DOI: 10.14464/gammas.v8i1.972 GAMMAS 2026, Vol. 8, No. 1, p. 83–97

International Association
of Applied Mathematics and Mechanics
– Archive for Students –

Toward Individualized Surgical Parameter Selection via Numerical Simulations and Bayesian Statistics

Felix J. Beckmann[a,†,*], Georg Eichler[b,†], Veit Gemmer[c,†], Anina A. Morgner[d,†], Hannah Willkomm[e,†], Elise Wolf[f,†], Marlon Esmeyer, Jan Grove, Klara M. Gutekunst, Helena Mehler, Philipp M. Siegler, Thilo Bergner, Florentine Domrös, Finn S. Faßbender, Karl Kindermann, Isabel Kurth, Johann Lieberwirth, Leonard Mantel, Hannah Müller, Eva Pöhlmann, Dominik Göddeke[g], Katja Ickstadt[h]

[a] Karlsruhe Institute of Technology, Karlsruhe, Germany
[b] Leipzig University, Leipzig, Germany
[c] Technical University of Munich, Munich, Germany
[d] Technical University of Berlin, Berlin, Germany
[e] Technical University of Darmstadt, Darmstadt, Germany
[f] University of Mannheim, Mannheim, Germany
[g] University of Stuttgart, Stuttgart, Germany
[h] TU Dortmund University, Dortmund, Germany

received 04.06.2026, accepted 28.09.2026, published 26.10.2026

* corresponding author: Felix J. Beckmann, uonhb@student.kit.edu
† equal contribution

Abstract: *Restoring natural movement and sensation after limb loss remains a central challenge in modern prosthetics. Beyond orthopedic changes, amputation disrupts the natural interaction between opposing muscles, which normally supports movement and proprioceptive feedback. The agonist-antagonist myoneural interface (AMI) surgically restores this coupling by connecting opposing muscles. However, the antagonist prestretch, which influences the resulting range of motion (ROM), is currently selected heuristically because clinical data is scarce and high-fidelity simulations are computationally expensive. This study presents an initial interdisciplinary approach toward addressing this gap. We first generate data using multiscale muscle-tendon simulations, then apply Bayesian Optimization (BO) to identify informative sampling points for estimating the optimal prestretch. A near-optimal prestretch is found within a few forward evaluations, revealing an interior maximum beyond which further prestretch reduces ROM. Finally, Bayesian Data Augmentation (DA) infers the underlying data distributions, aiming toward patient-specific, uncertainty-aware predictions of ROM. In essence, our study serves as a computational framework demonstrating a route toward overcoming the data sparsity and hardware limitations that currently obstruct individualized AMI planning.*

Keywords: Muscle prestretch optimization, agonist-antagonist myoneural interface (AMI), Bayesian Optimization (BO), Bayesian Data Augmentation (DA)

This is an open-access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

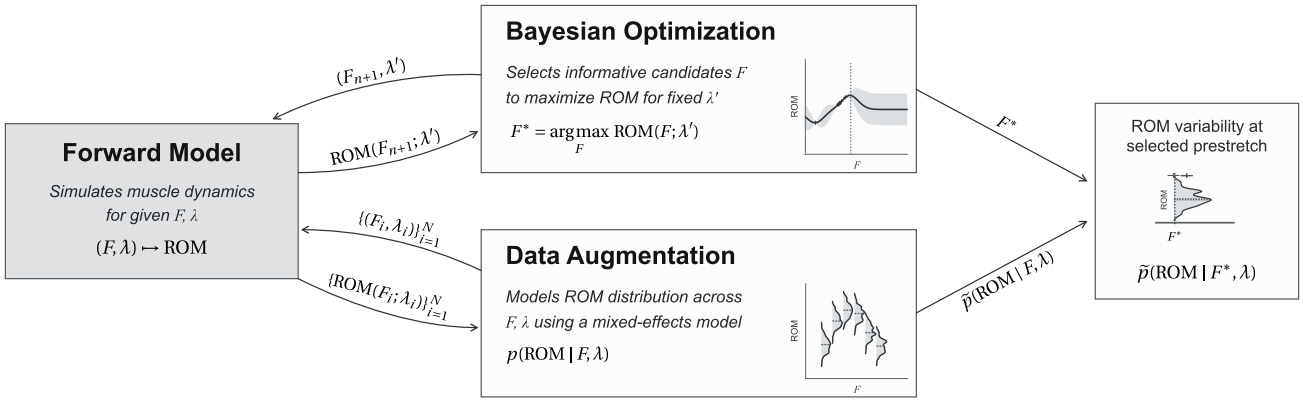


Figure 1 – Methodological workflow. A single forward model maps a prestretch force F and a set of physiological parameters λ to one range of motion (ROM) value. Bayesian Optimization (BO) holds the physiological parameters fixed at one representative set λ', repeatedly proposes a next force F_{n+1}, receives the simulated $\mathrm{ROM}(F_{n+1};\lambda')$, and on termination returns the best evaluated force F^*. Data Augmentation (DA) evaluates the prescribed sweep configurations $\{(F_i,\lambda_i)\}_{i=1}^N$ and fits a mixed-effects model to the returned ROM values, yielding the conditional distribution $p(\mathrm{ROM} \mid F,\lambda)$. Combining both BO and DA gives the estimated ROM variability at the selected prestretch, $\tilde{p}(\mathrm{ROM} \mid F^*,\lambda)$.

1 Introduction

1.1 Agonist-Antagonist Myoneural Interface for limb amputation

Limb amputation is a common procedure, with approximately 60 000 performed annually in Germany [10]. Patients are left with limited control over a prosthesis, one reason being loss of proprioception, the sense of limb position and force, which depends on sensory organs within the muscles and their tendons.

Voluntary movement depends on the coordinated action of muscle pairs. Since muscles can only contract by pulling their attachment points together, these pairs, known as agonist and antagonist muscles, act in opposite directions and are subject to regulation by both mechanical and neural feedback mechanisms [10]. The brain initiates the movement command, which is then transmitted to muscle tissue via central nervous system (CNS). Specialized sensory organs, including muscle spindles and Golgi tendon organs, are responsible for monitoring movement and provide CNS with information about the resulting position and force, allowing it to adjust accordingly [10].

In standard limb amputation, surgery disrupts the natural interactions of agonist-antagonist muscles and leaves transected nerves without end organs. This reduces proprioceptive feedback and thus prosthetic control. In several cases, this can lead to phantom pain [7, 10].

The agonist-antagonist myoneural interface (AMI), developed at MIT, addresses these issues by surgically reconstructing the muscle pair via tendon coupling [5]. The advantage of this is that muscle-tendon information is transmitted to the CNS by neural feedback as well as to the external prosthesis by artificial sensing. Consequently, it also provides a closed-loop feedback from the prosthesis into the CNS. In summary, AMI enhances proprioception and prosthesis control based on neuronal excitation [7, 21]. The first AMI amputation in a human patient was performed in July 2016 [7]. Clinical studies have demonstrated that AMI enhances movement precision, increases the movement range and reduces pain symptoms [22]. Moreover, imaging has confirmed increased fascicle strains [22], supporting improved muscle interaction [7]. However, limitations remain, including atypical sensations resulting from the unnatural force application, which causes the patient to feel the activation and dynamics of both muscles simultaneously [7].

As demonstrated in preclinical experiments using murine and caprine animal models, the AMI agonist muscle produces contractions that are correlated with motor nerve activation. This provides afferent neural feedback that increases with antagonist muscle stretch [7]. The critical surgery parameter of the antagonist muscle stretch describes the magnitude to which the remaining muscles are prestretched when connected. This parameter directly influences force production, range of motion (ROM), and sensory information [10]. No established procedure exists for selecting this patient-dependent variable, which therefore requires optimization. The aim of this paper is to find the optimal prestretch computationally to maximize the resulting ROM.

Neuromuscular models range from reduced single-fiber electrophysiology [4, 9] to coupled multiscale frame-

CC BY

works that integrate mechanics and electrophysiology across spatial scales [10, 15, 18]. A promising approach is provided by three-dimensional continuum-mechanical muscle models, which are able to represent physiological behavior, from mechanics to electrophysiology. In this context, OpenDiHu [15] provides a framework for parameterized multiscale modeling that enables a consistent and flexible representation of coupled neuromuscular processes.

This work builds on a forward model of the OpenDiHu framework [10] to optimize surgical prestretch selection in AMI. Since a single forward evaluation is computationally expensive, exhaustive exploration of the prestretch range is impractical. Bayesian Optimization (BO) is therefore used to identify the prestretch that maximizes ROM from a small number of evaluations. Since clinical data are scarce and physiological properties differ between individuals, Data Augmentation (DA) within a mixed-effects model is used to obtain ROM distributions rather than single-point predictions.

1.2 Motivation of Methods

In order to optimize the ROM in AMI simulations, detailed biophysical models that account for three-dimensional muscle, tendon and neuronal dynamics are required. The existing high-fidelity OpenDiHu model of Homs-Pons et al. [10] is accurate but computationally expensive. Consequently, it is impractical for iterative processes such as naive parameter optimization. Therefore, the process of surgical planning and configuration of AMI remains suboptimal, primarily due to two key factors. First, because of the sparsity of data from simulations and real-world measurements and second, the computational costs resulting from the complexity of modeling.

1.3 Structure

The methodological workflow (Figure 1) consists of three components. The process commences with the *Forward Model*, which simulates the mechanical response of the muscle–tendon system under prestretch. The full agonist-tendon-antagonist configuration is reduced to a single muscle cuboid model - simplifying the geometry and coarsening the spatial discretization, while retaining the essential electromechanical coupling between activation, fiber stretch, and deformation (Figure A1). This ensures that a sufficiently large set of parameter configurations can be simulated and provides the ROM as the parameter of interest as output.

Based on forward-model evaluations, *BO* constructs a probabilistic surrogate of ROM. Concentrating on simulations in regions where improvements are likely, the optimizer identifies prestretch configurations that are likely to yield maximal ROM while maintaining a minimal number of forward simulations.

Finally, the *DA* model fits a mixed-effects model to a separate parameter sweep of the same forward model, accounting for physiological variability between individuals and thereby enabling the generation of parameter distributions rather than single-point estimates. It extends the optimization results from the idealized forward model to the population level and enables clinically meaningful variability analyses.

1.4 Outline of the paper

The structure of the paper is designed to guide the reader from the biophysical foundations of the AMI modeling approach to the interpretation of its computational results. Since the study relies exclusively on simulated data and operates with a reduced model formulation, it is essential to clarify how each methodological step contributes to the overall reasoning and where its interpretational boundaries lie.

Section 2 establishes the modeling assumptions that shape subsequent results. It describes mechanical and electrophysiological simplifications that influence the behavior of the simulated system. Therefore, it defines the scope within which ROM predictions should be understood. Section 3 builds on these modeling foundations. It explains how information from the forward simulation is transformed into an exploratory strategy for identifying optimal prestretch values. This step is subject to its own methodological constraints, as the surrogate model reflects not only the ROM response but also the uncertainties introduced by model reduction and limited sampling. Section 4 extends the analysis by embedding the simulation outputs in a statistical framework that represents physiological variability. Since no clinical data are used, the study is not intended for clinical decision-making. Therefore, an explicit mechanism is required to distinguish between model-based variability and inter-individual differences that would occur in practice. Section 5 integrates all components - forward simulation, BO, and DA - into a joint modeling procedure allowing for a coherent interpretation of the model's behavior. Section 6 provides a thorough examination of the implications and limitations that emerge from this workflow. We identify which conclusions are supported by the modeling chain and which remain unresolved in the absence of experimental validation. Finally, Section 7 summarizes the methodological contributions and outlines potential future research linking computational results to eventual clinical evaluation.

CC BY

2 Forward Model

The simulations in this work are based on the AMI simulation in the OpenDiHu framework [15]. The canonical AMI configuration in OpenDiHu, which we refer to as the AMI simulation, consists of a left muscle, a tendon, and a right muscle, which are coupled bidirectionally in the simulations via *preCICE* [2], a coupling library for multiscale simulations. Homs-Pons et al. [10] formulate the AMI simulation as a multi-X hierarchical muscle model. The left muscle simulates electrophysiology through a 1D fiber simulation driven by a 0D neural simulation, with the underlying mathematical model described in CellML. Here, 0D refers to the absence of a spatial representation, where the neural dynamics are modeled as time-dependent ordinary differential equations at a single spatial location, rather than being resolved along a spatial domain. The resulting neural excitation is then used to drive the 1D propagation of the action potential along the muscle fibers. The left muscle additionally includes 3D mechanics, whereas the tendon and right muscle are modeled purely mechanically.

Although the original configuration provides physiologically accurate electromechanical behavior across the entire agonist-tendon-antagonist chain, the fine spatial resolution required for high-fidelity fiber propagation leads to runtimes that prohibit rapid parameter exploration. Therefore, we reduce the complexity of the model while preserving the electromechanical structure.

The reduction of the AMI configuration proceeds in two steps. First, the geometry is simplified from the full agonist-tendon-antagonist AMI structure to a single muscle with a tendon, referred to as the cuboid model. In the reduced setting considered here, the right muscle is omitted. In the full AMI configuration, its primary role is to provide the opposing mechanical action to the left muscle, while its independent electrophysiological activation is not considered. We therefore restrict the reduced model to the electromechanical behavior of the left muscle. Tendon deformation is assumed to be negligible compared with muscle deformation. This assumption is supported by the full AMI simulations, in which the tendon displacement remained below the reported numerical resolution (up to two decimal places 0.00 cm). Consequently, the effective axial deformation is dominated by the left muscle. All electromechanical couplings, material models, and solver configurations are retained in the simplified model geometry, thereby reducing coupling overhead.

In the second step, both the 3D mechanical discretization and the embedded 1D fiber architecture are coarsened to reduce computational cost. The original configuration utilizes a 3D mesh of $8 \times 8 \times 20$ linear hexahedral elements containing an embedded 6×6 fiber grid, with each fiber discretized by 60 points. In the reduced model, the 3D mesh is down-sampled to $4 \times 4 \times 10$ linear elements, resulting in a correspondingly reduced 3×3 fiber grid and 30 points per fiber. This yields approximately 75% fewer fibers and about 50% in the number of discretization points per fiber (see Table 1).

Table 1 – Mesh analysis: Computational cost for cuboid model at different spatial resolutions.

Discretization	Fibers	Elements	Points per fiber
8×8 (reference)	64	$8 \times 8 \times 20$ (lin.)	60
4×4 (reduced)	16	$4 \times 4 \times 10$ (lin.)	30

Simulations consist of a passive prestretch phase followed by electrically driven contraction. The inferior surface of the cuboid is fixed in the axial direction z, while a traction force of magnitude $F_{\text{prestretch}}$ is applied to the superior surface. Lateral boundaries are constrained to a minimal extent to prevent rigid body motion while permitting transverse deformation. Following the passive loading phase, activation is induced through predefined motor neuron firing sequences, which generate spatially distributed action potentials and subsequent active tension via the Hodgkin-Huxley-Razumova (HHR) electromechanical coupling [4, 9, 17].

The principal output metric is the ROM, defined as the axial excursion of the cuboid muscle during contraction. The muscle length L_t at time step t is computed from the superior surface position $z_{\text{sup}}(t)$ and inferior surface positions $z_{\text{inf}}(t)$ over the time interval $t \in [0, T)$ as follows:

$$L_t = z_{\text{sup}}(t) - z_{\text{inf}}(t),$$

The initial muscle length before prestretch is denoted by L_0, while L_f denotes the muscle length after the passive prestretch phase. The ROM is obtained as

$$\text{ROM} = \max_{t \in [0,T]} L_t - \min_{t \in [0,T]} L_t,$$

providing a scalar measure of the axial excursion induced by contraction for a given prestretch magnitude $F_{\text{prestretch}}$ and subsequent muscle shortening. This real scalar metric serves as the objective function in the optimization studies which will be described in Section 3.

The reduced mesh resolution has been shown to decrease runtime by 91.62%, from 88.06 min to 7.24 min. Therefore, the single-muscle cuboid model with reduced spatial resolution functions as a complexity-reduced model that retains similar accuracy in ROM values and

CC BY

is thereby used for further algorithmic implementations in Section 3 and 4.

The range of admissible prestretch values is guided by in vivo measurements of passive muscle–tendon elongation in humans. Hoang et al. [8] report that the medial gastrocnemius muscle-tendon unit can be elongated by approximately 6.7 cm above its original length of 42.3 cm. This corresponds to a maximal muscle-tendon-unit strain of about 16% under passive stretching. To establish an exploratory upper bound for our parametric studies, we target an elongation of approximately 20% (i.e., stretching a 35 cm muscle by 7 cm), acknowledging that this may exceed typical physiological conditions. However, it provides a conservative margin for investigating the full range of length-dependent contractile behavior.

Given the nonlinear passive mechanics of the Mooney-Rivlin material model, with which the material deformation of the cuboid model in OpenDiHu of Maier et al. [15] is simulated, the relationship between applied traction $F_{\text{prestretch}}$ and resulting elongation in the z-direction $\bar{z}_{\text{sup}}$ is material-dependent and cannot be described a priori. We therefore employ a force-controlled parameterization. The actual elongation, ΔL, achieved for a given force $F_{\text{prestretch}}$ is measured a posteriori from the converged quasi-static equilibrium via

$$\Delta L = \bar{z}_{\text{sup}} - \bar{z}_{\text{inf}} - L_0.$$

3 Bayesian Optimization

We consider the general optimization problem

$$x^* = \arg\max_{x \in \mathcal{X}} f(x), \tag{1}$$

where $x \in \mathcal{X} \subset \mathbb{R}$ is a candidate prestretch force, $\mathcal{X}$ denotes the admissible search space, x^* is the optimal prestretch to be identified and f is a black-box objective function. BO is well-suited to such problems where each evaluation of the objective f is costly, derivatives $\nabla_x f$ are unavailable and f is nonconvex or multimodal. BO can locate the solution x^* with relatively few objective function evaluations by exploiting all previously observed data to select informative new query points.

To this end, BO maintains a *probabilistic surrogate model* of f. Concretely, it assigns a prior over admissible objective functions and updates this prior using Bayes' rule after each observation. At iteration n, a candidate point $x_n \in \mathcal{X}$ is selected by maximizing an *acquisition function*

$$a(\cdot; \mathcal{D}_n) : \mathcal{X} \to \mathbb{R},$$

which scores the utility of evaluating f at x_n given the observations $\mathcal{D}_n$ accumulated so far. The function form itself is fixed; only the underlying Gaussian Process (GP) surrogate changes as new data are incorporated. This choice balances exploration, which is equivalent to sampling where posterior uncertainty is large, and exploitation, equivalent to sampling where the predicted value is high. For a high-level overview of the resulting BO framework, we refer to Algorithm 1 [19].

Algorithm 1 Bayesian Optimization

1: **Input:** initial data $\mathcal{D}_0$, acquisition $a(\cdot; \mathcal{D}_n)$, budget N
2: **for** $n = 1, 2, \ldots, N$ **do**
3: $\quad x_{n+1} \leftarrow \arg\max_{x \in \mathcal{X}} a(x; \mathcal{D}_n)$
4: $\quad y_{n+1} \leftarrow f(x_{n+1})$
5: $\quad \mathcal{D}_{n+1} \leftarrow \mathcal{D}_n \cup \{(x_{n+1}, y_{n+1})\}$
6: $\quad$ update surrogate with $\mathcal{D}_{n+1}$
7: **end for**

In the context of AMI, clinical data are sparse and each forward evaluation is expensive, so a finely resolved exhaustive sweep of the prestretch axis is costly. Therefore, we optimize the forward model from Section 2 using BO to identify the prestretch force $F_{\text{prestretch}}$ that maximizes the ROM:

$$\arg\max_{F_{\text{prestretch}} \in \mathcal{X}} \text{ROM}(F_{\text{prestretch}}). \tag{2}$$

In the following, we describe the surrogate model and acquisition function used in our implementation.

Surrogate Model. We place a GP [25] on f,

$$f \sim \mathcal{GP}\big(m(x),\, k(x, \tilde{x})\big), \tag{3}$$

with mean m and covariance kernel $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$, where $\tilde{x} \in \mathcal{X}$ denotes a second arbitrary input point. Given data $\mathcal{D}_n = \{(x_i, y_i)\}_{i=1}^n$ with noisy observations $y_i = f(x_i) + \varepsilon_i$, $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$, the GP posterior at x is Gaussian with mean $\mu_n(x)$ and standard deviation $\sigma_n(x)$:

$$\boldsymbol{K}_n = \big[k(x_i, x_j)\big]_{i,j=1}^n, \quad \mathbf{k}_n(x) = \big(k(x_1, x), \ldots, k(x_n, x)\big)^\top,$$
$$\mu_n(x) = m(x) + \mathbf{k}_n(x)^\top \big(\boldsymbol{K}_n + \sigma^2 \mathbf{I}\big)^{-1} (\mathbf{y} - \mathbf{m}),$$
$$\sigma_n^2(x) = k(x, x) - \mathbf{k}_n(x)^\top \big(\boldsymbol{K}_n + \sigma^2 \mathbf{I}\big)^{-1} \mathbf{k}_n(x),$$

where $\mathbf{y} = (y_1, \ldots, y_n)^\top$ and $\mathbf{m} = (m(x_1), \ldots, m(x_n))^\top$.

Acquisition Function. We use the Upper Confidence Bound (UCB) acquisition [20]:

$$a^{\text{UCB}}(x; \mathcal{D}_n) = \mu_n(x) + \sqrt{\beta}\, \sigma_n(x),$$

CC BY

where μ_n and σ_n come from the GP posterior. The parameter $\beta > 0$ controls the trade-off between exploration and exploitation: a larger β encourages sampling where uncertainty is high, whereas smaller β favors high predicted values.

We implemented BO in BoTorch [1] using a GP surrogate model and UCB acquisition function with $\beta = 0.1$. The initial dataset $\mathscr{D}_0$ contains one evaluation at a uniformly sampled prestretch force. Since each BO iteration adds one evaluation, n iterations correspond to $n+1$ forward-model evaluations in total. We terminated optimization using an exponential moving average stopping criterion with window size 2.

A limitation of BO in this setting is that its uncertainty quantification is tied to the surrogate model and sparse forward evaluations, rather than to physiological or patient-level variability. For sparse data regimes, which are precisely the setting in AMI simulations, this becomes an important methodological gap. This motivates the subsequent DA step, which provides a distributional model of ROM variability across prestretch values.

4 Data Augmentation

The limited availability of medical datasets for AMI processes [21], combined with high inter-subject variability due to anatomical and physiological differences, requires generalizable patient-specific models to obtain patient-individual optimal ROM. We address this issue by using Bayesian DA. The concept of augmented data is closely connected to that of latent variables, as it introduces variables Y_{aug} into the posterior distribution $p(\theta|Y)$ [23]. As used here, it conceptually differs from common machine-learning approaches, which typically extend observed data through transformations. Here it introduces statistically principled latent structures that directly reflect the underlying problem [23].

Given the observed data $Y \in \mathbb{R}^d$, the method facilitates efficient sampling from the posterior

$$p(\theta \mid Y) \propto p(Y \mid \theta)\, p(\theta), \tag{4}$$

where $p(Y \mid \theta)$ denotes the likelihood with respect to a reference measure μ on $(\mathbb{R}^d, \mathscr{B}(\mathbb{R}^d))$ and $p(\theta)$ is the prior density on $\Theta \subset \mathbb{R}^k$ [23].

Augmented data Y_{aug} is connected to the observed data Y via a many-to-one mapping $\mathscr{M} : Y_{\text{aug}} \to Y$ and must satisfy

$$\int_{\mathscr{M}(Y_{\text{aug}})=Y} p(Y_{\text{aug}}|\theta)\mu(dY_{\text{aug}}) = p(Y|\theta),$$

to ensure the augmented data does not distort the original model $p(Y|\theta)$ when combining the augmented data points mapped to a given observed data point by $\mathscr{M}$ [23].

In order to apply data augmentation to a given problem two design choices have to be made. First, one needs to define a parametrized model which maps covariates to observations. Second, one must choose an optimization strategy for the model parameters based on sampling augmented data Y_{aug}. We decided to use a mixed-effects model, which introduces a distributional perspective on the response. It quantifies uncertainty in both fixed effects (population-level trends) and random effects (subject-level variability). This model represents an appropriate initial approach, given the sparsity of the data and the clear distinction between controlled experimental parameters, such as muscle length, and physiologically varying quantities, such as muscle density, which can fluctuate over time within a given patient. The mixed-effects model is a combination of terms linear in the fixed and random effects together with a statistical error term. It takes the form

$$y_i = X_i\beta + Z_i\xi + Z_i b_i + e_i, \tag{5}$$

where $y_i \in \mathbb{R}^{n_i \times 1}$ is the observed response for all observations $i = 1,\dots,m$, which is composed of the covariates $X_i \in \mathbb{R}^{n_i \times p}$ and $Z_i \in \mathbb{R}^{n_i \times q}$ with $\beta \in \mathbb{R}^{p\times 1}$ and $\xi \in \mathbb{R}^{q\times 1}$ as the fixed effects and $b_i \in \mathbb{R}^{q\times 1}$ as the random effect. The random effect b_i and the error term e_i are drawn from the normal distributions $b_i \sim \mathscr{N}_q(0, T)$ with T being the covariance and $e_i \sim \mathscr{N}(0, \tau^2 I_{n_i})$ with $\tau \in \mathbb{R}$, where $b_i \perp e_i$ holds. For simplicity, $n_i \equiv n$ for all i will be assumed. The parameters of interest are $\theta = (\beta, \xi, \tau^2, T)$ which will be optimized.

Data augmentation treats the random effect as missing data introducing the working parameter $\omega := (\gamma, \Upsilon)$ where $\gamma \in \mathbb{R}^{q\times 1}$ and $\Upsilon \in \mathbb{R}^{q\times q}$. The augmented data is collectively denoted as

$$Y_{\text{aug}} = \{(y_i, X_i, Z_i, d_i := \Upsilon^{-1} b_i + \gamma), \quad i = 1,\dots,m\}.$$

This form of Y_{aug} is meaningful as it will later allow for an efficient and suitable optimization of the model's parameters: In our framework, the data augmentation of the mixed-effect model is conducted as proposed in the two-step approach from Van Dyk and Meng [23], which combines the Expectation Maximization (EM) algorithm and the linear regression model in combination with the normal-inverse Wishart distribution [6].

In the first step of the two-step sampling scheme, the random-effect coefficients b_i are drawn from a normal distribution. The second step combines the observed data, the augmented data from the previous iteration, and the sampled b_i: the variance of the error term, τ^2, is

CC BY

sampled first on this basis, after which the fixed-effects coefficient β, an auxiliary parameter and the working parameter Υ, are sampled jointly. A second auxiliary parameter is subsequently sampled, followed by the second working parameter γ, yielding the updated augmented data. The remaining mixed-effects parameters, ξ and T, are then updated on this basis. A detailed description of this procedure is provided in Appendix A.1.

This framework efficiently samples complex posteriors through improved Markov chain Monte Carlo (MCMC) mixing, addressing data sparsity while providing detailed distributions of patient-specific attributes related to ROM. Based on this framework, we define both fixed and random effects, enabling the estimation of physiological variability and uncertainty. The fixed effects comprise three parameters. First, the prestretch force describes the externally applied prestretch ($F_{\text{prestretch}}$) and serves as the primary parameter of interest for optimization, being systematically varied throughout the simulation protocol. Second, the geometric extent of the muscle (L_0) is held constant within each simulation batch but may vary across individuals. Third, a constant term equal to 1 is incorporated to introduce an intercept into the mixed-effects model. Effectively this allows the model to represent an affine linear mapping rather than a purely linear and thus making it more general. The random effects are chosen to represent physiological heterogeneity at both the muscle-unit and cellular levels. The surface area-to-volume ratio (A_m) captures variability in muscle fiber geometry and membrane surface density, both of which influence electrophysiological properties. Similarly, muscle density (ρ) reflects variability in tissue composition and depends on the hydration state of the muscle.

As this study serves as a proof of concept, random effects are assumed to follow a normal distribution within the specified ranges. For real-world applications, clinically informed prior distributions should be incorporated instead, alongside a problem-specific choice of interaction and higher-order terms.

5 Results

5.1 Data Creation

The data used throughout Sections 2–4 have been generated from the reduced cuboid muscle formulation introduced in Section 2. All simulations employ the single-muscle cuboid model with reduced spatial fiber resolution, which enables a substantial reduction in runtime while preserving the electromechanical structure of the original AMI simulation. A single forward run requires approximately seven to eight minutes, allowing the evaluation of sufficiently many prestretch configurations for Bayesian Optimization in Section 3 as well as a parameter sweep used for the mixed-effects model. Corrupted forward runs exhibiting nonphysical ROM values were excluded.

The simulation for BO assumes an initial muscle length of $L_0 = 34.97\,\text{cm}$ and an aspect ratio of the muscle of 9.04 as fixed effects, chosen as representative values from within the parameter ranges of the complete parameter sweep (Table A1).

The complete synthetic dataset incorporates all uncorrupted combinations of the parameter sweep summarized in Table A1, resulting in 3764 forward simulations. These serve as the basis for the posterior sampling in the Data Augmentation framework using a mixed-effects model.

In the mixed-effects model, each physiological parameter is defined together with its assumed distributional properties (see Section 4 for parameter explanation). The random effects describe physiologically meaningful heterogeneity at the muscle-fiber and cellular level, while fixed effects represent population-level mechanical response. All parameter ranges were revised for internal consistency between geometry, material properties, and electrophysiological characteristics. At least one effect included in the linear model must enter quadratically in order to represent an extremum, thereby making the optimization problem well-posed. A given effect can be incorporated quadratically by including it twice: once in its linear form and once in its squared form. Based on simulation data and expert consultation, the prestretch force was modeled quadratically, together with the cross terms of aspect ratio with prestretch force and with prestretch force squared, respectively.

5.2 Results of Each Methodological Component

Overall interpretation. Across all components of our pipeline, we observe that reduced model complexity, in combination with distributional reasoning, allows us to efficiently explore the AMI prestretch setting. The reduction of the forward model maintains physical meaningfulness while significantly lowering computational cost, enabling effective parameter studies even under sparse data availability. The Bayesian Optimization run concentrates its evaluations near an interior local maximum of the surrogate, while the mixed-effects model provides interpretable uncertainty quantification and allows predictions that account for patient-specific variability. The integration of these components results in a cohesive chain that encompasses electromechanical modeling, surrogate-based optimization, and population-level gen-

CC BY

eralization.

Forward Model. The reduced 4×4 cuboid discretization preserves the prestretch-induced deformation observed with the 8×8 reference discretization, while substantially reducing computational cost (cf. Table 2). For both discretizations, the muscle length increases from $L_0 = 34.97$ cm to $L_f = 41.96$ cm under a prestretch force of 31 N, corresponding to an elongation of $\Delta L = 6.99$ cm. Without prestretch, both discretizations remain at their initial length, yielding $\Delta L = 0$ cm. Thus, for the tested prestretch conditions, the coarser spatial discretization reproduces the same axial prestretch deformation as the reference discretization and the model complexity reduction preserves the physical interpretability of AMI simulations while enabling extensive parameter exploration.

discretization	$F_{\text{prestretch}}$ (N)	L_0 (cm)	L_f (cm)	ΔL (cm)	runtime (s)
4×4 (reduced)	0	34.97	34.97	0.00	434.61
4×4 (reduced)	31	34.97	41.96	6.99	454.17
8×8 (reference)	0	34.97	34.97	0.00	5283.52
8×8 (reference)	31	34.97	41.96	6.99	5319.26

Table 2 – Comparison for reduced and full reference cuboid discretizations of the passive prestretch response and the cost of one forward run. L_0 is the axial muscle length before any prestretch is applied, L_f the axial length reached at the end of the passive prestretch phase, and $\Delta L = L_f - L_0$ the resulting elongation, as defined in Section 2. Runtime is the wall-clock time of one complete forward run.

Bayesian Optimization. The Bayesian Optimization step is driven directly by forward-model evaluations of ROM. By using the cuboid model within the OpenDiHu framework, the same optimization workflow can be integrated with other neuromuscular simulation setups, requiring minimal modification.

The reported Bayesian Optimization run comprised one initial evaluation followed by 13 BO iterations, resulting in 14 forward-model evaluations. The largest observed ROM was obtained at a prestretch force of $F_{\text{prestretch}} = 45.67$ N, first sampled at the 12th forward-model evaluation.

Unlike a naive uniform grid search on the search space $\mathscr{X}$, BO allocates evaluations non-uniformly by selecting each new prestretch force based on the current GP surrogate. For reference, a uniform endpoint-inclusive grid with 14 evaluations over $\mathscr{X} = [0, 70]$ N would have a fixed spacing of

$$\frac{70\,\text{N}}{14-1} = 5.38\,\text{N}.$$

In contrast, the final three BO evaluations were performed at 45.67, 46.40, and 46.39 N. All three evaluation points lie within $[45.67, 46.40]$ N, an interval of width 0.73 N. Since a uniform grid applies the same spacing throughout $\mathscr{X}$, achieving a comparable resolution would require a substantially larger set of uniformly distributed prestretch values and, consequently more forward-model evaluations. Specifically, taking $\Delta F = 0.73$ N as the maximum grid spacing, an endpoint-inclusive uniform grid over $\mathscr{X}$ would contain 97 prestretch values and hence require 97 forward-model evaluations.

At each iteration, the next prestretch force is chosen by maximizing the UCB acquisition function over $\mathscr{X}$. With $\beta = 0.1$, the uncertainty term is weighted less strongly than the predicted ROM, so the acquisition function primarily favors exploitation while retaining some exploration. The stages shown in Figure 2 illustrate how the GP surrogate and the selected evaluation points evolve as forward-model results are incorporated. Following each update, the GP surrogate provides a posterior mean prediction of the ROM and an associated posterior standard deviation throughout $\mathscr{X}$, including at prestretch forces not yet evaluated by the forward model. Figure 2A shows the broad posterior uncertainty after the initial evaluation, while Figure 2B shows the increasing concentration of evaluations in regions favored by the acquisition function after seven completed evaluations of the forward model.

The final surrogate (Figure 2C) indicates that ROM increases with prestretch up to approximately 45 N, representing a local maximum found by BO. The observed monotonic trend up to this value is broadly consistent with the clinical finding that antagonist stretch enhances proprioceptive feedback and movement range [7, 22].

Beyond this value, the predicted ROM decreases despite further increases in prestretch, indicating that larger prestretch does not necessarily lead to larger ROM.

Data Augmentation. Forward-simulation outputs act as model inputs, linking physical simulation and statistical inference.

The model then predicts ROM distributions for three different sample sizes which are shown in Figure 3. The resulting distributions reflect the quadratic dependence of ROM on prestretch force, which has been explicitly incorporated in the mixed-effects model. It exhibits a maximum around 45 N across all sample sizes. For a given prestretch, the ROM distribution originates from

CC BY

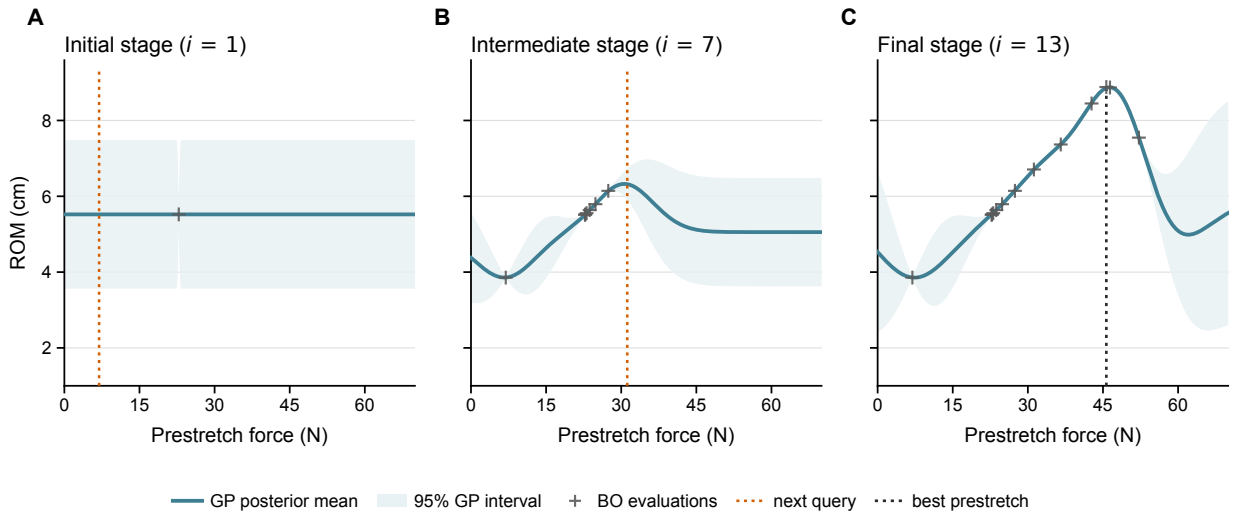


Figure 2 – Bayesian Optimization (BO) surrogate predictions of range of motion (ROM) at the initial **(A)**, intermediate **(B)** and final **(C)** surrogate stages of the reported run. The index i printed in each panel is the number of completed forward-model evaluations the surrogate was fitted to; it counts the initialization, so $i = 1$ in **(A)**, and it is not the BO iteration index. Panel **(C)** shows the last surrogate that was fitted, at $i = 13$; the 14th and final forward-model evaluation was selected from it and is therefore not among the crosses shown. The solid line and shaded band denote the posterior mean and approximate 95% posterior interval, respectively. Crosses mark completed forward-model evaluations. In **(A)** and **(B)** the orange dashed line marks the next prestretch selected by the acquisition function; in **(C)** the black dashed line marks the best evaluated prestretch.

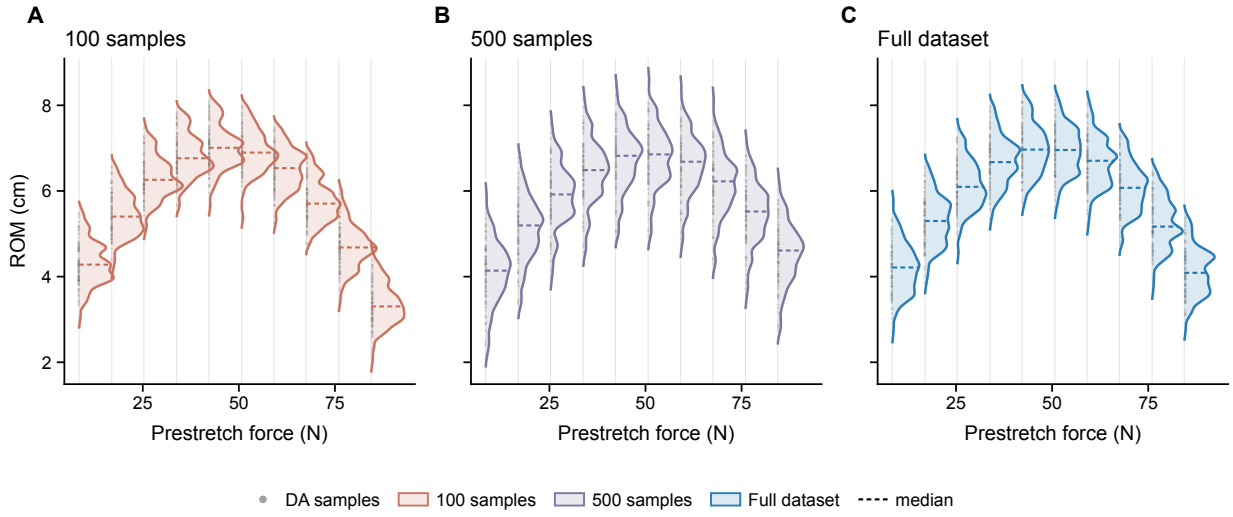


Figure 3 – All panels show the range of motion (ROM) distributions produced by Data Augmentation (DA) integrated over random variables. Subfigures **(A)**, **(B)** and **(C)** show model predictions for training with 100, 500 and the full dataset, i.e., 3764 samples, respectively.

CC BY

the integration over the random variables. The spread of the ROM distribution overlaps across all prestretch values and sample sizes. While the distribution characteristics are similar across all prestretch values for 100 samples (Figure 3A), the ROM distributions vary across prestretch values for the full dataset (Figure 3C). Comparing the distributions for 100 and 500 samples, the 500 sample distribution shows lower peaks and longer tails, reflecting more outliers but a lack of sufficient data to mitigate their negative impact (Figure 3A-B). For the full dataset, the effect of outliers diminishes: the distributions show a heavy centered bulk with thin tails.

5.3 Optimized Prestretch under ROM Variability

BO identifies the best evaluated prestretch force from a limited number of forward-model evaluations. Data augmentation can extend this result by providing the corresponding ROM distributions for the evaluated prestretch forces across multiple physiological settings.

Figure 4 combines the results of both approaches. The BO evaluations broadly follow the trend of the DA predictions, and both approaches indicate a maximum at a similar prestretch force. The DA results therefore add distributional information to the BO result, allowing the selected prestretch force to be interpreted in light of the predicted ROM variability.

Nevertheless, it is evident that the DA model includes only terms of the prestretch force up to quadratic order effectively restricting the possible form of predictions. The reasoning for this is provided in Section 5.1. This model choice reduces the complexity at the cost of the model's versatility. This hard restriction of the slope of ROM medians worsens the capability to model all BO predictions accurately. Nevertheless, the agreement in the location of the maximum supports the consistency of the two results.

6 Discussion

6.1 Limitations

The results of this study should be interpreted in light of several methodological constraints. First, all analyses rely exclusively on simulated data from a reduced single-muscle cuboid model rather than the full agonist-antagonist AMI configuration. While the reduction preserves the essential electromechanical coupling, it does not capture the full bidirectional interaction between agonist and antagonist in the full AMI configuration. The forward model also does not represent sensory pathways or proprioceptive feedback. Second, the absence of in vivo calibration and clinical measurements means that the quantitative outputs cannot be directly transferred to patient care. Third, the mixed-effects model relies on a quadratic approximation of the prestretch-ROM relationship, which limits its flexibility. Finally, the reported BO result is based on one run for a fixed physiological configuration λ' and a fixed UCB exploration weight $\beta = 0.1$. With 14 forward-model evaluations, the

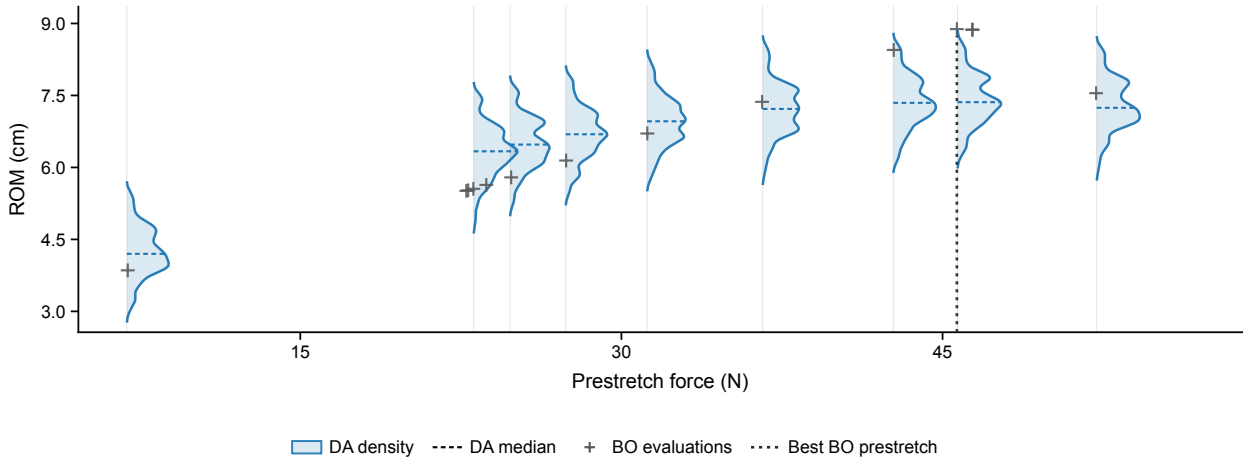


Figure 4 – The figure combines the point results from Bayesian Optimization (BO) and distributional results from Data Augmentation (DA). The grey dashed line marks the best evaluated prestretch found by BO. The blue distributions show the corresponding range of motion (ROM) distributions produced by DA integrated over the random variables.

CC BY

run identifies the best evaluated prestretch within the investigated range but cannot guarantee that this value is the global optimum of the underlying objective.

6.2 Outlook

Future research should include an empirical evaluation of the full AMI configuration, which is expected to be more computationally demanding. We did not measure the runtime of the full AMI configuration. However, for context, Maier and Schulte [14] report that simulating 1 s with a high-resolution OpenDiHu muscle–tendon simulation required a runtime of approximately 7 h 20 min using 172 processes on four compute nodes of the Hawk supercomputer at HLRS Stuttgart. Although these results are not directly comparable to the present model, they illustrate the computational demands of high-resolution OpenDiHu simulations. The computational demands also motivate a systematic investigation of further electrophysiological model reduction. Potential approaches include removing selected ordinary differential equation components or simplifying action-potential propagation, although their computational benefits and ability to preserve the relevant neuromuscular dynamics would need to be evaluated.

Furthermore, future work could examine adaptive control of the exploration–exploitation balance within each BO run. Our reported BO run used a fixed UCB parameter, $\beta = 0.1$, to control the balance between exploration and exploitation. A fixed coefficient may cause the search to concentrate prematurely around the current estimated best prestretch or, if chosen too large, to allocate too many evaluations to uncertain but less promising regions. Wang et al. [24] propose an alternative estimation-based acquisition strategy and show that, at each iteration, it selects the same point as GP–UCB with a posterior-dependent exploration weight, thereby automatically and adaptively balancing exploration and exploitation. For prestretch selection, such an approach could reduce sensitivity to a manually chosen value of β and allow the limited evaluation budget to shift between examining uncertain force ranges and refining promising regions as information accumulates.

Beyond an individual BO run, information could also be reused across different physiological configurations in patient-specific settings. In the present implementation, the GP models ROM as a function of prestretch force for one fixed physiological configuration $\boldsymbol{\lambda}'$. Optimizing the prestretch for another configuration would therefore require a separate BO run with new forward-model evaluations. Contextual BO could instead model $f(F, \boldsymbol{\lambda})$ jointly over prestretch force F and the physiological parameters $\boldsymbol{\lambda}$, allowing evaluations from related configurations to inform the search for a new configuration [11]. This could reduce the number of additional forward-model evaluations required for a new physiological configuration.

Finally, assessing the transferability of the results to patient-specific surgical planning requires in vivo calibration and clinical validation using measurements of the applied prestretch and resulting ROM. The availability of such data remains limited because AMI is still a relatively uncommon surgical procedure. Carty et al. [3] reported recently that about 50 lower extremity AMI procedures have been performed in 41 patients. Therefore, initial patient-specific applications would likely need to combine sparse clinical measurements with simulation data and physiologically informed prior distributions.

7 Conclusion

This work presents a three-stage computational framework for investigating the AMI under varying prestretch configurations. First, a reduced forward model enables efficient simulation of the muscle-tendon system across a physiologically plausible prestretch range. Second, BO identifies a local optimal prestretch using only few forward evaluations, showing that optimal ROM is not achieved by the strongest possible prestretch and therefore reveals a non-linear relationship between prestretch and ROM. Third, a mixed-effects data augmentation framework provides patient-conditioned ROM distributions that account for inter-individual physiological variability. Collectively, these components address the key challenges of data sparsity and computational cost in AMI simulation, serving as a proof of concept for future work toward clinically informed and individualized prestretch selection.

Code Availability: The source code and simulation scripts used to produce the results presented in this paper are openly available at `https://github.com/AG1-NUIW-Kolleg/Toward-Individualized-Surgical-Parameter-Selection.git`.

CRediT authorship statement: F. J. Beckmann, G. Eichler, V. Gemmer, A. A. Morgner, H. Willkomm, E. Wolf: Conceptualization, Methodology, Software, Validation, Investigation, Writing – original draft, Writing – review & editing; M. Esmeyer, J. Grove, K. M. Gutekunst, H. Mehler, P. M. Siegler, T. Bergner, F. Domrös, F. S. Faßbender, K. Kindermann, I. Kurth, J. Lieberwirth, L. Mantel, H. Müller, E. Pöhlmann: Methodology, Software; D. Göddeke, K. Ickstadt: Conceptualization, Supervision;

For the methodological development, the team was divided into subgroups: Forward Model (G. Eichler, E.

CC BY

Wolf, F. Domrös, I. Kurth, H. Müller, E. Pöhlmann, F. S. Faßbender); Bayesian Optimization (H. Willkomm, V. Gemmer, J. Grove, L. Mantel, T. Bergner, K. Kindermann, J. Lieberwirth); Data Augmentation (F. J. Beckmann, A. A. Morgner, M. Esmeyer, K. M. Gutekunst, H. Mehler, P. M. Siegler). However, the first authors were responsible for writing and reviewing the manuscript across all topic areas.

Acknowledgments: We would like to thank our supervisors, Prof. Dr. Katja Ickstadt and Prof. Dr. Dominik Göddeke, for their continuous support and patient guidance throughout this project. We are also very grateful to Carme Homs-Pons, who shared her expertise on the AMI and on the forward model with us, and to Robin Lautenschlager for many helpful discussions. Finally, we thank the Studienstiftung des deutschen Volkes, and in particular the Naturwissenschaftliches Kolleg, for making this format possible and for the support we received throughout.

References

[1] Maximilian Balandat, Brian Karrer, Daniel R. Jiang, Samuel Daulton, Benjamin Letham, Andrew Gordon Wilson, and Eytan Bakshy. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In *Advances in Neural Information Processing Systems 33*, 2020.

[2] Hans-Joachim Bungartz, Florian Lindner, Bernhard Gatzhammer, Miriam Mehl, Klaudius Scheufele, Alexander Shukaev, and Benjamin Uekermann. precice–a fully parallel library for multi-physics surface coupling. *Computers & Fluids*, 141:250–258, 2016.

[3] Matthew J. Carty, Michael Fernandez, Corey L. Sullivan, Rachael Chiao, Lori Berger, Tawnee L. Sparling, Jason M. Souza, Benjamin K. Potter, Luis Beltran, and Hugh M. Herr. Human implementation of upper extremity amputation incorporating agonist–antagonist myoneural interface construction. *Plastic and Reconstructive Surgery*, 157(6):874e–886e, 2026.

[4] CellML. Hodgkin-huxley model (cellml repository). `https://models.cellml.org/exposure/c94a51d2f7bc05f4221e1413a26b63ed`, 2025. Accessed: 15 May 2025.

[5] Tyler R Clites, Matthew J Carty, Shriya Srinivasan, Anthony N Zorzos, and Hugh M Herr. A murine model of a novel surgical architecture for proprioceptive muscle feedback and its potential application to control of advanced limb prostheses. *Journal of Neural Engineering*, 14(3):036002, 2017.

[6] Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. *Bayesian Data Analysis*. CRC Press, Boca Raton, FL, 3rd edition, 2021.

[7] Hugh Herr and Matthew J. Carty. The agonist-antagonist myoneural interface. *Techniques in Orthopaedics*, 36(4):337–344, 2021.

[8] PD Hoang, RD Herbert, G Todd, RB Gorman, and SC Gandevia. Passive mechanical properties of human gastrocnemius muscle–tendon units, muscle fascicles and tendons in vivo. *Journal of Experimental Biology*, 210(23):4159–4168, 2007.

[9] A. L. Hodgkin and A. F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4):500–544, 1952.

[10] Carme Homs-Pons, Robin Lautenschlager, Laura Schmid, Jennifer Ernst, Dominik Göddeke, Oliver Röhrle, and Miriam Schulte. Coupled simulations and parameter inversion for neural system and electrophysiological muscle models. *GAMM-Mitteilungen*, 47(3):e202370009, 2024.

[11] Andreas Krause and Cheng Soon Ong. Contextual Gaussian process bandit optimization. In *Advances in Neural Information Processing Systems*, volume 24, pages 2447–2455. Curran Associates, Inc., 2011.

[12] Kaitlyn C Leonard, Nikole Worden, Marissa L Boettcher, Edwin Dickinson, Kailey M Omstead, Anne M Burrows, and Adam Hartstone-Rose. Anatomical and ontogenetic influences on muscle density. *Scientific Reports*, 11(1):2114, 2021.

[13] Jan Lexell, David Downham, and Michael Sjöström. Distribution of different fibre types in human skeletal muscles: fibre type arrangement in m. vastus lateralis from three groups of healthy men between 15 and 83 years. *Journal of the neurological sciences*, 72(2-3):211–222, 1986.

[14] Benjamin Maier and Miriam Schulte. Mesh generation and multi-scale simulation of a contracting muscle–tendon complex. *Journal of Computational Science*, 59:101559, 2022.

[15] Benjamin Maier, Dominik Göddeke, Felix Huber, Thomas Klotz, Oliver Röhrle, and Miriam Schulte. Opendihu: an efficient and scalable framework for biophysical simulations of the neuromuscular system. *Journal of Computational Science*, 79:102291, 2024.

[16] TD O'Brien, ND Reeves, V Baltzopoulos, DA Jones, and CN Maganaris. Muscle-tendon structure and dimensions in adults and children. *Journal of Anatomy*, 216(5):631–642, 2010.

[17] Maria V. Razumova, Anna E. Bukatina, and Kenneth B. Campbell. Stiffness-distortion sarcomere model for muscle simulation. *Journal of Applied Physiology*, 87(5):1861–1876, 1999.

[18] Oliver Röhrle, John B Davidson, and Andrew J Pullan. A physiologically based, multi-scale model of skeletal muscle structure and function. *Frontiers in physiology*, 3:358, 2012.

[19] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104 (1):148–175, 2015.

[20] Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: no regret and experimental design. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ICML'10, page 1015–1022, Madison, WI, USA, 2010. Omnipress.

[21] S. S. Srinivasan, G. Tuckute, J. Zou, S. Gutierrez-Arango, H. Song, R. L. Barry, and H. M. Herr. Agonist-antagonist myoneural interface amputation preserves proprioceptive sensorimotor neurophysiology in lower limbs. *Science Translational Medicine*, 12(573):eabc5926, 2020.

[22] Shriya S. Srinivasan, Samantha Gutierrez-Arango, Ashley Chia-En Teng, Erica Israel, Hyungeun Song, Zachary Keith Bailey, Matthew J. Carty, Lisa E. Freed, and Hugh M. Herr. Neural interfacing architecture enables enhanced motor control and residual limb functionality postamputation. *Proceedings of the National Academy of Sciences*, 118(9):e2019555118, 2021.

[23] David A Van Dyk and Xiao-Li Meng. The art of data augmentation. *Journal of Computational and Graphical Statistics*, 10(1): 1–50, 2001.

[24] Zi Wang, Bolei Zhou, and Stefanie Jegelka. Optimization as estimation with Gaussian processes in bandit settings. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine*

CC BY

Learning Research, pages 1022–1031. PMLR, 2016.

[25] Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006.

CC BY

A Appendix

A.1 Two step optimization of the mixed-effects model in Bayesian Data Augmentation

This appendix details the two steps applied in the Data Augmentation section in order to efficiently sample suitable parameters for the mixed effect model. It follows the method introduced by Van Dyk and Meng [23].

Step 1: The random effects coefficients b_i are drawn independently for $i = 1, \ldots, m$ from a normal distribution

$$b_i \sim \mathcal{N}(\hat{b}_i(\theta), \mathrm{Var}_{b_i})$$

where the mean $\hat{b}_i(\theta)$ describes the random effect in the observed data, the variance Var_{b_i} includes contributions from random noise and the covariance of random effects. Mean and variance are scaled by observed and augmented data, respectively.

Step 2: Given the output from Step 1, the parameters $\tilde{\theta} = (\beta, \tilde{\xi}, \tau^2, \tilde{T})$ and the working parameter ω for the augmented data are sampled [23]:

$$[\tau^2]^{(t+1)} \sim \frac{Y^\top \left(I_n - \tilde{X}(\tilde{X}^\top \tilde{X})^{-1} \tilde{X}^\top\right) Y}{\chi^2_{n-p-q-\nu(\Upsilon)}},$$

where $\tilde{X}$ combines the augmented and observed data (see (7)) and I_n is the unit matrix. Intuitively, τ^2 is the updated random error-term, which is the part of the observed labels that is orthogonal to the hyperplane spanned by observed and augmented data. The χ^2 acts as a normalization factor with respect to the degrees of freedom.

Both the fixed effects coefficients β and $\tilde{\xi}$ as well as the augmented data parameter Υ are drawn from a normal distribution via

$$\begin{pmatrix} \beta^{(t+1)} \\ \tilde{\xi} \\ \mathrm{vec}(\Upsilon) \end{pmatrix} \sim \mathcal{N}\left((\tilde{X}^\top \tilde{X})^{-1} \tilde{X}^\top Y, [\tau^2]^{(t+1)} (\tilde{X}^\top \tilde{X})^{-1}\right). \tag{6}$$

This normal distribution builds upon the linear regression model of augmented and observed data $\tilde{X}$ and Y. Furthermore, the covariance of the distribution is additionally scaled by the random error term sampled before. The second augmented data parameter γ is generated from a normal distribution centered around $\mu_\gamma = \sum_{i=1}^m \frac{b_i}{m}$

$$\gamma \sim \mathcal{N}\left(\mu_\gamma, \frac{1}{m}\tilde{T}\right),$$

where the covariance is the multidimensional generalization of the χ^2 method, thus introducing regularization for γ

$$\tilde{T}^{-1} \sim \mathrm{Wishart}_{m-q-1}[(\sum_{i=1}^m (b_i - \mu_\gamma)(b_i - \mu_\gamma)^\top)^{-1}].$$

Finally, we update the augmented data and the quantities $\tilde{X}$, $\xi^{(t+1)}$, $T^{(t+1)}$

$$\tilde{X} = \begin{pmatrix} X_1 & Z_1 & \tilde{Z}_1\big(d_1(b_1)\big) \\ \vdots & \vdots & \vdots \\ X_m & Z_m & \tilde{Z}_m\big(d_m(b_m)\big) \end{pmatrix}, \tag{7}$$

$$\xi^{(t+1)} = \tilde{\xi} + \Upsilon\gamma, \tag{8}$$

$$T^{(t+1)} = \Upsilon \tilde{T} \Upsilon^\top, \tag{9}$$

where $\tilde{Z}$ is the augmented version of Z. It depends on the combined augmentation parameter $d_i := \Upsilon^{-1} b_i + \gamma$ and is defined as

$$\tilde{Z}(d_i) = (Z_{i1} d_i^\top, \ldots, Z_{iq} d_i^\top)_{n_i \times q^2}.$$

CC BY

Altogether, this gives an updated version of the augmented data $Y_{\text{aug}} = \{(y_i, X_i, Z_i, d_i := \Upsilon^{-1} b_i + \gamma), \quad i = 1, \ldots, m\}$ as well as auxiliary variables which will then be used in the next iteration. The algorithm iterates until convergence or the occurrence of a stopping criterion.

A.2 Forward-model discretization and parameter-sweep configuration

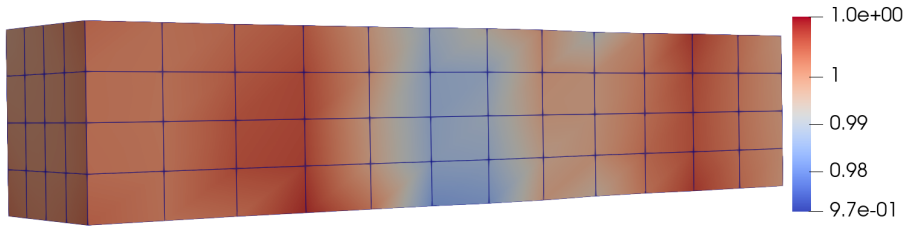


Figure A1 – Discretization of the reduced cuboid muscle model used in the forward simulations. The three-dimensional mesh represents the mechanical muscle domain, while embedded one-dimensional fibers inside the muscle encode the electrophysiological activation structure. The colors represent the stretch factor of the fiber tissue from the resting position.

Parameter	Physical Meaning	Range	N	Values
F	Prestretch traction force	0–45 N	16	0, 3, 6, 9, 12, 15, 18, 21, 24, 27, 30, 33, 36, 39, 42, 45
V	Muscle volume [16]	421.6–738.0 cm^3	5	421.6, 455.9, 573.5, 691.2, 738.0
$\text{AR} = b/a$	Aspect ratio [16] (length / cross-section)	6.8800–10.1000	4	6.8800, 7.9533, 9.0267, 10.1000
a	Cross-section side length	3.47–4.75 cm	—	
b	Muscle length	27.1–42.2 cm	—	
A_m	Membrane surface-to-volume ratio [13, 18]	450–550 cm^{-1}	3	450, 500, 550
ρ	Tissue density [12]	10.493–10.616 $\times 10^{-4}$ kg/cm^3	4	10.493, 10.534, 10.575, 10.616
Simulated physical duration per sample:				200 ms
Total combinations ($4 \times 5 \times 16 \times 3 \times 4$)*:*				**3,840 simulations**

Table A1 – Parameter sweep configuration for synthetic training data generation. Muscle dimensions are derived from volume V and aspect ratio $\text{AR} = b/a$: cross-section side length $a = (V/\text{AR})^{1/3}$ and muscle length $b = \text{AR} \cdot a$ (cuboid geometry $[a, a, b]$). Target force F is defined at a reference surface area and scaled to the actual geometry surface area for application.

CC BY