationships derived from it should also be true to the underlying LBM model. The meta-model is substantially faster to evaluate, with a median evaluation time between 0.04 and 1.3 CPU seconds, depending on the polynomial order, compared to 81 CPU seconds for the LBM model (Table 1). The majority of the cost of evaluating the meta-model is caused by the overhead of accessing the polynomial. If multiple parameter sets are passed to the meta-model at once, the average evaluation time decreases by a factor of

Figure 3: Selection of pulse parameters. Green boxes represent the subset of pulse profiles that were selected for by each dMRI metric. Blue boxes represent pulse profiles that were selected by at least one, but not all, dMRI metrics. Red boxes represent the five final pulse profile parameters. Only pulse parameter combinations that were within the allowed b-range are shown.

200. Additionally, the analytical nature of the meta-model allows computing derivatives, which may aid in local sensitivity analysis and model fitting using gradient descent methods. As the polynomial order increases, the accuracy of the meta-model increases, but so too does the evaluation time. For the rest of this paper, we consider the 6th order polynomial meta-model, as it has a good combination of accuracy and short evaluation time.

3.2 PGSE Pulse Selection

Because the meta-model can be evaluated much more efficiently than the underlying LBM model, questions that would have been computationally prohibitive to answer using the original LBM model can now be addressed. One such question is which PGSE pulse profiles can encode the most information about the microstructure? When measuring the apparent diffusion tensor, the shape of the diffusion-weighted pulse is important [9]. For the PGSE pulse in particular, two key parameters are the gradient strength and diffusion time. Both enable the dMRI signal to encode different aspects of the microstructure. For example, long diffusion times allow water molecules to travel further distances, thus sensitizing them to structures at larger length scales [7]. By utilizing multiple pulse combinations, it is possible to increase the amount of microstructural information that can be extracted from the combined dMRI signals. While it is broadly known how different pulse profiles affect the signal [9], it is not known which combinations of pulses encode the most structural information about the tissue. Because the use of additional pulses increases imaging time, encoding the most possible information into the fewest possible pulses is a priority.

Figure 4: Example simulations of histology image and corresponding periodic REV with same microstructural values.

In the meta-model, there are nine independent inputs. Seven inputs are related to the microstructure and are intrinsic to the tissue. The other two are related to the PGSE pulse profile (gradient and diffusion time) and can be controlled by the MRI operator. To determine how different pulse profiles affect the signal, an $N \times N$ array consisting of combinations of the two pulse parameters was defined and the meta-model was sampled at each pulse profile for multiple combinations of microstructural parameters. For a uniform sampling (M) of each of the seven microstructural parameters, M^7 model evaluations are required at each point on the $N \times N$ array leading to $M^7 N^2$ evaluations. These evaluations were used to determine which combinations of pulse profiles give the most information about the microstructure using two measures: the variance of values within each vector and the covariance between the vectors.

If a pulse profile has a low variance, it will yield similar signal values for any microstructural combination, making it difficult to extract unique information about underlying microstructure. If two vectors have a high covariance, they do not contain independent information, leading to redundant measurements. After sampling the meta-model, the results were filtered to only examine pulse profiles with a b-value between 200 and 1500 s/mm^2 . The b-value is a measure of the cumulative magnetization of the pulse and is a useful predictor of how much signal attenuation is expected. This b-range is typical of what is feasible in clinical imaging experiments. For each dMRI metric, the vector with the largest variance was selected, and the 25 vectors with the lowest average correlation to already selected vectors were iteratively added. This resulted in 31 pulse profiles, 19 of which were selected for by all four dMRI metrics. From these 19, the five pulse profiles with the lowest average correlation coefficient were selected. The pulse profiles selected and the relative variance of all pulse profiles is shown in Figure 3.

It is here that the value of the meta-model, compared to the original LBM model, becomes apparent. The meta-model was evaluated on a 20×20 array of pulse profiles with 12 uniform samples for each of the seven microstructural parameters. This created a set of 20^2 vectors each with 12^7 elements requiring 14×10^9 model evaluations. For a median LBM model run-time of 81 CPU seconds, such an investigation would be prohibitively expensive (36,800 CPU years).

4 HISTOLOGY-BASED SIMULATIONS

The LBM model's REV geometry is a necessary abstraction of muscle's complex topology in order to allow parameterization of the geometry. To test how well it represents actual muscle structure, three

Figure 5: Relative error of periodic REV model compared to simulations of histological informed geometries using.

histological images of skeletal muscle were procured [14, 17, 24] and manually segmented into intra- and extracellular domains. Different diffusion coefficients and permeability values were used to calculate a range of possible scenarios, while the diameter was measured from the histology image with the Feret diameter plugin in ImageJ [36] and the volume fraction was determined from the segmented image. These histologically informed domains were used with the LBM model to solve for the dMRI signal using 21 gradient directions for each of the five pulse profiles identified with the meta-model. For each pulse, the apparent diffusion tensor was calculated and the dMRI metrics were calculated. Simulations over an REV with the same microstructural parameters were also performed.

There exist preferential directions in terms of cell shape and organization in the histology images. This lead to transverse anisotropy of the apparent diffusion tensor ($\lambda_2 > \lambda_3$), something that the current periodic REV does not reproduce due to its symmetries. dMRI of muscle does contain transverse anisotropy [8], and some authors have hypothesized that it is related to elliptical cross-sections of muscle cells [18]; however, this has not been proven conclusively and other evidence suggests that the anisotropy may originate from effects at larger length scales [30]. Because the REV was unable to represent this anisotropy, an attempt was made to minimize local geometric effects of the histology images by simulating the domain in its original orientation as well as rotated 90° . The resulting signals were averaged using only the middle half of the domain's signal in each direction to avoid edge effects where the geometry does not match across the periodic boundary.

Figure 4 compares the domain maps of a histologically informed geometry with a simplified REV geometry with the same parameters. The extracellular space of the histology is much more tortuous and uneven than the consistent thickness of the REV's extracellular space. The error of the REV in representing the dMRI signal of the histological images is shown in Figure 5. λ_1 is most accurately reproduced by the current REV, while the other five metrics are substantially less accurate. Because λ_1 represents diffusion in the axial direction, where there are fewer restrictions to free diffusion, it is closely related with the volumetric average of the two diffusion

Table 2: Range of input parameters used in meta-model and inversion scheme. Parameters used in the meta-model but not during the inversion scheme are marked by *.

Parameter	Range
Muscle Cell Diameter	10 - 80 μm
Volume Fraction	0.70 - 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 $\mu\text{m}^2/\text{ms}$
*Intracellular T_2	20 - 40 ms
*Extracellular T_2	80 - 140 ms
*Gradient Strength	10 - 80 mT/m
*Diffusion Time	10 - 750 ms

coefficients, which the REV is capable of reproducing. In contrast, the other dMRI metrics are additionally influenced by diffusion in the transverse plane, where differences in the cell geometry between the histological image and the REV affect the signal. Most of the disagreement between the histology images and the REV comes from dMRI pulses with a long diffusion time. In these cases, differences in the extracellular topology lead to large dMRI signal differences.

More work is needed to improve the ability of the REV (or some similarly reduced order, parameterized model) to match the expected signal from more complex muscle tissue. Additionally, the employed histology images have a small field of view ($\sim 200 \mu\text{m}$ a side). Using images with a field of view on the order of a dMRI voxel ($\sim 1 - 2 \text{ mm}$) may eliminate some of the local geometrical artifacts; however, solving over such a large domain while still maintaining spatial resolution to resolve muscle structure ($dx \sim 1 \mu\text{m}$) will require increased HPC resources.

5 INVERSION SCHEME

Up until this point we have examined the forward problem, which consists of solving for the dMRI signal starting from a known microstructural parameter set. However, one does not typically have this microstructural information. Rather, the dMRI signal is measured with the goal of determining the microstructure that generated it. This inversion of the forward problem is challenging because there is no direct way to model this inverse relationship. Instead, a search algorithm is implemented to try and find a microstructural parameter set that corresponds to the measured dMRI signal.

In a previous sensitivity study [28], the effect of two of the seven microstructural parameters was found to be negligible (intra- and extracellular T_2 time) and are not considered here, leaving five microstructural parameters to be the unknowns for the inverse problem. Using the five pulse profiles previously selected, a candidate solution vector of the six calculated dMRI metrics from each pulse was constructed. This candidate solution vector was compared with the reference vector of dMRI data, from which we wish to infer microstructural parameters. An objective function consisting of the normalized L2 error between the reference vector and the candidate solution vector was defined. For each candidate solution,

Table 3: Average L2 error of objective function for inversion attempts of three different test cases.

	Muscle Cell Diameter	Volume Fraction	Membrane Permeability	Intracellular Diffusion	Extracellular Diffusion
REV with no noise	2.46%	11.3%	3.73%	3.40%	9.97%
REV with SNR = 50	10.6%	11.6%	12.6%	10.73%	13.3%
Histology-based images	14.1%	44.4%	27.5%	7.14%	15.56%

all five pulse parameters were simulated for an REV constructed from microstructural parameters of the candidate solution.

The objective function was minimized using a fitting procedure based on the method developed by Knysh and Korkolis [21], which consists of a global search followed by a local search. During the global search, 50 model evaluations are performed using a Latin hypercube to examine the entire input space. The algorithm was modified to save the results of every model evaluation during this global search. At the beginning of each new case, previously saved iterations were loaded to allow them to be reused for the global search. This allowed an initial set of solutions to be available without needing to be recomputed. As more global searches are performed, this list of previous solutions grows, allowing for fewer simulations to be required during the global search portion of the algorithm. The local search used radial basis functions to construct a response surface from a modified CORS algorithm [33]. 100 model evaluations were performed during this local search method. At the end of the local search, the best candidate solution was returned. Microstructural parameters were restricted to the range of values given in Table 2. The global and local searches each allow for multiple simultaneous function evaluations enabling parallel computation. LBM simulations were performed using GPU-accelerated code on a local server with 8 Nvidia TITAN X GPUs and on two of SDSU's *Comet* GPU nodes with 4 Nvidia Pascal P100 GPUs [39]. Five simulations were performed on each GPU in parallel allowing for 40 or 20 concurrent simulations at a time, respectively.

Three different cases were examined. The first consisted of simulations of a periodic REV with no noise, the second was a periodic REV with noise added at SNR=50, and the third case was noise-free simulations of histology images. For the REV, four different microstructural cases were generated. For each of the three histology cases, four microstructural cases were generated with different permeability and diffusion coefficients (diameter and volume fraction are fixed by the image). Parameters for these cases are given in Table A1. As shown in section 4, comparing the periodic REV results with the histology simulations for the forward problem indicates that the former is not a realistic representation of the histology images. By examining how well the model can fit simulations for the REV and the REV with noise, we can gain an understanding of how well the dMRI signal can be inverted without the issues of the mismatch between the REV and actual muscle tissue. This leaves the problem to that of improving the REV's ability to represent the muscle signal. If an appropriate reduced-order computational model of muscle can be developed, it should be possible to invert the signal in order to infer microstructural parameters.

Results of the fitting procedure are shown in Table 3. The most accurate results are for the periodic REV geometry with no added

noise while the REV geometry with added noise had approximately twice the relative error. The accuracy of the results from the histologically informed geometries is comparable to that of the periodic REV geometry with noise except for volume fraction and permeability. These parameters were substantially less accurately fit, a result likely due to the inability of the REV to perfectly represent the dMRI signal from the histology images as discussed above. The poor results for the volume fraction are partially due to the method of normalizing results. Results are normalized by the allowable parameter range, which is relatively small for volume fraction (0.25) compared to its absolute value (0.70 - 0.95); such a relationship does not exist for the other microstructural parameters. The success in fitting the REV, even when there is noise, shows that the inverse problem can be solved if the model represents the underlying physics. Improving the ability of the REV to represent the dMRI signal of muscle through better morphological parameterization will improve this ability further.

We focused on developing a scheme to fit the model using parallel computations run on GPUs. A similar method could have been performed using multiple CPUs, or both could be combined allowing for full utilization of available computing resources. The results presented here are generally concerned with inferring microstructural parameters from a single voxel, however, future work will focus on the scalability of the method in terms of extracting parameters for every voxel of a dMRI image. The fitting method took approximately 45 minutes per case. Much of this time was computing in domains that had small diameters and high volume fractions. Although the extracellular domains of these cases require the increased spatial discretization, the accuracy of the intracellular domain does increase substantially. This leads to wasted resources as the increased resolution in the intracellular space (up to 90% of the domain) is not necessary. Implementation of multi-domain mesh refinement [22] could substantially reduce these simulation times by only refining the mesh over the extracellular domain where the increased resolution is necessary. Additionally, future improvements of the algorithm could use the meta-model for the global search to establish regions of potential candidate solutions before the more expensive, but accurate, LBM model is used for a local search. These improvements would yield faster evaluation times, enabling either more model evaluations during the fitting procedure or higher throughput of individual cases, something necessary if an entire dMRI image is to be examined.

Future work will also examine how machine learning, in particular deep learning, can aid in increasing accuracy and speed of the inversion process. Implementing machine learning is, in some ways, essentially just expanding the meta-model to be a more complicated response surface that does not have an analytical expression.

Machine learning methods were initially considered in place of the fitting procedure presented. However, because there are several issues that are extrinsic to the inversion algorithm (in particular how well the REV matches actual muscle tissue), an explicit model fitting method was chosen in order to allow better understanding of the methodological error before adding on potential additional error from a machine learning algorithm.

6 CONCLUSION

The aim of this work is to present techniques used to leverage forward problem simulations to improve the inverse problem consisting of extracting muscle cell microstructure from dMRI signal. These techniques allow for efficient model evaluation, enabling investigations that are too computationally expensive with the underlying models. We showed how it is possible to construct a meta-model of dMRI in muscle, allowing for fast model evaluation, and how such a model can be used to determine dMRI pulse parameters to increase sensitivity to the muscle microstructure. Additionally, we showed that it is possible to solve the inverse problem and recover the microstructural parameters of the muscle from the dMRI signal using a simple REV model. We found that the REV model performs poorly when compared to simulations based on histology images of muscle, suggesting future work is required to improve this REV abstraction approach. The methods presented here alleviate the computational burden of evaluations of a numerical model for dMRI, allowing the interpretation of the dMRI signal of muscle that had previously only been performed with analytical methods.

ACKNOWLEDGMENTS

Funding for this work was provided by an NSF Graduate Research Fellowship for NMN and by the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant ACI-1548562 and provided access to the SDSC Comet Cluster under allocation TG-MCB180044. Additional support was provided by the R.A. Pritzker endowed chair and NSF grant CMI-1762774.

REFERENCES

- [1] Angelos Barmpoutis and Baba C Vemuri. 2010. 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, 1385–1388.
- [2] Peter J Basser, James Mattiello, and Denis LeBihan. 1994. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal* 66, 1 (1994), 259.
- [3] Leandro Beltrachini, Zeike A Taylor, and Alejandro F Frangi. 2015. A parametric finite element solution of the generalised Bloch–Torrey equation for arbitrary domains. *Journal of Magnetic Resonance* 259 (2015), 126–134.
- [4] Matthew J Delmonico, Tamara B Harris, Marjolein Visser, Seok Won Park, Molly B Conroy, Pedro Velasquez-Mieyer, Robert Boudreau, Todd M Manini, Michael Nevitt, Anne B Newman, and Bret H. Goodpaster. 2009. Longitudinal study of muscle strength, quality, and adipose tissue infiltration. *American Journal of Clinical Nutrition* 90, 6 (2009), 1579–1585. <https://doi.org/10.3945/ajcn.2009.28047>
- [5] Richard HT Edwards. 1981. Human muscle function and fatigue. In *Ciba Found Symp*, Vol. 82. Wiley Online Library, 1–18.
- [6] Jonathan Feinberg and Hans Petter Langtangen. 2015. Chaospy: An open source tool for designing methods of uncertainty quantification. *Journal of Computational Science* 11 (2015), 46–57.
- [7] Els Fieremans, Gregory Lemberskiy, Jelle Veraart, Eric E Sigmund, Soterios Gyftopoulos, and Dmitry S Novikov. 2017. 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* 30, 3 (2017), e3612.
- [8] Craig J. Galbán, Stefan Maderwald, Kai Uffmann, Armin de Greiff, and Mark E. Ladd. 2004. Diffusive sensitivity to muscle architecture: a magnetic resonance diffusion tensor imaging study of the human calf. *European Journal of Applied Physiology* 93, 3 (dec 2004), 253–262. <https://doi.org/10.1007/s00421-004-1186-2>
- [9] Denis S Grebenkov. 2010. Use, misuse, and abuse of apparent diffusion coefficients. *Concepts in Magnetic Resonance Part A: An Educational Journal* 36, 1 (2010), 24–35.
- [10] Matt G. Hall and Chris A. Clark. 2017. Diffusion in hierarchical systems: A simulation study in models of healthy and diseased muscle tissue. *Magnetic Resonance in Medicine* 78, 3 (sep 2017), 1187–1198. <https://doi.org/10.1002/mrm.26469>
- [11] Anneriet M Heemskerk and Bruce M Damon. 2007. Diffusion Tensor MRI Assessment of Skeletal Muscle Architecture. *Current medical imaging reviews* 3, 3 (2007), 152–160. <https://doi.org/10.2174/157340507781386988> arXiv:15334406
- [12] Jonathan D Herman and Will Usher. 2017. SALib: An open-source Python library for Sensitivity Analysis. *J. Open Source Software* 2, 9 (2017), 97.
- [13] Steven B. Heymsfield, M. Cristina Gonzalez, Jianhua Lu, Guang Jia, and Jolene Zheng. 2015. Skeletal muscle mass and quality: Evolution of modern measurement concepts in the context of sarcopenia. *Proceedings of the Nutrition Society* 74, 4 (nov 2015), 355–366. <https://doi.org/10.1017/S0029665115000129>
- [14] M.A. Hill. 2017. Embryology Skeletal muscle histology 003.jpg. https://embryology.med.unsw.edu.au/embryology/index.php/File:Skeletal_muscle_histology_003.jpg. Accessed: 2019-04-05.
- [15] Scott N Hwang, Chih-Liang Chin, Felix W Wehrli, and David B Hackney. 2003. An image-based finite difference model for simulating restricted diffusion. *Magnetic Resonance in Medicine* 50, 2 (2003), 373–382.
- [16] F. M. Ivey, B. L. Tracy, J. T. Lemmer, M. NessAiver, E. J. Metter, J. L. Fozard, and Ben F. Hurley. 2000. Effects of strength training and detraining on muscle quality: Age and gender comparisons. *Journals of Gerontology - Series A Biological Sciences and Medical Sciences* 55, 3 (mar 2000), B152–B157. <https://doi.org/10.1093/gerona/55.3.B152>
- [17] R. Jennings and C. Premanandan. 2017. Veterinary Histology Chapter 4: Skeletal Muscle. <https://ohiostate.pressbooks.pub/vethisto/chapter/4-skeletal-muscle/>. Accessed: 2019-04-05.
- [18] Dimitrios C. Karampinos, Kevin F. King, Bradley P. Sutton, and John G. Georgiadis. 2009. Myofiber ellipticity as an explanation for transverse asymmetry of skeletal muscle diffusion MRI in vivo signal. *Annals of Biomedical Engineering* 37, 12 (dec 2009), 2532–2546. <https://doi.org/10.1007/s10439-009-9783-1>
- [19] Robert T Kell, Gordon Bell, and Art Quinney. 2001. Musculoskeletal fitness, health outcomes and quality of life. *Sports Medicine* 31, 12 (2001), 863–873.
- [20] Henrik Klitgaard, M Mantoni, S Schiaffino, S Ausoni, L Gorza, C Laurent-Winter, P Schnohr, and B Saltin. 1990. Function, morphology and protein expression of ageing skeletal muscle: a cross-sectional study of elderly men with different training backgrounds. *Acta Physiologica Scandinavica* 140, 1 (1990), 41–54.
- [21] Paul Knysh and Yannis Korkolis. 2016. Blackbox: A procedure for parallel optimization of expensive black-box functions. *arXiv preprint arXiv:1605.00998* (2016).
- [22] Daniel Lagrava, Orestis Malaspinas, Jonas Latt, and Bastien Chopard. 2012. Advances in multi-domain lattice Boltzmann grid refinement. *J. Comput. Phys.* 231, 14 (2012), 4808–4822.
- [23] Amy E Latimer-Cheung, Lara A Pilutti, Audrey L Hicks, Kathleen A Martin Ginis, Alyssa M Fenuta, K Ann MacKibbin, and Robert W Motl. 2013. Effects of exercise training on fitness, mobility, fatigue, and health-related quality of life among adults with multiple sclerosis: a systematic review to inform guideline development. *Archives of physical medicine and rehabilitation* 94, 9 (2013), 1800–1828.
- [24] Lesela. 2011. Cross section of normal skeletal muscle from mouse quadriceps. <https://www.flickr.com/photos/agaphd/5923238105/in/photostream/>. Accessed: 2018-11-06.
- [25] Jamie S McPhee, James Cameron, Thomas Maden-Wilkinson, Mathew Piasecki, Moi Hoon Yap, David A Jones, and Hans Degens. 2018. The contributions of fibre atrophy, fibre loss, in situ specific force and voluntary activation to weakness in sarcopenia. *The Journals of Gerontology: Series A* 00, 00 (2018), 1–8. <https://doi.org/10.1093/gerona/gly040>
- [26] Noel M Naughton, Nicolas R. Gallo, Aaron T. Anderson, and John G Georgiadis. 2019. Comparison of dMRI Models for Skeletal Muscle Microstructure Estimation with Numerical Simulations and Porcine Phantom. In *Proceedings of the 27th Annual Meeting of the International Society for Magnetic Resonance in Medicine*. ISMRM.
- [27] Noel M Naughton and John G Georgiadis. 2019. Comparison of two-compartment exchange and continuum models of dMRI in skeletal muscle. *Physics in Medicine & Biology* 64, 15 (2019), 155004. <https://doi.org/10.1088/1361-6560/ab2aa6>
- [28] Noel M Naughton and John G Georgiadis. 2019. Global Sensitivity Analysis of Skeletal Muscle dMRI: Effects of Microstructural and Pulse Parameters. *Magnetic Resonance in Medicine (submitted)* (2019).
- [29] Noel M Naughton and John G Georgiadis. 2019. Lattice Boltzmann method for simulation of diffusion magnetic resonance imaging physics in heterogeneous tissue models. *arXiv preprint arXiv:1907.00908* (2019).

- [30] Noel M Naughton, Anthony Wang, and John G Georgiadis. 2019. 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*. ISMRM.
- [31] Van-Dang Nguyen, Johan Jansson, Johan Hoffman, and Jing-Rebecca Li. 2018. A partition of unity finite element method for computational diffusion MRI. *J. Comput. Phys.* 375 (2018), 271–290.
- [32] Peter P Purslow. 2002. The structure and functional significance of variations in the connective tissue within muscle. *Comparative Biochemistry and Physiology Part A: Molecular & Integrative Physiology* 133, 4 (2002), 947–966.
- [33] Rommel G Regis and Christine A Shoemaker. 2005. Constrained global optimization of expensive black box functions using radial basis functions. *Journal of Global optimization* 31, 1 (2005), 153–171.
- [34] Greg Russell, Kevin D Harkins, Timothy W Secomb, Jean-Philippe Galons, and Theodore P Trouard. 2012. A finite difference method with periodic boundary conditions for simulations of diffusion-weighted magnetic resonance experiments in tissue. *Physics in Medicine and Biology* 57, 4 (2012), N35.
- [35] Dinesh Samuel, Philip Rowe, Victoria Hood, and Alexander Nicol. 2011. The relationships between muscle strength, biomechanical functional moments and health-related quality of life in non-elite older adults. *Age and ageing* 41, 2 (2011), 224–230.
- [36] Caroline A Schneider, Wayne S Rasband, and Kevin W Eliceiri. 2012. NIH Image to ImageJ: 25 years of image analysis. *Nature methods* 9, 7 (2012), 671.
- [37] Isaac Selva Raj, Stephen R. Bird, and Anthony J. Shield. 2017. Ultrasound Measurements of Skeletal Muscle Architecture Are Associated with Strength and Functional Capacity in Older Adults. *Ultrasound in Medicine and Biology* 43, 3 (mar 2017), 586–594. <https://doi.org/10.1016/j.ultrasmedbio.2016.11.013>
- [38] Edward O Stejskal and John E Tanner. 1965. Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *The Journal of Chemical Physics* 42, 1 (1965), 288–292.
- [39] J. Towns, T. Cockerill, M. Dahan, I. Foster, K. Gaither, A. Grimshaw, V. Hazlewood, S. Lathrop, D. Lifka, G. D. Peterson, R. Roskies, J. R. Scott, and N. Wilkins-Diehr. 2014. XSEDE: Accelerating Scientific Discovery. *Computing in Science & Engineering* 16, 5 (Sept.-Oct. 2014), 62–74. <https://doi.org/10.1109/MCSE.2014.80>
- [40] Dang Van Nguyen, Jing-Rebecca Li, Denis Grebenkov, and Denis Le Bihan. 2014. A finite elements method to solve the Bloch–Torrey equation applied to diffusion magnetic resonance imaging. *J. Comput. Phys.* 263 (2014), 283–302.
- [41] B. H. Wokke, J. C. van den Bergen, M. J. Versluis, E. H. Niks, J. Milles, A. G. Webb, E. W. van Zwet, A. Aartsma-Rus, J. J. Verschuuren, and H. E. Kan. 2014. Quantitative MRI and strength measurements in the assessment of muscle quality in Duchenne muscular dystrophy. *Neuromuscular Disorders* 24, 5 (may 2014), 409–416. <https://doi.org/10.1016/j.nmd.2014.01.015>
- [42] Dongbin Xiu and George Em Karniadakis. 2002. The Wiener–Askey polynomial chaos for stochastic differential equations. *SIAM journal on scientific computing* 24, 2 (2002), 619–644.
- [43] Junzhong Xu, Mark D Does, and John C Gore. 2007. Numerical study of water diffusion in biological tissues using an improved finite difference method. *Physics in Medicine and Biology* 52, 7 (2007), N111.

A MICROSTRUCTURAL PARAMETERS FOR INVERSION RESULTS

Table A1 provides the parameters for the simulations that were used with the inversion scheme for both the periodic REV and histology simulations.

Table A1: Microstructural Parameters for fitting cases.

REV Geometry	Case 1	Case 2	Case 3	Case 4	Case 5
Diameter (μm)	15	30	50	75	50
Volume Fraction	0.80	0.85	0.72	0.80	0.90
Permeability ($\frac{\mu m}{s}$)	15	50	80	50	25
Intracellular Diff. ($\frac{\mu m^2}{ms}$)	1.0	1.5	2.0	1.2	1.0
Extracellular Diff. ($\frac{\mu m^2}{ms}$)	2.0	2.2	1.5	2.0	1.8
Histology Geometry	Case 1	Case 2	Case 3	Case 4	
Permeability ($\frac{\mu m}{s}$)	10	10	85	50	
Intracellular Diff. ($\frac{\mu m^2}{ms}$)	1.0	1.0	2.0	1.0	
Extracellular Diff. ($\frac{\mu m^2}{ms}$)	2.0	2.0	2.0	1.0	